

Designing Policies to Make Cars Greener: A Review of the Literature

Soren T. Anderson¹ and James M. Sallee²

¹Michigan State University, East Lansing, MI, USA; email: sta@msu.edu

²University of California, Berkeley, Berkeley, CA, USA; email: sallee@berkeley.edu

Xxxx. Xxx. Xxx. Yyyy. AA:1–27

This article's doi:
10.1146/((please add article doi))

Copyright © Yyyy by Annual Reviews.
All rights reserved

Keywords

fuel economy regulation, fuel taxation, automobiles, greenhouse gas emissions

Abstract

We review what is known about the economic efficiency of fuel taxes relative to efficiency standards aimed at mitigating environmental externalities from automobiles. We present a simplified model of car choice that allows us to emphasize the relationships between fuel economy, other car attributes, and miles traveled. We focus on greenhouse gas emissions, although we note how other environmental externalities affect our conclusions. Our main conclusion—that standards are substantially less efficient than a fuel tax—is already familiar. Less familiar are points we make about the relative importance of the rebound effect, on the effects of attribute-based policies, and the implications of behavioral biases. We point to areas where we believe future research can have the greatest contribution, including work on uncertainty, heterogeneity, and empirical work in low and middle-income countries.

Contents

1. Introduction	2
2. Theoretical framework	4
2.1. Consumer welfare	4
2.2. Externalities and internalities	5
2.3. Policy options	6
2.4. Consumer behavior	6
2.5. First-best corrective policies	7
2.6. Limitations of our model	7
3. How do alternative policies affect behavior?	7
3.1. The effects of a fuel tax	8
3.2. The effects of a traditional fuel-economy standard	10
3.3. Is a feebate the same as a fuel-economy standard?	13
3.4. What are the effects of an attribute-based fuel-economy standard?	14
4. Additional considerations	17
4.1. Behavioral biases do little to justify fuel-economy standards	17
4.2. Heterogeneity in vehicle longevity and on-road fuel economy strengthen the case for a fuel tax	18
4.3. The implications of imperfect competition for energy policy are not well-understood	18
4.4. Cheating erodes the efficiency of standards	19
4.5. Safety-related externalities are large	20
4.6. Accounting for other un-priced externalities reinforces the efficiency advantage of a fuel tax	20
4.7. Knowledge spillovers may justify complementary policies	21
4.8. Barriers limiting the adoption of biofuels may justify complementary policies	21
5. Conclusion	22
A. Alternative model with fuel economy given as a function of car attributes and technology	26
A.1. Consumer welfare	26
A.2. Externalities and internalities	26
A.3. Policy options	27
A.4. Consumer behavior	27

1. Introduction

Few policy challenges loom larger than climate change. Given the high stakes, policies designed to mitigate carbon emissions should be both effective and efficient. Economists generally advocate the use of carbon pricing for achieving efficient mitigation. Yet command-and-control regulations proliferate. Can such regulations mimic the incentives of carbon pricing? If not, what are the costs of inefficient regulations?

In this review, we attempt to answer these questions for the transportation sector. Transportation is an important piece of the climate change puzzle, accounting for about 13% of global greenhouse gas emissions. Moreover, there is a pronounced relationship between vehicle ownership and wealth both over time and across countries, suggesting that this sector will become more important as low and middle-income countries grow. China—the prime example—boasted 22 million cars sold in 2013 and has been the world’s largest car market since 2009. Are car fleets and carbon emissions destined to grow in lock-step with economic growth?

Even absent an economy-wide carbon tax, an efficient Pigouvian tax for transportation emissions is available in the form of a tax on fuel. Are existing fuel taxes already inducing

the efficient level of mitigation? While fuel taxes are high in many countries, typical policies remain far from the chalkboard ideal. The United States has the lowest fuel taxes among industrialized countries (Parry & Small 2005), and many oil-producing countries actually subsidize the price of fuel, presumably for political reasons (Parry, Veung & Heine 2014a; Davis 2014). Meanwhile, the European Union explicitly exempted the transportation sector from its carbon pricing scheme. Moreover, despite decades of research pointing to the superiority of fuel taxes, fuel-economy standards predominate as the main environmental policy in the car market. Some countries, such as the United States and Japan, have long-standing regulatory regimes, but similar policies have recently been adopted by the European Union, China, India, and others.

Our aim is to understand how these apparently entrenched regulations compare to a fuel tax in terms of the incentives they create and the behaviors they induce. We first consider a traditional fuel-economy standard, which requires manufacturers to sell vehicles that have a minimum average fuel economy. We compare this standard to a feebate, which explicitly taxes or subsidizes vehicles according to their fuel economy, as well as to attribute-based standards that provide a sliding fuel-economy target based on a car's size or weight.

We focus on economic efficiency, though we note where equity concerns are most likely to alter our findings. In addition, we focus on carbon emissions, but we discuss how our conclusions are affected by considering local pollution, congestion, and safety externalities—in brief, our conclusions are reinforced. We concentrate on the literature related to the United States, but we make attempts to point to recent work in other countries. The relative efficiency of fuel-economy standards versus a fuel tax is well-trod ground and addressed in several recent reviews (see Anderson, Parry, Sallee & Fischer 2011, Proost & Van Dender 2011, and Gillingham 2013). Thus, we do not attempt to provide a comprehensive review of this literature. Instead, we focus on the most recent findings and issues and, with an eye to future research, we attempt to identify the most important unanswered questions.

Our main conclusions can be summarized as follows:

1. No regulatory alternative achieves the efficiency of a fuel tax.
2. Theory and empirics around the rebound effect have often been mismatched. It is critical to understand the interaction between vehicle quality and miles traveled in determining which empirical estimates are relevant for a particular policy analysis.
3. Biases in consumer choice are unlikely to be strong enough to overturn the traditional wisdom that a fuel tax is preferred to standards on efficiency grounds.
4. In practice, attribute-based standards function as technology-forcing policies by limiting other compliance options, thereby raising mitigation costs.
5. Attribute-based standards are poorly suited for addressing behavioral biases, since they preserve existing tradeoffs between fuel economy and other car attributes.

In our view, despite the substantial volume of research on these topics, a number of important questions remain, including the following:

1. The world's largest car markets all feature attribute-based standards, but this topic has received little attention in the academic literature.
2. A priority for future research should be to work on rapidly growing car markets, where both the car market and government regulation are inchoate.
3. As compared to traditional standards, both attribute-based standards and feebates have the potential to stabilize regulatory incentives in the context of volatile fuel

- prices. Research quantifying the importance of this feature would be useful.
4. Even if biases in consumer choice are small on average, they could nevertheless be quite large for some consumers. Too little is known about the effects of fuel-economy labels, information programs, and nudges on vehicle choice in this context.
 5. Much of the