Attribute Production and Technical Change in Automobiles*

Asa Watten† Soren Anderson‡ Gloria Helfand§

May 23, 2023

PRELIMINARY: PLEASE DO NOT CITE OR DISTRIBUTE

Abstract

We develop a theoretical model of technical change and attribute production for light-duty vehicles. We apply our theory to survey data on car purchases for 1995–2017 and estimate technical change equivalent to a 1.5% annual increase in fuel economy holding drivetrain-related costs and other attributes fixed. For comparison, the EPA recently proposed a 9% annual increase in greenhouse gas emissions standards starting in 2027. We estimate that, in terms of foregone fuel economy, the opportunity cost of acceleration has tripled during our sample period, while the opportunity cost of car size has increased by half. Yet size and acceleration have both grown unabated. We credit these trends to growing consumer demand for car performance, which has overpowered rising fuel prices and biased technical change. Our estimates imply that gas taxes and fuel-economy standards today provide a stronger incentive to add fuel-saving technology and a weaker incentive to reduce size and acceleration than 20 years ago. In particular, a permanent 10% increase in gas prices leads to a 4.1% increase in fuel economy in the late 1990s and a 5.6% increase in the 2010s. Meanwhile, every 10% gain in fuel economy induced by policy leads to a 15% decrease in size and acceleration in the late 1990s but only an 8% decrease in the 2010s. We emphasize the dual role of consumer preferences and technology in driving trends in car attributes over time. By formalizing a theory of attribute production we provide tools for interpreting prior empirical estimates of attribute trade-offs and technical change in the literature.

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

*Acknowledgements: We thank Kevin Bolon and Andrew Moskalik for engineering insights, Stephanie Webber, and seminar participants at the University of Alaska Anchorage, the University of California Davis and NBER Transportation in the 21st Century 2023 Workshop 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

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

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

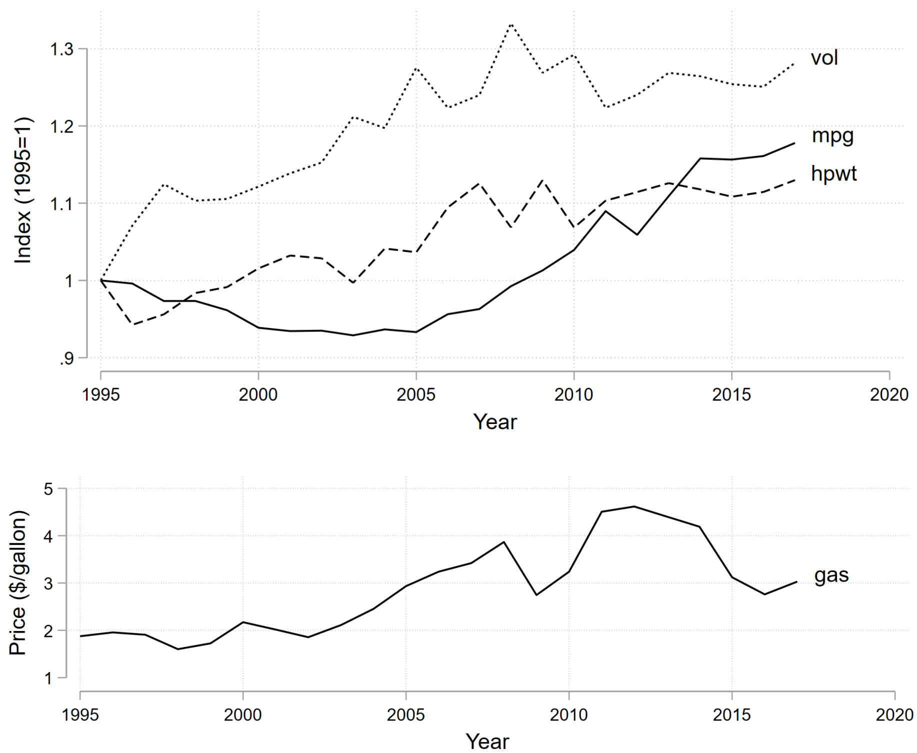

Note: Ratio changes in mean car attributes (sales-weighted) and gasoline prices compared to 1994. The horsepower-to-weight ratio (HP/weight) is a close proxy for acceleration capacity. Gas prices are core inflation adjusted.

What explains these trends? Cars have gotten bigger, faster, more efficient, and arguably more affordable, so clearly there have been large gains in technology. What has been the nature of this technical change?

Performance fell in the 1980s but not the 2000s when oil prices increased. Does this imply an altered relationship between fuel economy and performance? Car technology does not exist in a vacuum, of course. What has been the role of shifting consumer preferences for size and acceleration? How about the role of gas prices and government policy? Finally, looking to the future: What does a better understanding of these trends imply for fuel-economy standards and for the looming shift to electric cars? We address these questions using a combination of economic theory, empirical tests, and counterfactual simulation.

We begin by discussing the physical determinants of car attributes to ground our analysis in engineering considerations. We observe that the relationship between engine size, acceleration, and fuel economy has shifted dramatically over time. We then develop a theory of car attribute production that accounts for technology, preferences, and policy. On the supply side, we model a competitive car industry facing a Cobb-Douglas-like variable cost function that depends on car size, acceleration, and fuel economy. We show precisely how these attributes must enter the cost function to ensure engineering relationships that align both with engineering principles and economic intuition. On the demand side, we assume heterogeneous marginal willingness-to-pay for car attributes across consumer types, leading to an equilibrium distribution of differentiated car models. We solve for a consumer's optimal choice of fuel economy conditional on other car attributes (as in Knittel 2011), as well as optimal attributes conditional only on model primitives. We then probe the model to derive comparative statics for the equilibrium effects of higher gasoline prices, tighter fuel-economy standards, and various forms of technical change.

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 adopted uniformly or are targeted at bigger, more powerful cars as our theory predicts. Our data sources include three waves of the National Household Transportation Survey (2001, 2009, and 2017) covering cars originally purchased in 1980–2017, along with data on car attributes from the Environmental Protection Agency, and detailed data on car size from the Canadian government. Finally, we simulate counterfactual trends in size, acceleration, and fuel economy to better understand how rising gas prices, technical change, and shifting preferences for car performance have all collided to produce the trends we observe in car attributes.

Our analysis generates four key findings. First, we estimate overall technical change equivalent to a 0.7% annual increase in all car attributes for a given cost during our sample period. Yet we also find evidence for biased technical change favoring fuel economy. In particular, we estimate that a 10% reduction in acceleration leads to a 23% gain in fuel economy in the late 1990s but a 64% gain in the 2010s, holding drivetrain costs fixed. Likewise, a 10% reduction in car size leads to a 34% fuel-economy gain in the 1980s but a 52% gain in the 2010s. Because fuel economy improvements are so much cheaper in recent years, reducing other car attributes even a little bit while holding costs fixed can provide large efficiency gains. We show theoretically that, by itself, this form of biased technical change leads to gains in fuel economy and reductions in car performance.

Second, we show theoretically that this form of biased technical change makes fuel economy more sensitive to gas prices and efficiency standards, while making car size and acceleration less sensitive. Indeed, we estimate that a 10% increase in gasoline prices leads to a 4.1% gain in fuel economy in the late 1990s and a 5.6% gain in the 2010s. Meanwhile, for every 10% increase in fuel economy induced by higher gas prices or tighter standards, we estimate that size and acceleration fall by 15% in the late 1990s but just 8% in the 2010s. These results might 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 performance now yield much larger fuel-economy gains than in the past, making performance less sensitive to shifts in the demand for fuel economy.

Third, technical change in practice often manifests as a cost reduction for specific attribute-enhancing technologies, e.g. turbochargers. We show theoretically and confirm empirically that this form of technical change is inherently attribute-biased: costly technologies are first installed on cars that are bigger, faster, or more efficient overall, since the drivers of such cars by definition have the highest willingness to pay for car attributes. Such technologies only diffuse to lesser cars later, as installation costs decline. This empirical pattern holds for all of the technologies we consider, including turbo-chargers, direct fuel injection, continuously variable transmissions, and hybrid electric drivetrains. These results imply that we cannot regress fuel economy on size and acceleration for car model offerings and hope to recover the slopes of iso-cost curves, as in recent literature (Knittel 2011). Bigger, more powerful cars inevitably feature more advanced technology, implying more costly drivetrains. Controlling for these technologies in an attempt to limit the bias is also no solution, for such controls absorb the very mechanism of technical change we seek to estimate. Instead, we show how to recover costs by controlling for gasoline prices, since gas prices shift the demand for fuel economy relative to other attributes in an easily quantifiable way.

Finally, our simulations reveal that gas prices, technology, and preferences are all important for explaining trends in car attributes during 1995–2017. Rising gas prices have boosted fuel economy by 29% over this period, while keeping size and acceleration 28% lower. Technology has pushed in the same direction, boosting fuel economy by 21%, while cutting size by 15% and halving acceleration. This seemingly paradoxical result—technical change acting to constrain car performance—follows from the rising opportunity cost of size and acceleration due to biased technical change. Meanwhile, rising consumer preferences for size and acceleration have directly counteracted these trends, boosting these attributes by 29% and 53% while halving fuel economy. Overall, our simulations tell a nuanced story of surging consumer demand for size and acceleration battling rising gasoline prices and biased technical change to yield a three-decade span in which all three attributes—size, acceleration, and fuel economy—have managed to make solid gains.

A gigantic 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), who estimate shifts over time in the relationship between product cost and energy efficiency for air conditioners and water heaters, and relate this technical change to energy prices and policy. We contribute to this literature by providing theoretical micro-foundations, a new set of technical terms, and new empirical methods to better understand technical change for product attributes, including energy efficiency. Such theory is key to understanding the effects of regulation on equilibrium attributes, welfare, and incidence. Using this model, we show that failing to control for costs will lead to biased estimates of technical change. However, we show that controlling for consumer preferences can help compensate for a lack of accurate cost data. Our application to cars reveals attribute-biased technical change 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 Knittel (2011), who also uses market data to understand trade-offs among car attributes. Dozens of papers cite Knittel (2011) as the basis for their understanding of technical change and attribute trade-offs in the car industry, whether studying cars specifically (MacKenzie and Heywood 2012; Klier and Linn 2015; Whitefoot, Fowlie, and Skerlos 2017) or technical change more broadly (Newell, Jaffe, and Stavins 1999). We provide explicit micro-foundations 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, our model clarifies that different bundles of car attributes emerge from heterogeneous consumer preferences for size, power, and fuel economy. Thus, costly engine technologies are not applied at random, since their application is driven by the same consumer preferences that determine car attributes in equilibrium. Indeed, our theory predicts and our empirical results confirm that all of the engine technologies considered in Knittel (2011) are first applied to bigger, more powerful cars. These results imply that we cannot estimate iso-cost curves and technical change using market data as in Knittel (2011) without either controlling for consumer preferences or observing engine costs directly.[1] Lacking accurate cost data, we instead control for consumer preferences using microdata on car choices and state-level gasoline prices. Consistent with our theory, reduced-form correlations among car attributes and trends in fuel economy do not change drastically with these controls. But our interpretation is quite different. Iso-cost curves are

[1]Knittel (2011) proxies for costs using suggested retail prices, finding that the results are unchanged, but of course prices for cars—unlike water heaters, air conditioners, and other utilitarian appliances—reflect many unrelated add-ons (e.g., luxury trims). Knittel (2011) also controls directly for the presence of turbochargers and superchargers. He omits these controls in some models to allow his estimates of technical progress to reflect the increased penetration of these technologies, but his results change little, perhaps because these technologies never rise above 10% penetration in his sample.

steeper and the average rate of technical change is much larger than implied by the reduced form estimates. In addition, we emphasize the role of improved durability as an important mechanism of technical change, and the role of falling interest rates as a potential confounder.

Our paper also contributes to a large literature studying how fuel economy, size, weight, horsepower, and other car attributes respond to gasoline prices and fuel-economy standards. See Anderson and Sallee (2016) for a conceptual model and review. Many of these papers use reduced-form methods, implicitly capturing attribute trade-offs but failing to disentangle demand and supply. Some papers estimate discrete-choice structural models of demand, which they then complement with estimates of supply. But almost all such papers take the set of car models and their attributes, often including cost, as given. Thus, they do not consider cost and attribute trade-offs on the supply side. Exceptions include Whitefoot and Skerlos and Lin and Linn (2023) who make car attributes endogenous, at the expense of much complexity. Yet even these papers fail to consider long-run technical change. Our paper lies somewhere in-between these reduced-form and structural papers. We provide a structure that allows us to disentangle supply and demand. But we abstract away from naming specific car models and their makers, which allows us to consider both long-run attribute trade-offs and technical change. Our theory follows Ito and Sallee (2018), who study weight-based fuel-economy standards in the Japanese car market. Like us, they assume a perfectly competitive car industry with zero cost to introduce new car models, such that every consumer type gets a car model optimized to its own preferences. They then quantify short-run trade-offs between fuel economy and weight using a local quadratic approximation for welfare, which reflects both costs and preferences. We provide global functional forms for costs and preferences. In addition, we explicitly model the role of gasoline prices, interest rates, miles traveled, and depreciation. Thus, we are able to characterize the full distribution of car models and how this distribution relates to several key observables. We show that this simple model provides a rich set of predictions that align closely with empirical evidence.

The rest of this paper proceeds as follows. Section 2 discussing engineering and design trade-offs among car attributes. Section 3 develops a model of car attribute choices, showing how consumer preferences, gasoline prices, production costs, and technical change interact to drive equilibrium size, acceleration, and fuel economy. Section 4 leverages this this theory to estimate production cost parameters and rates of technical change using household-level microdata on car attribute choices. Section 2 presents several counterfactual simulations to isolate the impact of gas prices, preferences, and technical change on trends in car attributes. Section 7 concludes.

2 The 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

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.

decreasing engine size, and tuning the engine to achieve different goals. Table 1 illustrates how individual technologies and design choices affect acceleration and fuel economy. The table illustrates three categories of technologies that can be added to a vehicle: (1) those that only improve acceleration (turbo-chargers and super-chargers); (2) those that improve both fuel economy and acceleration (engine efficiency, light-weighting, aerodynamics, advanced transmissions)[2]; and (3) those that improve acceleration at the expense of fuel economy (engine displacement and tuning). The figure assumes that the size of the car is held fixed. Of course, increasing size and therefore weight lowers both fuel economy and acceleration, assuming no other changes to the drivetrain. We created this table in discussion with automotive engineers at the U.S. Environmental Protection Agency who have extensive experience testing and modeling design choices and technology options on cars. Not shown on the figure are energy-consuming technologies which may reduce fuel economy, but which have little or no effect on power, such as stereo systems, air conditioning, and computers for self-driving technologies.

Engine displacement is a particularly important design choice. Increasing engine size increases power and acceleration at the expense of fuel economy and decreasing engine size does the opposite. Adding any of these technologies, with the exception of engine tuning, add costs to manufacturing. So, to shift the attribute mix while holding cost fixed requires adding some technology and removing others or adding some performance-enhancing technology and decreasing engine size.

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 contemporary engine

[2]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.

of a given size is much more powerful than it was in 1975 while its fuel consumption is relatively unchanged. The result is that for a given horsepower, engines are both smaller and more efficient. Additionally, compared to the past, carmakers can add power with relatively small changes in engine size. The combined effect is a trend toward faster, more efficient engines that can be made faster by sacrificing relatively little fuel economy. Figure 12 (appendix) illustrates these shifts. It reproduces EPA engineering simulations of representative drivetrain technologies over time and shows that engines have improved in both 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 types of cost information to trace an iso-cost curve: (1) the cost of adding technology to marginally increase efficiency or power; and (2) the cost savings from reducing the size of the engine. 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 huge 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 have both improved vehicles and 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, gasoline prices, and policy. We use the model to derive optimal attribute choices in equilibrium comparative statics with respect to gasoline prices and technical change. Using the model, we show how to infer cost parameters and rates of technical change from observed market data on car attribute choices. We also demonstrate that discrete engine technologies are adopted on bigger, faster, more fuel-efficient cars, implying that unobserved costs are correlated with car attributes.

3.1 Modeling optimal attribute choice

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

3.1.1 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 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 their continuous choice of car attributes and spending on other goods. Utility is given by: $u(s, a, g) = v(s, a) + y - pgm - r(s, a, g)$, where $v(s, a)$ is utility derived from size and acceleration, y is income, pgm is lifetime expenditures on fuel, and $r(s, a, g)$ is the equilibrium 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 lifetime miles traveled (m) is exogenous, which is consistent with empirical evidence that demand for miles is inelastic; relaxing this assumption would not substantially change our conclusions. However, miles traveled (m) 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: $c(s, a, g)$, which depends only on size, acceleration, and fuel economy. We focus on these three attributes, since they are bound together in fundamental engineering relationships. Bigger cars are heavier, which slows them down. Larger engines can compensate but use more fuel. Hybrid engines and other technologies can save fuel but cost money. We assume that carmakers are price takers.[3] We further assume that there is free entry and zero fixed cost to develop and produce a new car model. Taken together, these assumptions imply that, in equilibrium, every consumer type chooses a car model custom-tailored to its own unique preferences, with equilibrium prices equal to costs: $r(s, a, g) = c(s, a, g)$. Note that the assumption of zero fixed costs will hold approximately, so long as the number of consumers of each type is sufficiently large.

Finally, we model a fuel-economy standard σ, which constrains average gallons-per-mile across each carmakers's fleet. We assume that this standard is implemented via a perfectly competitive credit-trading system with zero transaction costs, leading to equilibrium credit price τ. Thus, the fuel-economy standard can be modeled as an implicit tax on each car's gallons-per-mile relative to the standard: $\tau[g - \sigma]$. Note that this tax becomes an implicit subsidy if the car is more efficient than the standard: $g < \sigma$.[4] 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.[5]

[3] The presence of market power and markups should change the analysis little so 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.

[4] 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.

[5] Note that future iterations of this paper will consider the possibility, following recent policy changes, that the fuel-economy

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

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

which is utility from size and acceleration plus spending on other goods. 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:

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

where k is a cost shifter (the residual in our empirical model below), 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.[6] 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 respect to each of these attributes ($c_{ss}, c_{aa}, c_{gg} > 0$). Thus, it is increasingly costly to make a car bigger, faster, or more efficient. Finally, the cross-partials imply that making a car more efficient is costlier when the car is big or fast ($c_{ag}, c_{sg} < 0$) and likewise that making a car faster is costlier when the car is big ($c_{as} > 0$). To summarize succinctly: making a car better in any dimension is more costly when the car is already better in any dimension.

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:

$$\max_{s,a,g} y + \tau\sigma - \beta_s s^{-1} - \beta_a a^{-1} - \beta_g g - ke^{-\theta} s^{\alpha_s} a^{\alpha_a} g^{-\alpha_g}, \tag{3}$$

where preference parameters $\beta_s, \beta_a, \beta_g > 0$ capture the marginal benefit from improved car attributes. This formulation implies that utility is nonlinear in car size, which implies a strong desire to avoid tiny cars but a diminishing marginal benefit for ever-larger cars, and likewise for acceleration.[7] Our derivations above imply that the marginal benefit from a reduction in gallons-per-mile relates precisely to fuel prices, miles, and policy: $\beta_g \equiv pm + \tau$. However, our formulation here based on a generic β_g would permit various behavioral extensions, such as undervaluation of fuel economy due to myopia, credit constraints, or other factors.

Before moving to our necessary conditions, we pause briefly to explore the cost function in equation (2). Define the **marginal rate of technical substitution of attributes (MRTSA)** as the slope of the iso-cost

standard is an increasing function of car size. 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.

[6] 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.

[7] 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. the second-order conditions for utility maximization are satisfied.

curve in two-dimensional attribute space. The MRTSA measures how much one attribute must change to increase another attribute, while holding production costs constant. The MRTSA is simply the (negative) marginal cost ratio for a given pair of attributes:

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

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

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

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 holding costs fixed. Likewise, $MRTSA_{ga}$ in the second row tells us how much gallons-per-mile (a bad) increases with a marginal increase in acceleration holding costs fixed. Conversely, these values tell us how much fuel economy (a good) decreases with a marginal increase in size or acceleration.

Meanwhile, the elasticity versions are given by:

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

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

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

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, holding costs fixed, and likewise for the other two elasticities. The elasticity formulation is convenient, because the relationships among these attributes are equivalent when gallons-per-mile is replaced with its inverse (miles-per-gallon, a good). Signs flip but nothing else changes.

3.1.3 Necessary conditions

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

$$\text{s:}\ \beta_s = \frac{\alpha_s c(s,a,g)}{s^{-1}} \tag{10}$$

$$\text{a:}\ \beta_a = \frac{\alpha_a c(s,a,g)}{a^{-1}} \tag{11}$$

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

where we express size and acceleration as bads (s^{-1}, a^{-1}) in the denominators of (10) and (11), to make clear their symmetry with gallons-per-mile. These conditions say that the marginal benefit from improving a car attribute (left side) should equal the marginal cost (right side). Second-order sufficient conditions are shown in the appendix.

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

$$-\frac{\beta_s}{\beta_g} = -\frac{c_s(s,a,g)}{c_g(s,a,g)} = -\frac{\alpha_s}{\alpha_g}\frac{g}{s^{-1}} \tag{13}$$

$$-\frac{\beta_a}{\beta_g} = -\frac{c_a(s,a,g)}{c_g(s,a,g)} = -\frac{\alpha_a}{\alpha_g}\frac{g}{a^{-1}} \tag{14}$$

$$-\frac{\beta_a}{\beta_s} = -\frac{c_a(s,a,g)}{c_s(s,a,g)} = -\frac{\alpha_s}{\alpha_a}\frac{a^{-1}}{s^{-1}}, \tag{15}$$

where we continue to express size and acceleration as bads (s^{-1}, a^{-1}). 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 (13) as points of tangency between level sets of utility (indifference curves) and level sets of production costs (iso-cost curves) in fuel economy and size space.

3.1.4 Optimal attribute choices

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

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

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

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

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

To illustrate, suppose that the preference parameters on 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 parameter on gallons-per-mile $(\ln\beta_g)$. Inspection of equations (16)–(18) reveals

Figure 2: *Optimal attribute choices for two consumers*

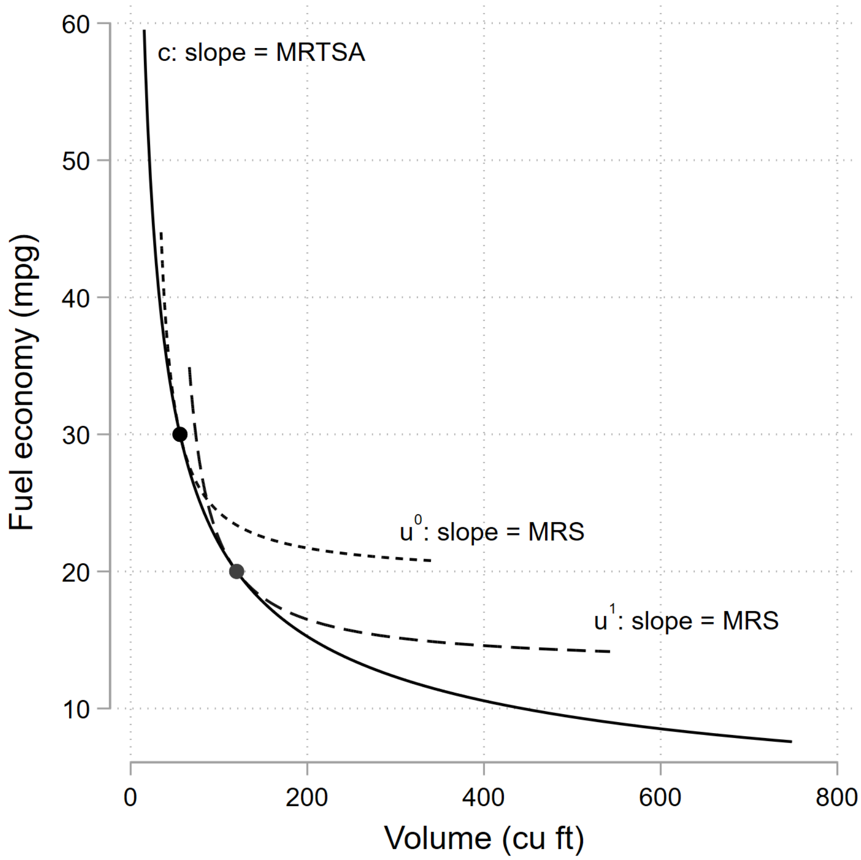

Note: This figure illustrates optimal attribute choices in two-dimensional (fuel-economy and size) space as points of tangency between consumers' indifference curves (u's) with slopes equal to the MRS, and the iso-cost curve in production (c) with slope equal to the MRTSA.

that consumers with stronger preferences to save fuel (larger $\ln \beta_g$) will choose smaller, less powerful, and more efficient cars. Further, because equations (16)–(18) are linear, any pair of logged attributes will be perfectly correlated across car models. In practice, we do not see a perfect correlation of car attributes across car models, 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 analysis below, and to connect to Knittel (2011) and other papers that regress log fuel economy on car attributes. Substitute for the cost function in equation (12) using the specific functional form in (2), solve for gallons-per-mile (g), and then take logs to yield:

$$\ln g = \frac{1}{1+\alpha_g}\ln k - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\ln \beta_g + \frac{\alpha_s}{1+\alpha_g}\ln s + \frac{\alpha_a}{1+\alpha_g}\ln a \tag{19}$$

$$\approx \frac{1}{1+\alpha_g}\ln k - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g}\ln pm + \frac{\alpha_s}{1+\alpha_g}\ln s + \frac{\alpha_a}{1+\alpha_g}\ln a, \tag{20}$$

where the approximation in the second row follows from $\beta_g = pm + \tau$ and the assumption that $\tau \approx 0$. This approximation is consistent with back-of-the-envelope calculations, which suggest that pm is at least an order of magnitude larger than τ.[8] We return to this issue in our empirical analysis below, where we show that our results are robust to controlling for the relatively small variation in β_g due to τ. This equation suggests that we can recover cost parameters $(\alpha_s, \alpha_a, \alpha_g)$ and the annual rate of attribute-neutral technical change $(\partial\theta/\partial t)$ from a regression of log fuel economy on log size, log acceleration, and a time trend. However, we must control for fuel-economy preferences ($\ln \beta_g \approx \ln pm$). Why? Even holding size and acceleration fixed, this equation implies that consumers with a strong preference to save fuel will choose more fuel-efficient and therefore more costly cars. Meanwhile, equations (16)–(17) show that consumers with strong preference to save fuel will choose smaller, less powerful cars. Thus, the fuel-economy preferences term $(\ln \beta_g/(1+\alpha_g))$ will be correlated with the other car attributes. Failing to control for this term will remand it to the error term, leading to biased coefficient estimates.[9]

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

[8] 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 non-compliance penalty under U.S. fuel economy standards is \$55 per mile-per-gallon per car, which translates to $5 \cdot 55 = \$225$ for a 5 mile-per-gallon improvement. Note that the non-compliance penalty is an upper bound on marginal compliance costs. Indeed, Anderson and Sallee (2011) show that marginal costs are under \$30 in many years. These calculations imply that pm is at least 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.

[9] Ignoring this term will also make it impossible to identify β_g, which is necessary to infer the rate of technical change the other cost parameters.

3.2.1 Knittel's (2011) regressions are likely biased

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

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

which is the form of Knittel's (2011) baseline regressions: log fuel economy on log attributes and year dummies to capture technical progress (θ).[10] Thus, our functional form for costs provides explicit micro-foundations for Knittel (2011) and related papers.

Note that the coefficient on log size (α_s/α_g) captures the slope of the iso-cost curve in two-dimensional product space: percent changes in g that hold costs fixed given a percent change in s. Likewise for the coefficient on log acceleration ($-\alpha_a/\alpha_g$). 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 20), putting them on a higher iso-cost curve. Meanwhile, these same consumers tend to choose different attributes (equations 16 and 17), which implies different marginal costs for fuel economy (equation 2). Thus, costly fuel-saving technologies will not be added to cars at random. We show this point theoretically below in our analysis of discrete technology adoption, and we confirm it empirically through an examination of penetration rates for various engine technologies.

Knittel (2011) concedes (pg. 3372) that this regression may not yield unbiased estimates for the slopes of iso-cost curves at a point in time or literal shifts in iso-cost curves over time due to technical change. But he argues that his estimates remain valid for predicting how fuel economy would have evolved in the presence of technical change had car attributes remain fixed. In effect, he argues that he is estimating the reduced-form time trend ($\theta/(1+\alpha_g)$) in our equation (20). 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 iso-cost 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 equilibrium responses to fuel economy forcing mecha-

[10] 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 signs on the coefficients appropriately flipped.

nisms (gas prices and standards) is to trade-off much less acceleration and size in equation (26) by adding fuel-saving technology. The iso-cost 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 iso-cost curves as proxies for equilibrium response to standards is increasing in the estimated iso-cost 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 consumer preferences.

3.3 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 (16)–(18) with respect to log fuel-economy preferences ($\ln \beta_g$) yields:

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

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

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

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

Note that the percentage changes in size and acceleration are given by α_g/ψ, which reflects the relative steepness of the cost function in the direction of gallons-per-mile. The costlier it is to improve fuel economy relative to other attributes, the more size and acceleration decrease.[11] Meanwhile, the percentage change in gallons-per-mile is given by $(1+\alpha_s+\alpha_a)/\psi$, which is inversely related to the steepness of the cost function in the direction of miles-per-gallon. Thus, the costlier it is to improve fuel economy relative to other attributes, the smaller the improvement.

These changes are proportional to the baseline attribute mix. This implies that consumers who value fuel savings more will make larger improvements relative to those who value it less, while those who value acceleration and size more will decrease these attributes more. This creates a spreading out of fleet attributes along the fuel economy axis and a contraction on the acceleration and size axes.

From these attribute changes—equations (22)–(24)—we can find the ratio of equilibrium changes to attributes, or the equilibrium attribute trade-off elasticity. This has an interesting relationship to the iso-cost

[11]Note that the percentage change is 0 when $\alpha_g = 0$ and approaches 1 as α_g grows to infinity.

elasticities—equations (8)–(7). Taking ratios of the pairs, we get the change in one attribute relative to the change in another resulting from a change in gasoline prices

$$\frac{\frac{\partial \ln g}{\partial \ln \beta_g}}{\frac{\partial \ln a}{\partial \ln \beta_g}} = \frac{\frac{\partial \ln g}{\partial \ln \beta_g}}{\frac{\partial \ln s}{\partial \ln \beta_g}} = \frac{1+\alpha_s+\alpha_a}{\alpha_g} \tag{25}$$

$$= \frac{1}{\alpha_g} + \text{elast.}_{g,a} + \text{elast.}_{g,s}. \tag{26}$$

Recall that these elasticities are positive and utility is decreasing in g. This relationship implies that the equilibrium effect of a change in gas prices on fuel consumption relative to other attributes is larger than the iso-cost elasticity. Intuitively, when the price of gas goes up, consumers add fuel-saving technology, moving to a higher iso-cost curve, in addition to trading-off some other attribute. As the iso-cost elasticity gets large, the equilibrium deviation from the iso-cost curve gets larger—consumers achieve their new fuel consumption level increasingly by adding technology rather than trading off other attributes. For example, if both cost elasticities are 0.5 (or for a 1% increase in fuel economy, there is a 2% decrease in acceleration or size) then the equilibrium consumer elasticity is more than 1 (for a 1% increase in fuel economy there is a 1% decrease in acceleration or size).

Note that increasing the stringency of the fuel-economy standard (τ) has qualitatively similar effects. To see this, observe that $\partial \ln \beta_g/\partial\tau \approx 1/pm$ when evaluated at $\tau \approx 0$. Thus, just re-scale each of the effects in equations (22)–(24) by $1/pm$ to yield the marginal percentage effects. 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 differences in lifetime miles across consumers.

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

$$\frac{\partial \ln g}{\partial \ln \beta_g} = -\frac{1}{1+\alpha_g} < 0, \tag{27}$$

which shows that fuel economy increases conditional on size and acceleration. Thus, consumers optimally choose cars that are smaller and slower overall (equations 22 and 23). 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 (27) 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.

3.4 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.

3.4.1 Attribute-neutral technical change

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

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

$$\frac{\partial \ln g}{\partial \theta} = -\frac{1}{\psi}, \tag{28}$$

and likewise for the other attributes (with a sign-flip). Thus, unbiased technical change leads to proportional improvements in size, acceleration, and fuel economy. These gains are inversely proportional to the sum of the exponent parameters in the cost function ($\psi = 1 + \alpha_s + \alpha_a + \alpha_g$). These parameters (the α's) equal 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. Since attribute changes are proportional to baseline attributes, unbiased technical change spreads out new attribute levels in all attribute dimensions; consumers use technical change to buy more of the attributes that they value most. We can see this spreading-out effect clearly in figure (8) in the next section.

Now consider how a marginal increase in parameter θ affects fuel economy conditional on other attributes in equation (20). Differentiating with respect to θ yields:

$$\frac{\partial \ln g}{\partial \theta} = -\frac{1}{1+\alpha_g} < 0. \tag{29}$$

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 cost parameter on fuel economy (α_g). The larger this parameter, the smaller the shift. Knittel (2011) implicitly assumes that this parameter is zero—that costs do not depend on fuel economy given other car attributes—such that the rate of technical change can be inferred directly from a regression of log fuel economy on a time trend while controlling for other car attributes. We show empirically that this assumption is violated ($\alpha_g > 0$). 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.

3.4.2 Attribute-biased technical change

Now consider attribute-biased technical change. We model this form of technical change using a linear add-on to our original cost function:

$$\tilde{c}(s, a, g) = c(s, a, g) + \mu(g - g_0), \tag{30}$$

where $\mu = 0$ implies our original cost function (c) and $\mu > 0$ 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:

$$\tilde{c}_g = -\frac{\alpha_g c}{g} + \mu = -\frac{\alpha_g c - \mu g}{g}, \tag{31}$$

which implies the following MRTSA between fuel economy and acceleration:

$$-\frac{\tilde{c}_a}{\tilde{c}_g} = -\frac{\alpha_a}{\alpha_g - \frac{\mu g}{c}} \frac{g}{a} < -\frac{\alpha_a}{\alpha_g} \frac{g}{a}, \tag{32}$$

where the last inequality assumes $\mu > 0$ with $\alpha_g > \mu g/c$. 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 (μ) 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 attributes 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 ($\mu = 0$). In the presence of biased technical change ($\mu > 0$), 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. Equation (26) shows that a steeper iso-cost elasticity reduces the equilibrium trade-off response to a change in gas prices. So, fuel taxes, standards, and price shocks will have a much lower impact on size and acceleration.

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.

3.4.3 Discrete technology adoption

Carmakers periodically invent new technologies that may be added to cars to improve one or more attributes. For example, turbochargers temporarily increase the performance of an internal combustion by forcing additional compressed air into the combustion chamber. Adding a turbocharger to a car's drivetrain allows it to accelerate faster. Carmakers often add turbochargers and simultaneously decrease engine displacement ("downsizing"), thereby improving fuel economy while holding acceleration constant (Shahed and

Figure 3: *Biased technical change*

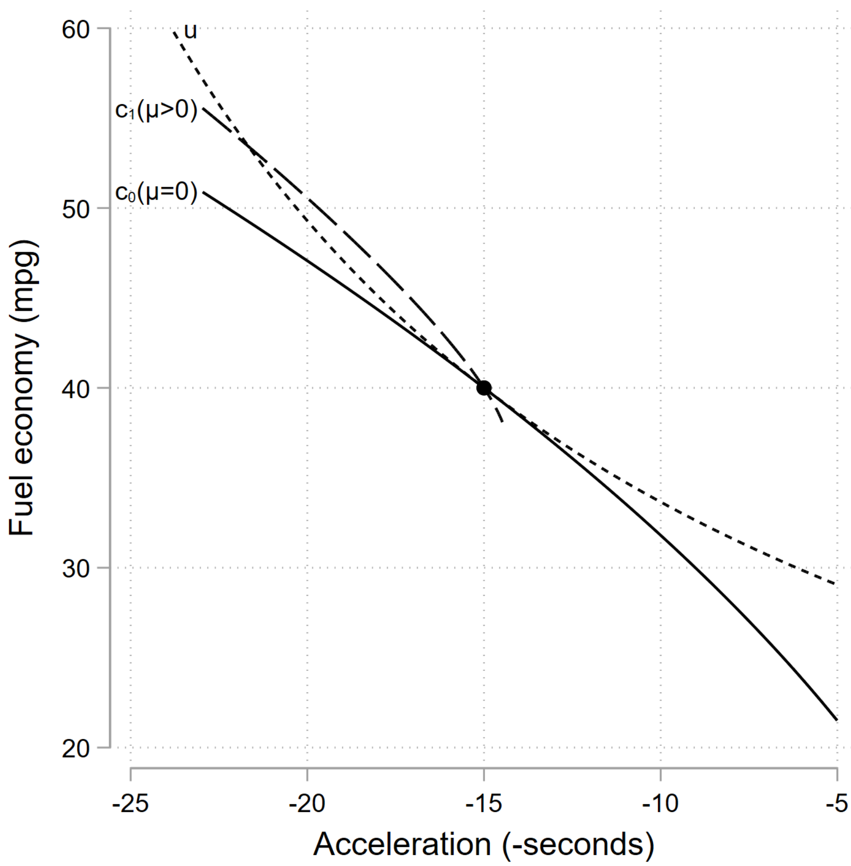

Note: This figure illustrates the effects of biased technical change in two-dimensional (fuel-economy and acceleration) space. The figure shows an optimal attribute bundle as the point of tangency between a consumer's indifference curve (u) and an iso-cost curve prior to technical change ($c(\mu = 0)$).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

Bauer 2009). Over time, the cost of adding a turbocharger has fallen—a form of technical change—leading to increased adoption. Other examples include continuously variable transmission, direct injection (replacing carburetor), and electric or hybrid-electric 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 a specific car attribute—acceleration, to use the turbocharger example—that consumers may add to their car at a cost. This technology has two important features. First, 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 the targeted attribute. Thus, cars with a higher baseline level of the targeted attribute will see a larger absolute improvement.

We consider a technology that boosts acceleration as perceived by the consumer to $a/(1-\omega) > a$, where $\omega > 0$. Conditional on adopting the technology, and ignoring installation cost, the consumer now experiences utility given by:

$$u(s, a, g; \omega) = y + \tau\sigma - \beta_s s^{-1} - \beta_a (1-\omega) a^{-1} - \beta_g g - c(s, a, g), \tag{33}$$

where we have scaled inverse acceleration by $1 - \omega < 1$, implying a fractional reduction of ω in this bad attribute. By the envelope theorem, the gain in utility for a marginal increase in ω at the optimum is given by: $\partial u(s, a, g; \omega)/\partial\omega = \beta_a a^{-1}$. Thus, the net benefit from adopting the discrete technology is given by the following approximation:

$$u(s, a, g; \omega > 0) - u(s, a, g; \omega = 0) - \kappa \approx \beta_a \omega a^{-1} - \kappa, \tag{34}$$

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

It can be shown that adoption benefits ($\beta_a \omega a^{-1} - \kappa$) are increasing in consumer preferences for size (β_s) and fuel economy (β_g), as well as acceleration (β_a). Intuitively, because car attributes are in technical competition, consumers with strong preferences for size and fuel economy get stuck in slow cars. It's exactly these consumers who would benefit most from a boost in acceleration—especially if they also crave speed. Note, there is nothing special about acceleration in our analysis—these same qualitative results hold for technologies targeting size and fuel economy. Thus, it is generally the consumers who want more of everything who are most likely to adopt an attribute-enhancing technology. Only as costs (κ) gradually fall over time will others adopt. This result conforms with the stylized fact that new car technologies typically first make their appearance in luxury segments, and only gradually filter down to the less driven drivers.

A key implication of these results is that technology is not uniformly adopted in the fleet. Instead, attributes that are not directly affected by the technology predict adoption. This result has implications for estimation. To estimate technical trade-offs by fitting a line through attribute data as in Knittel (2011), we must assume that unobserved costs are uncorrelated with car attributes. However, we have shown that discrete technologies are more likely to be adopted by consumers with extreme preferences, who choose bigger, faster, more fuel-efficient cars. Thus, unobserved costs are inevitably correlated with attributes. We confirm this prediction empirically in the following section.

What happens to equilibrium attributes for consumers that do adopt? Consider equations (16)–(18) and note that adopting the acceleration-boosting technology is equivalent to replacing acceleration preferences

$(\ln \beta_a)$ with modified preferences $(\ln(1-\omega)\beta_a)$ in each of these equations. Differentiating each of these equations with respect to ω and evaluating at $\omega = 0$ then yields:

$$\left.\frac{\partial \ln s}{\partial \omega}\right|_{\omega=0} = \frac{\alpha_a}{\psi} > 0 \tag{35}$$

$$\left.\frac{\partial \ln(1-\omega)^{-1}a}{\partial \omega}\right|_{\omega=0} = \frac{\alpha_a}{\psi} > 0 \tag{36}$$

$$\left.\frac{\partial \ln g}{\partial \omega}\right|_{\omega=0} = -\frac{\alpha_a}{\psi} < 0, \tag{37}$$

where note that the relevant attribute in the second row is log acceleration as perceived by the consumer. Strikingly, consumers that adopt the discrete technology experience improvements in all attributes, i.e. even those not targeted by the technology. Even more strikingly, the proportional improvement is the same across all three attributes. Thus, the impact of the technology is the same as for attribute-neutral technical change—but only among the adopters, who are selected on size, acceleration, and fuel economy.

These same results all obtain when we instead assume that that the discrete technology proportionally reduces the cost of supplying the targeted attribute: $c(s, a(1-\omega), g) = (1-\omega)^{\alpha_a} c(s, a, g)$, 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 facing high technology costs—the ones choosing big, fast, and fuel-efficient cars. Further, since adoption is equivalent to attribute-neutral technical change, it will lead to equal proportional improvements in size, acceleration, and fuel economy among the adopters—same as a technology affecting the perceived benefit of the targeted attribute.

3.5 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 β's) can be interpreted as the present-discounted values:

$$\beta_s = \int_{t=0}^{\infty} e^{-rt} b_s(t), \;\; \beta_a = \int_{t=0}^{\infty} e^{-rt} b_a(t), \;\text{ and }\; \beta_g = \int_{t=0}^{\infty} e^{-rt} b_g(t), \tag{38}$$

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:

$$\beta_g = p \int_{t=0}^{\infty} e^{-rt} m(t) dt + \tau, \tag{39}$$

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:

$$\beta_g = \frac{pm(0)}{r + \rho + \delta}, \tag{40}$$

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 (20) for the optimal choice of fuel economy conditional on other attributes becomes:

$$\ln g \approx \frac{k}{1+\alpha_g} - \frac{\theta}{1+\alpha_g} - \frac{1}{1+\alpha_g} \ln pm + \frac{1}{1+\alpha_g} \ln(r + \rho + \delta) + \frac{\alpha_s}{1+\alpha_g} \ln s + \frac{\alpha_a}{1+\alpha_g} \ln a, \tag{41}$$

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 (38) 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.

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.

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 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.

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 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 EPA prior to matching to NHTS data.

Our fuel economy and acceleration data come from the U.S. Environmental Protection Agency (EPA). We measure fuel-economy as the combined city-highway measure reported by EPA. The EPA measures gallons-per-mile in a simulator designed to mimic both city and highway driving. EPA calculates a weighted 55% city and 45% highway average of miles per gallon.[12] 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).[13] 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.[14]

We match our car choice data to state-level retail gasoline prices. These data are based on Davis and Kilian (2011) who start with pre-tax retail prices reported by EIA and then painstakingly add percentage-based (ad-valorem) and constant per-gallon (specific) gasoline taxes from myriad sources. Bates and Kim (2022) updated 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 (2011) 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, 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. So, 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

[12] The EPA revised the measure for car labeling requirements at dealerships starting in 22012, 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.

[13] 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) shows that the relationship between of horsepower-to-weight and 0–60 time changed little from 1995 to 2010 (the end of their study period).

[14] 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.

NHTSA (2006), along with our own estimates for the annual decay in miles for new vs. used cars based on the 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, 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 8 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, holding cost fixed. But our theory shows that the slope of these relationships also depend on the cost of adding fuel-saving technology. Second, cars have gotten bigger, faster, and less efficient over time (comparing the red and black histograms). However, cars of a given size and speed have become more efficient over time (comparing the red and black scatter diagrams). These shifts are consistent with growing consumer preferences for size and acceleration, along with improved car technology that has allowed even biggest and fastest cars to become more efficient—interpretations emphasized by Knittel (2011). But again, 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

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

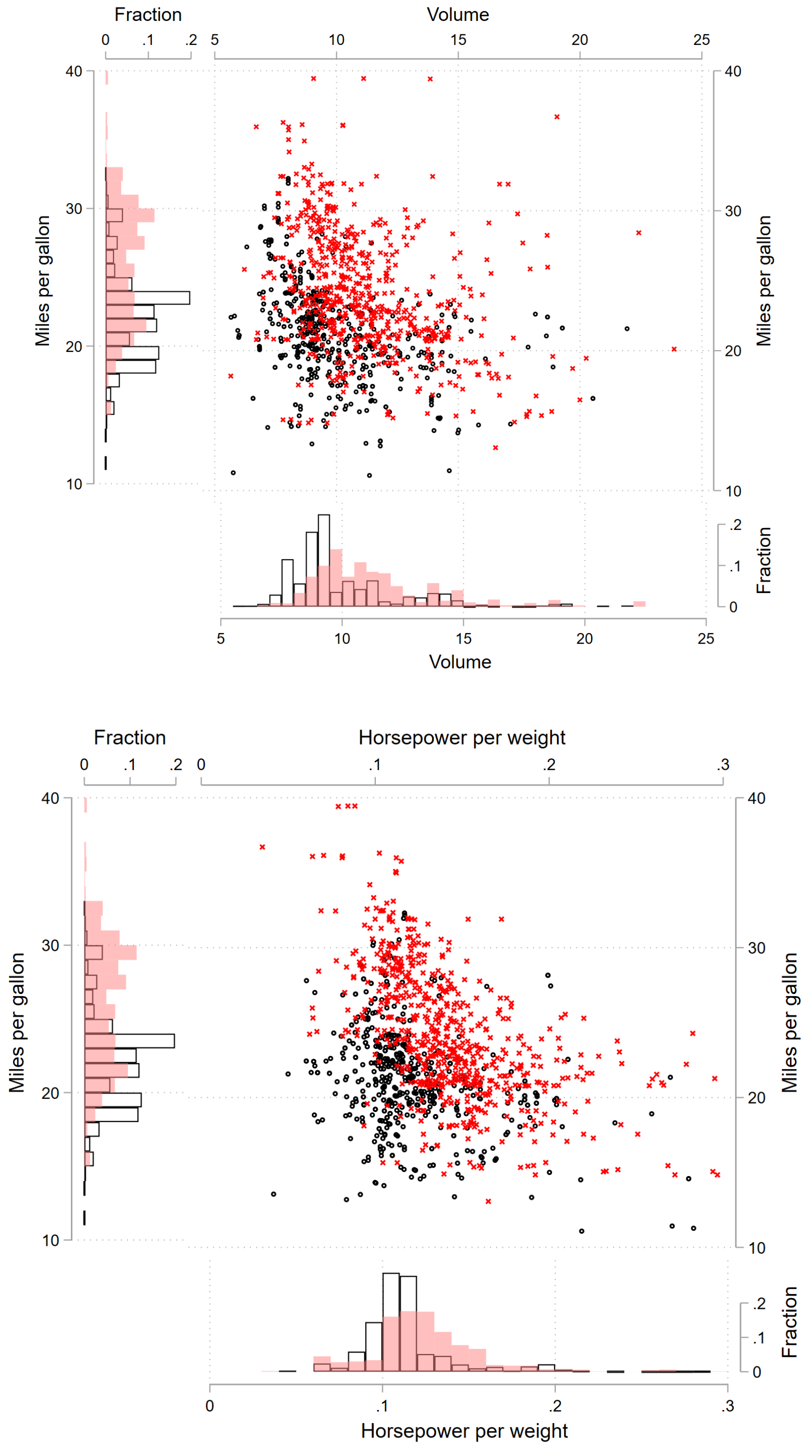

Note: This figure plots the distribution of new car fuel economy (miles per gallon), volume (cubic meters based on $l \times w \times h$ of car), and acceleration (as proxied by the horsepower per weight-in-pounds ratio). Attributes for model years 1999–2001 are shown in black while attributes for 2015–2017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

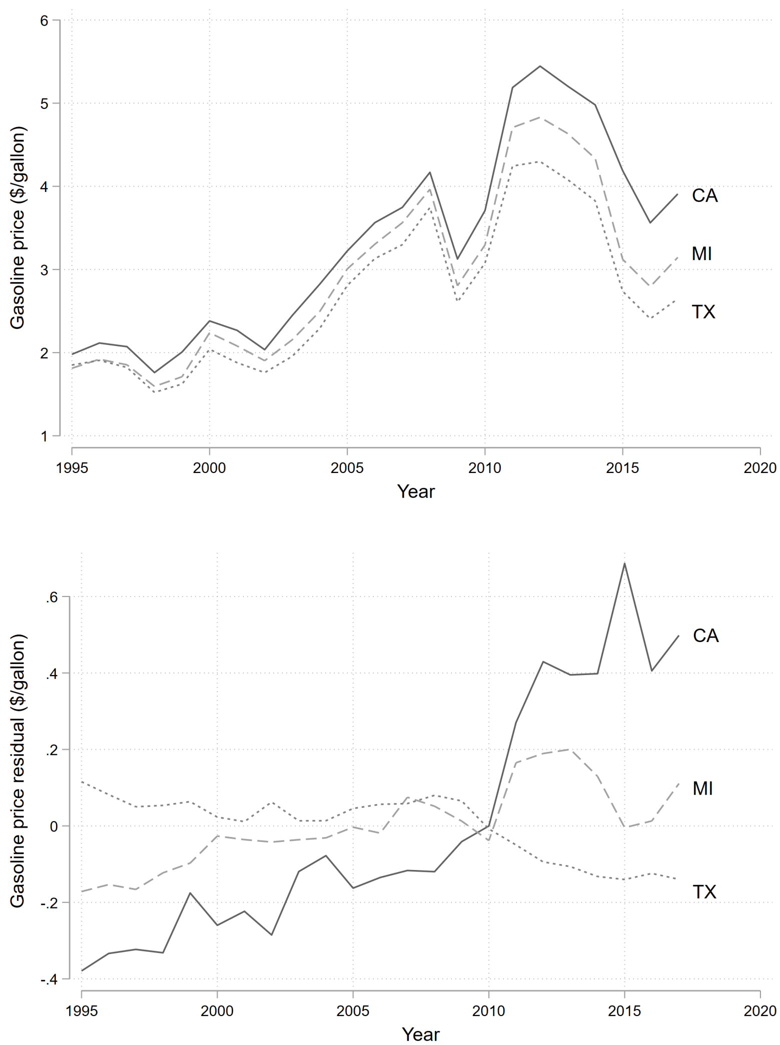

Note: The top panel shows real ,monthly retail gasoline prices in California, Michigan, and Texas from 1995–2017. The bottom panel shows the residual variation in gasoline prices for these three states, following an OLS regression of gasoline prices on state and month fixed effects for all 50 states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

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 driven by global crude oil prices, with persistent gaps across states. However, the bottom panel shows

Figure 6: *Trends in interest rates on new-car loans*

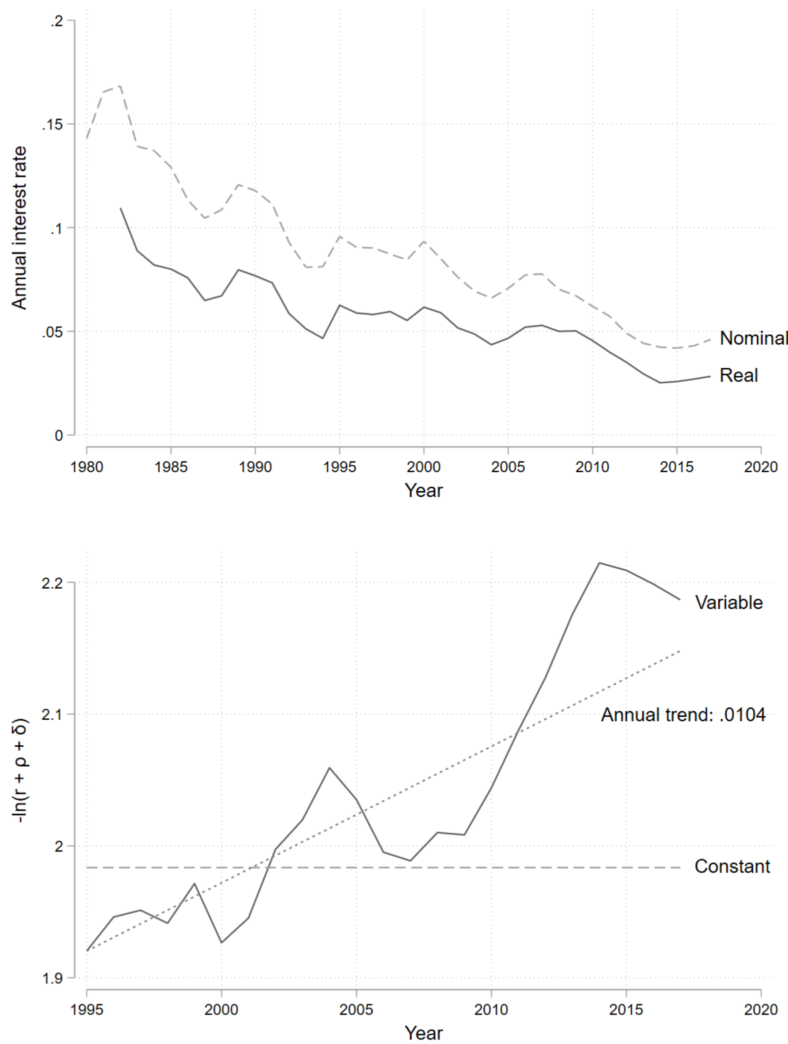

Note: The top panel plots the trend in nominal and real interest rates on 48-month new auto loans. The bottom panel plots the trend of the present value multiplier, which is a function of the real interest rate.
Data source: Federal Reserve Bank of Saint Louis.

that there is substantial variation in gasoline prices, even after removing state and month fixed effects, with California's price rising relative to Michigan and Texas over time. This residual variation is 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 6 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 have fallen by a factor of 2 over our sample period. Meanwhile, the bottom panel shows the trend in (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 has risen 24% over our sample period due to 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 NHTSA estimates that scrap rates (ρ) have fallen over time, we find no evidence that the decay in miles (δ) has fallen across the three NHTS waves, nor do we find that first-year miles ($m(0)$) has changed substantially. Thus, we control only for time-series variation driven by interest rates.

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 (41) above:

$$\ln mpg_{ist} = f_\theta(t) + \gamma_p \ln p_{ist} + \gamma_s \ln volume_{ist} + \gamma_a \ln hpwt_{ist} + \phi_s + \varepsilon_{ist}, \tag{42}$$

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; $\ln p_{ist}$ is log gas prices in state s and year t; $f_\theta(t)$ captures attribute-neutral technical change; ε_{ist} is an error term; and the γ's 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, # household members, # adults, # vehicles, home ownership, urban vs. rural, and population density.

We estimate this equation via 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 state and year dummies. There are two main concerns. First, unobserved shifts in a state's demand for fuel economy over time might be positively correlated 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 as well as miles traveled interacted with car age; but note we only observe miles in survey years 2001, 2009, and 2017.[15] 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 iso-cost 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 iso-cost curves (MRTSA values). The coefficient on log gas prices is also key to recovering trends in the index of attribute-neutral technical change, whether we capture this index via a linear time trend or year dummies ($\Delta\theta \equiv \Delta f_\theta(t)/\gamma_p$).[16] We do not observe the precise month in which a car is purchased. Thus, we include two years of lagged gasoline prices, 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.

Our theory implies that the coefficient on log gas prices $\times$ 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. 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, to capture cross-sectional and temporal variation in demand for fuel economy, analogous to our demographic controls.

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

[16] 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.

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 figure 8, 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 7 plots the coefficients on the year dummies and shows large increases in fuel economy conditional on attributes (gray triangles with short-dashed line). Following Knittel, column (2) then replaces year dummies with a linear time trend. Consistent with the dummy year variables, we estimate a 1.7% 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 annual gasoline 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 non-zero coefficient on gas prices implies that these correlations are not literally the slopes of iso-cost curves. To recover these slopes, we must divide by $1 - \gamma_g \approx 0.9$. Thus, iso-cost curves are approximately 1.1 times steeper than the equilibrium relationships shown in figure 8, as fuel-saving technology is disproportionately added to bigger 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.

Figure 7 again plots the coefficients on the year dummies. Model (3) controls for gas prices. The year coefficients from this model (gray diamonds with 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.[17] Model (4) additionally controls for car age and log miles, along with their squares and interactions. The year coefficients from this model (gray squares with solid line) are slightly higher than those from model (3). Finally, model (5) additionally controls for demographics.

[17]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 7.

Table 2: *Fuel economy conditional on attributes*

	(1)	(2)	(3)	(4)	(5)
Annual trend (γ_θ)		0.017***			
		(0.001)			
Log gas price t			0.019	0.016	0.016
			(0.019)	(0.018)	(0.017)
Log gas price t-1			0.029*	0.027*	0.026*
			(0.011)	(0.011)	(0.011)
Log gas price t-2			0.047*	0.046*	0.046*
			(0.023)	(0.020)	(0.021)
Log HP/weight (γ_a)	-0.509***	-0.517***	-0.509***	-0.509***	-0.507***
	(0.017)	(0.016)	(0.017)	(0.017)	(0.017)
Log volume (γ_s)	-0.469***	-0.480***	-0.469***	-0.469***	-0.467***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Pickup	-0.068***	-0.064***	-0.068***	-0.067***	-0.065***
	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)
Year FEs	X		X	X	X
State FEs	X	X	X	X	X
VMT x Age				X	X
Demographics					X
Sum gas price (γ_g)			0.096	0.089	0.088
			0.041	0.038	0.038
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}}\vert_{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g}\right)$			-0.563	-0.559	-0.556
			0.043	0.041	0.041
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}}\vert_{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g}\right)$			-0.519	-0.515	-0.512
			0.030	0.028	0.028

Note: This table presents coefficient estimates from 42. Standard errors in parentheses are clustered by state. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively. Data source: NHTS, EPA, and Transport Canada.

The year coefficients from this model (black circles with solid line) are notably higher, crossing over the coefficients from model (1) but with a shallower slope.

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 model (5), we estimate $r_\theta = 0.15$ over our 22-year sample period. This rate applies to theoretical drivetrain costs ($c(s, a, g)$), 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

Figure 7: *Estimated coefficients on year dummies*

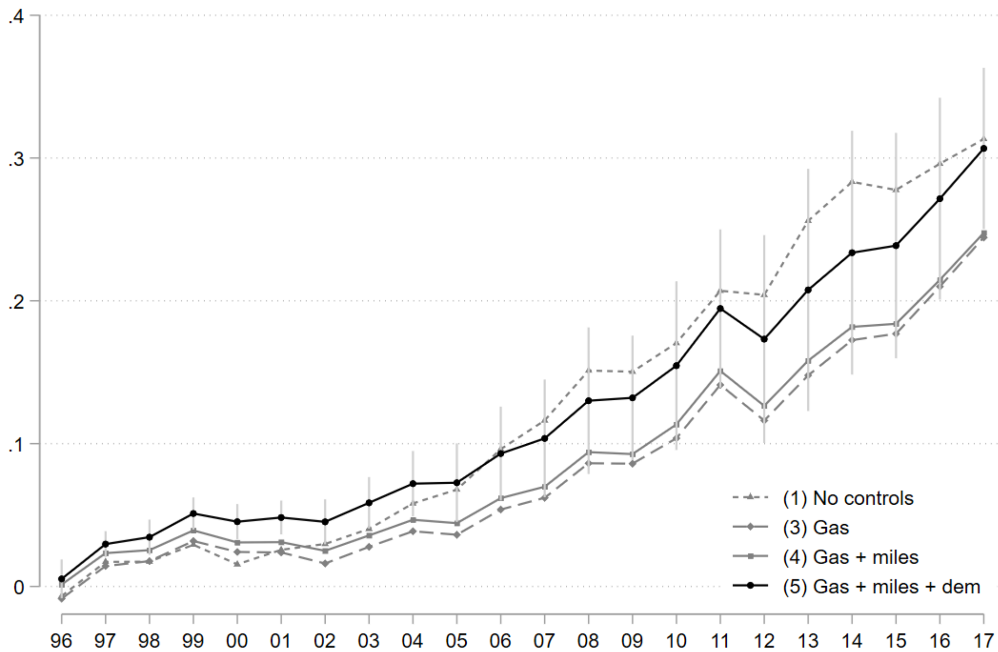

Note: This figure plots estimated coefficients on the year dummies from regressions (1), (3), (4), and (5) from table 2. Vertical bars are the 95% confidence intervals from regression (5) based on standard errors that are clustered by state.

to equation (41). Based on this approach, we estimate annual technical change in attributes of 0.7%.[18] 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%.

5.1.2 Consumer preferences

Our theory shows in equations (13)-(15) that consumers optimize by equating attribute trade-offs in utility (MRS) to slopes of iso-cost curves (MRTSA). Thus, we use these theoretical conditions, along with the estimated cost coefficients reported in table 2 above, to infer each household's log MWTP for size in terms of forgone fuel economy ($\ln(\beta_s/\beta_g)$), and likewise for acceleration ($\ln(\beta_a/\beta_g)$).[19] Figure 8 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 respondents (histograms). The figure shows wide

[18]Equation (28) 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 get 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.

[19]Given an estimate of for the present-discounted value of lifetime fuel expenditures, we could estimate WTP for size and acceleration in dollars.

Figure 8: *Consumer preferences*

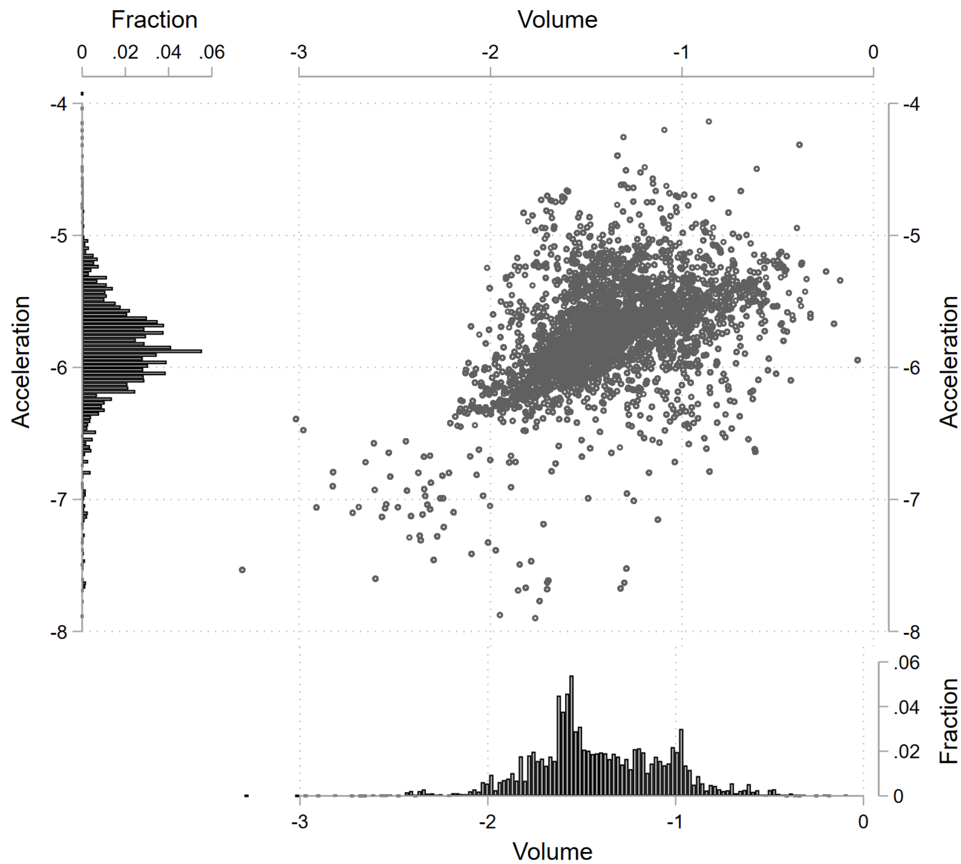

Note: This figure plots estimated preference parameters for acceleration ($\ln \beta_a/\beta_g$) and size ($\ln \beta_s/\beta_g$) for every car purchase in our final estimation sample, covering NHTS waves 2001, 2009, and 2017. The scatter diagram in the center plots preferences for size vs. acceleration, 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 (13)–(15) and do not depend on gas prices (which vary) or lifetime miles (which we do not observe). We calculate $MRS = MRTSA$ using on 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.

dispersion in preferences, where a one log-point change in implies $\exp(1) \approx 2.72$ times higher MWTP for size and acceleration relative to fuel economy. The figure also shows a strong positive correlation between preferences for size and power. Appendix figure 11 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 bigger 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

Table 3: *Biased technical change*

	(1) 1995-01	(2) 2002-09	(3) 2010-17
Log gas price t	0.006	0.006	0.006
	(0.013)	(0.013)	(0.013)
Log gas price t-1	0.023	0.023	0.023
	(0.013)	(0.013)	(0.013)
Log gas price t-2	0.047*	0.047*	0.047*
	(0.021)	(0.021)	(0.021)
Log HP/weight (γ_a)	-0.212***	-0.527***	-0.594***
	(0.015)	(0.016)	(0.012)
Log volume (γ_s)	-0.341***	-0.461***	-0.517***
	(0.016)	(0.006)	(0.009)
Pickup	-0.060***	-0.060***	-0.060***
	(0.009)	(0.009)	(0.009)
Sum log gas price (γ_g)	0.076*	0.076*	0.076*
	(0.037)	(0.037)	(0.037)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{HP/Wt}}\vert_{\Delta c=0} \left(\frac{\alpha_a}{\alpha_g} = \frac{\gamma_a}{1-\gamma_g}\right)$	-0.229***	-0.571***	-0.642***
	(0.022)	(0.040)	(0.039)
$\frac{\partial \ln \text{mpg}}{\partial \ln \text{volume}}\vert_{\Delta c=0} \left(\frac{\alpha_s}{\alpha_g} = \frac{\gamma_s}{1-\gamma_g}\right)$	-0.369***	-0.499***	-0.560***
	(0.027)	(0.024)	(0.032)

Note: This table presents coefficient estimates from equation 42 in which we interact acceleration and size 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. Note *, **, and *** indicate statistical significance at the 5%, 1%, and 0.1% levels respectively.
Data source: NHTS, EPA, Transport Canada.

have the highest willingness to pay.

5.1.3 Biased technical change

Have the slopes of the iso-cost curves shifted over time? Table 3 shows OLS results from equation (42) in which we allow the coefficients on acceleration and size to differ across three different eras: 1995–2001, 2002–2009, and 2010–2017. This regression is based on model (5) from table

The results show increasingly steep iso-cost curves over time in both fuel-economy vs. acceleration space and fuel economy vs. size space. These results are consistent with technical change that is substantially biased toward fuel economy.

Figure 9: *Discrete technology adoption*

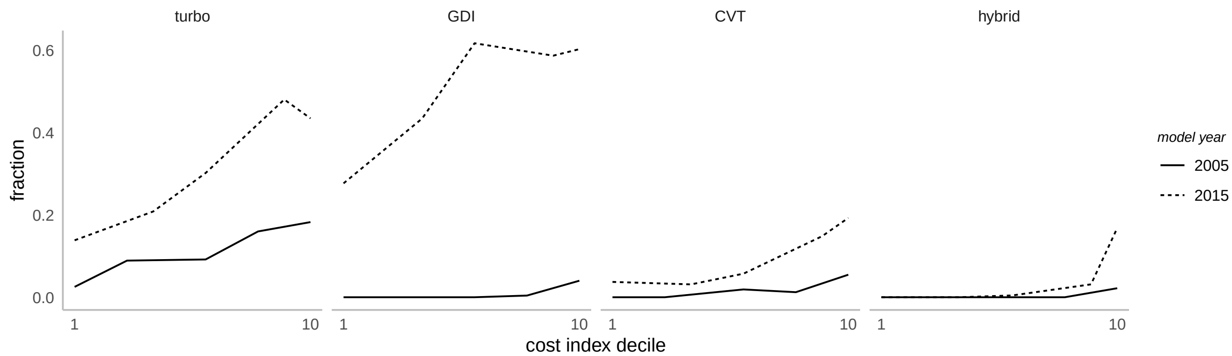

Note: sales-weighted fractions of discrete technology adoption for cost index decile for 2005 model years and 2015 model years. Deciles are computed within each year and the estimated cost index uses coefficients from our main specification.

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

Figure 9 the fraction of adoption for deciles of an estimated drivetrain cost index. The estimated drivetrain cost index is computed using our main specification and is a weighted sum of vehicles attributes:

$$cost\ ind. = \hat{\alpha}_s \ln vol + \hat{\alpha}_a \ln hpwt + \hat{\alpha}_g \ln mpg. \tag{43}$$

We compute the index at the trim-level using Wards Automotive data, which contains 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 and electric motors. We show these for two years, 2005 and 2015 to show that that is both greater adoption and the pattern of adoption at an early and later year. The figure shows that both for earlier and later years, costlier vehicles—those with greater 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:

$$technology_{it} = \phi_{bt} + \gamma_s \ln volume_{jt} + \gamma_a \ln hpwt_{jt} + \gamma_g \ln mpg_{jt} + \varepsilon_{jt}, \tag{44}$$

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 micro-data as we would loose trim-level detail and instead use trim-level sales weights from EPA. Vehicle body styles here are serving to capture larger consumer groups that have distinct preferences for acceleration and fuel economy.

The OLS results are given in table 4. While acceleration and fuel economy are physically affected by all of these technologies, volume is not which is highly correlated with all technologies within a vehicle body style.

Table 4: *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 (43) 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, sedands/waggons, hatchbacks, coupes/convertables.
+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

These results further support our hypothesis that drivers who want large vehicles will adopt technologies at a higher rate holding other preferences fixed. Additionally, the results support our concern that the correlation of attributes alone cannot estimate iso-cost curves. Inestead, 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, interest rates, 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 10 plots the year dummies ($\theta/(1+\alpha_g)$) and trend in MRTSA values (α_a/α_g and α_s/α_g) from this regression (upper-left panel). Note that sum of the gas price coefficients is $\gamma_g = 0.064$, which implies $\alpha_g = 14.6$.

Calibrating preferences. Given our estimated cost parameters, we calculate the preference parameters associated with each car choice in our sample according to the $MRTAS = MRS$ conditions in equations (13)–(14). 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:

$$\ln(\beta_g) = \ln p + \ln m(0) + \ln(r_t + \rho + \delta), \tag{45}$$

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 six-year moving average of log gas prices, which yields preferences for fuel economy, size, and acceleration that all evolve gradually over time.[20] We calibrate $\rho = 0.05$ based on data in NHTSA (2006).[21] We estimate (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 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 has shifted over time. Thus, first-year miles $m(0)$ only varies 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 10 plots the resulting trends in preferences for fuel economy, size, and acceleration, all relative to their 1995 levels.

Calibrating the cost residual. Our final step is to calibrate the cost residual (k) based on equation 20, given our above calibrations for costs (the α's and θ) and preferences (the β's). 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 observed baseline fuel economy. Meanwhile, since our counterfactual simulations hold k and $m(0)$ fixed, they play zero further role. Figure 10 replicates the baseline trends in car attributes as generated by our simulation model (middle-left panel).

[20] Recall that we infer preferences for size and acceleration based on their values relative to fuel economy. Thus, basing fuel-economy preferences on just one year of gas prices would result in wildly fluctuating preferences for size and acceleration, which seems unreasonable. Recall that our econometric model uses the current year's gas price plus two lags.

[21] Annual scrap rates conditional on survival change over time as a car ages: lower than 0.05 when a car is young, higher than 0.05 during middle age, and then low again in the golden years. These patterns presumably reflect a complex mix of heterogeneity across cars and drivers, e.g. beloved old cars that are barely driven. These dynamics are not central to our analysis. Thus, we pick $\delta = 0.05$, which roughly corresponds to the weighted-weighted average scrap rate, with 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.

6.2 Counterfactual simulations

We perform four counterfactual simulations that respectively hold gas prices ($\ln p$), interest rates (r), technology (θ & α's), 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 (16)–(18), which give optimal attribute choices as a function of model primitives. Finally, we calculate mean attributes in every year. Figure 10 plots the resulting trends in fuel economy (middle-right panel), 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, while interest rates play a minor role. When we hold gas prices fixed at their 1995 levels ("no gas"), fuel economy in 2017 is 29% lower than baseline, while size and acceleration are 28% higher. These effects are totally unsurprising.

When we hold car technology constant ("no tech"), fuel economy is 21% lower, while size is 15% higher and acceleration is more than twice as high. These results illustrate the multi-dimensional nature of technical change that we estimate in our model. Unbiased technical change ($\Delta\theta$) leads to proportional increases in all attributes. Meanwhile, biased technical change favoring fuel economy, as reflected in a rotating MRTSA (upward trends in α_a/α_g and α_s/α_g), leads to an increase in fuel economy and a decrease in size and acceleration. Intuitively, the opportunity cost of performance is increasing over time, leading 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.

When we hold consumer preferences for size and acceleration constant ("no prefs"), fuel economy almost doubles, while size decreases by 29% and acceleration decreases by 53%. These results illustrate the tremendous growth in consumer preferences over time, especially for acceleration, along with the trade-off that consumers are willing to make: lower fuel economy for better performance.

Finally, when we hold interest rates constant ("no i-rate"), car attributes barely change.

7 Conclusion and discussion

Why has fleet average fuel economy increased with acceleration and size 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 towards fuel economy. This has made the opportunity cost of forgone fuel economy very large while the desire for size and acceleration are also large. These two trends have worked together such that consumers are buying more of all attributes instead of trading-off

Figure 10: *Trends in technology, preferences, baseline car attributes, and counterfactual changes*

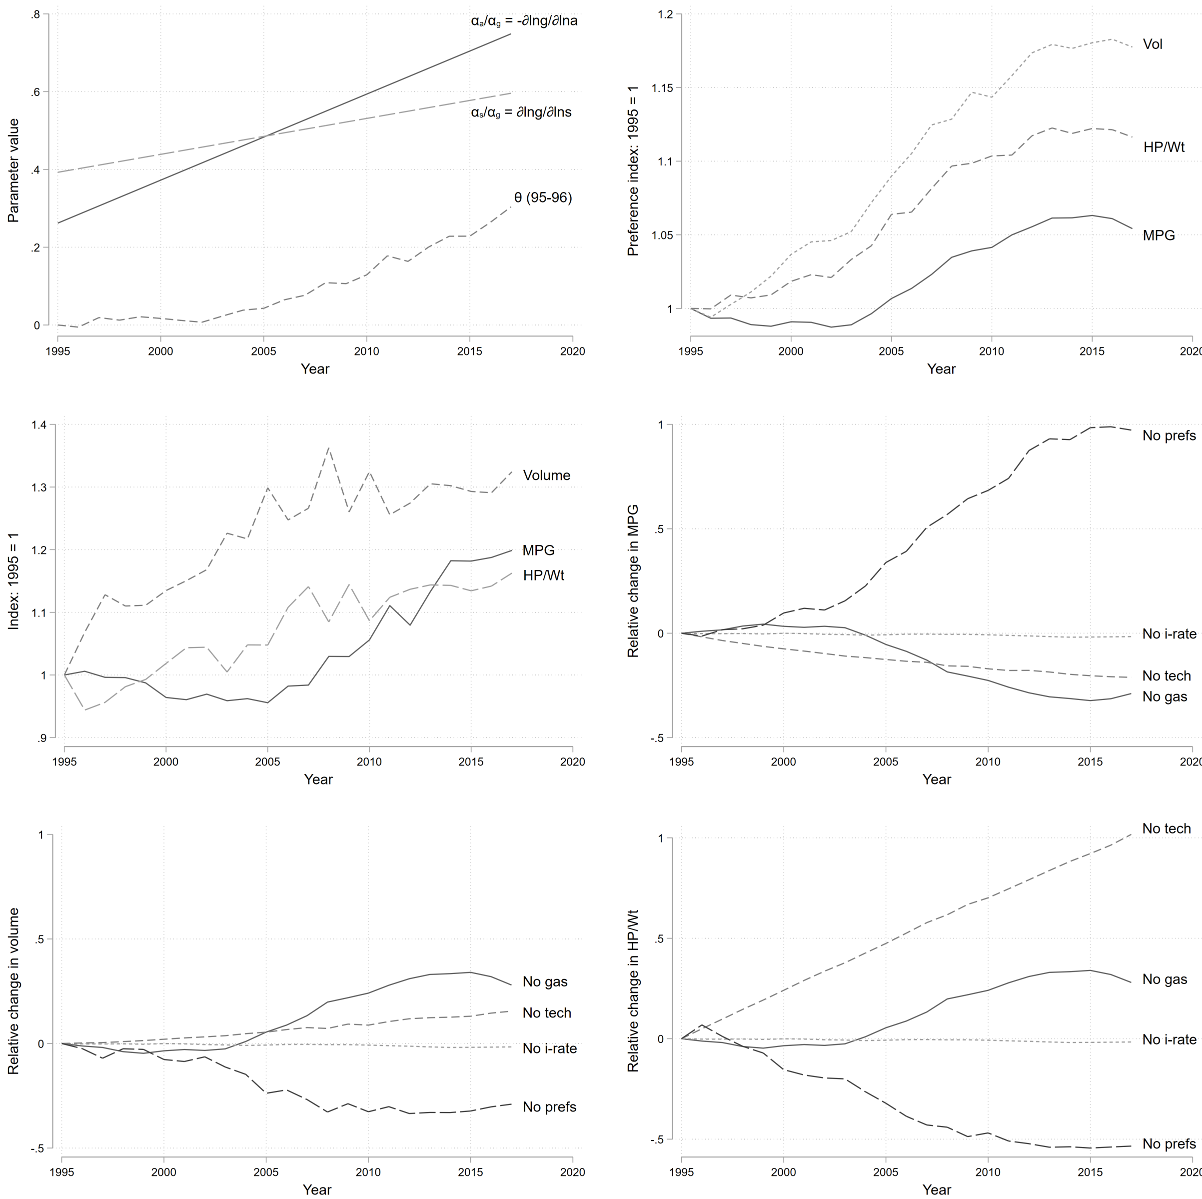

Note: This figure presents the parameter inputs to our counterfactual simulations, along with our results. The top-left panel shows the trend in attribute-neutral technical change (θ) and MRTSA values (α_a/α_g & α_s/α_g). The top-right panel shows the trends in average consumer preferences for attributes. The middle-left panel shows trends in baseline attributes relative to their 1995 values as generated by our simulation model (same as actual trends in figure 1). The remaining panels show changes relative to baseline 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: gasoline prices (labeled ”no gas”), interest rates (“no i-rate”), car technology (“no tech”), and consumer preferences (“no prefs”). See text for details.

attributes. At the same time, a driver with identical preferences responds to increases in gas prices, gas taxes, and standards by buying more fuel economy than they had in the past and reducing acceleration and size less.

7.1 Implications for EVs

To include electric vehicles (EVs) in our model, we would need to include range in the cost. Range is a function of battery size and inertia weight. So, increasing range interacts with the marginal cost of other attributes. Charging characteristics (speed of charging) on the other hand do not meaningfully interact with the marginal cost of other attributes.

While we do not analyze EVs, our model has some qualitative predictions. Electric drivetrains are a technology that both improves cost per mile and acceleration. As technology improves and costs come down, EVs make these attributes cheaper. To the extent that EVs reduce the cost of improving efficiency more than they reduce the cost of improving acceleration relative to today, EVs will reduce the equilibrium response trade-offs to standards and fueling costs. That is, as standards become more stringent, drivers will sacrifice even less acceleration and possibly size than they had in the past.

Our comparative statics tell us that EVs will first be adopted by consumers with higher preferences for lowering driving costs, by consumers with high preferences for acceleration, and, to a lesser extent, those who prefer large cars. Electric vehicles appear to be replacing highly efficient gasoline cars and sports cars. This is important for considering marginal changes in standards. As standards increase and encourage greater adoption of EVs, the gas car that the EV replaces will be a larger, less efficient one. This implies that, to the extent that EV incentives force technological improvement, there is the knock-on effect of increasing social benefit per EV. This is a departure from the standard model that assumes decreasing marginal benefits in the level of incentives. Studies that measure the marginal benefit of EVs (Holland, Mansur, Muller, and Yates 2019) would need a very different interpretation.

We plan to test these implications in future work.

7.2 Implications for regulatory analysis

EPA analysis aims to hold all attributes constant and only consider the cost of adding sufficient fuel or energy-saving technology to achieve the standards. We show that not only have improvements in fuel economy become cheaper relative to other attributes, but the equilibrium response to standards is much closer to EPA analysis than in the past.

Our model is relatively simple and transparent. This makes it a good candidate to test complex regulatory analyses such as EPA's OMEGA model. EPA already produces a cost-surface for producing attributes using their ALPHA model. These data could easily be deployed using our theoretical model to estimate the cost parameters and predict the impact of regulation on equilibrium attribute production including distributional effects.

References

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. 2011. “Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards.” *American Economic Review* 101 (4):1375–1409.

———. 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.”

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/.

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, 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.

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. ???? "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

$$u_{ss} = -2\frac{\beta_s}{s^3} - \alpha_s(\alpha_s - 1)\frac{c}{s^2} \tag{46}$$
$$u_{aa} = -\alpha_a(\alpha_a + 1)\frac{c}{a^2} \tag{47}$$
$$u_{gg} = -\alpha_g(\alpha_g + 1)\frac{c}{g^2} \tag{48}$$
$$u_{sa} = \alpha_s\alpha_a\frac{c}{sa} \tag{49}$$
$$u_{sg} = \alpha_s\alpha_g\frac{c}{sg} \tag{50}$$
$$u_{ag} = -\alpha_a\alpha_g\frac{c}{ag}. \tag{51}$$

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

$$u_{ss} = -\alpha_s(\alpha_s + 1)\frac{c}{s^2} \tag{52}$$

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

$$D_2 = u_{ss}u_{gg} - u_{sg}^2 \tag{53}$$
$$= \alpha_s\alpha_g\frac{c^2}{s^2g^2}\Big((\alpha_s + 1)(\alpha_g + 1) - \alpha_s\alpha_g\Big) > 0. \tag{54}$$

The determinant of the third leading principle minor is

$$D_3 = u_{ss}u_{aa}u_{gg} - u_{ss}u_{ag}^2 + 2u_{sa}u_{sg}u_{ag} - u_{sa}^2u_{gg} \tag{55}$$
$$= -\alpha_s\alpha_a\alpha_g\frac{c^3}{s^2a^2g^2}\Big(\alpha_s + \alpha_a + \alpha_g + 2\alpha_s\alpha_a\alpha_g + 2\alpha_a\alpha_g\Big) < 0. \tag{56}$$

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

B Tables and figures

B.1 Engineering simulation

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

Figure 11: *Correlates of consumer preferences*

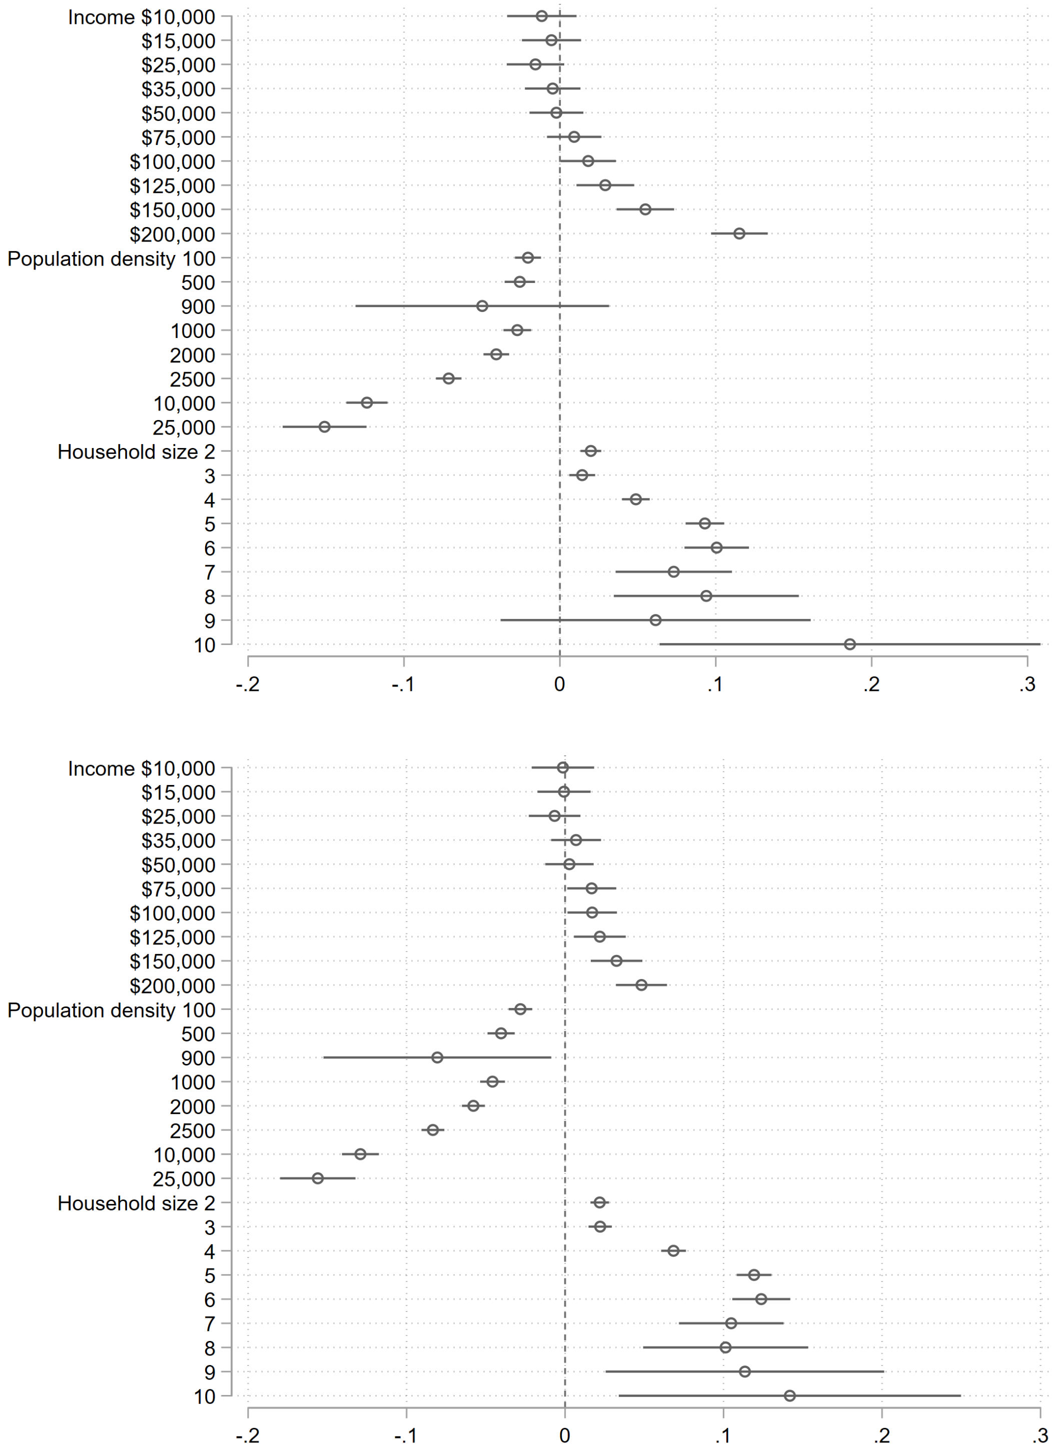

Note: This figure plots the estimated coefficients (hollow circles) and 95% confidence intervals (horizontal bars) from an OLS regression of estimated logged consumer preference ratios ($\ln \beta_a/\beta_g$ and $\ln \beta_s/\beta_g$) on dummy variables for household income, population density, and # 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 12: *Simulated fuel economy vs. acceleration (via engine displacement)*

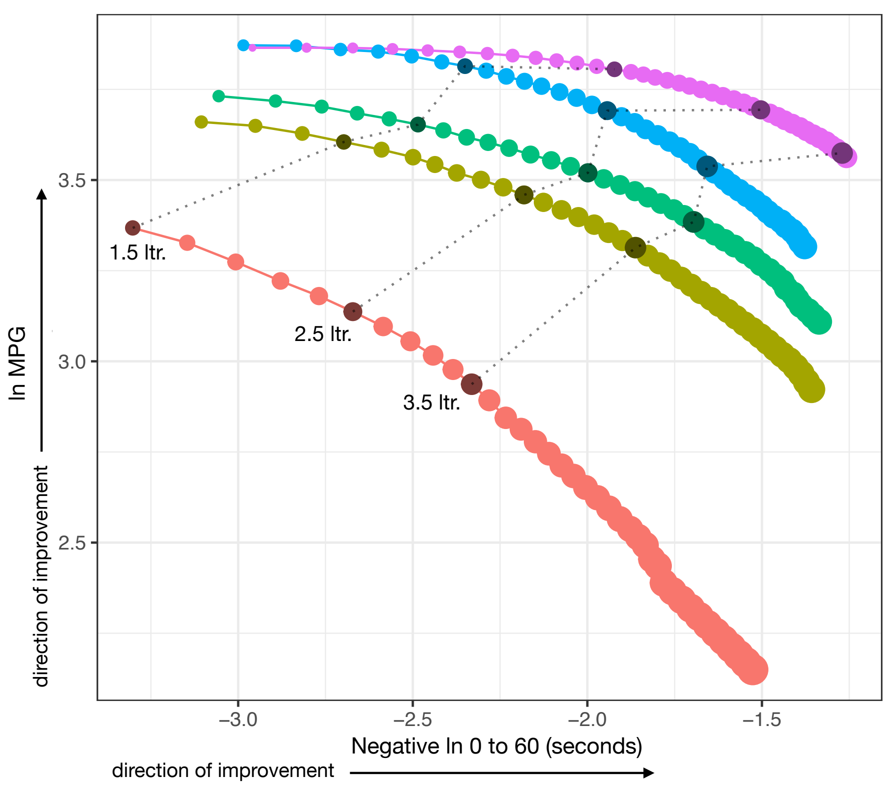

Note: This figure shows simulated log fuel economy vs. log acceleration for five distinct drivetrains, with different combinations of fuel economy and acceleration achieved by changing engine displacement. From bottom to top the five engines and transmissions are: 1980s carbureted (3-speed), 2007 Toyota PFI (5-speed), 2013 GM GDI (6-speed), 2017 Honda turbo (8-speed), future Ricardo 24 bar EGR (8-speed).

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 12, 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.[22] The five power-trains include: (1) a 1980s carbureted engine with a three-speed automatic transmission; (2) a 2007 Toyota port fuel-injected (PFI) engine coupled to a five-speed transmission; (3) a 2013 GM gasoline direct injection (GDI) engine coupled to a six-speed transmission; (4) a 2017 Honda turbo-charged engine coupled 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

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

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 big 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.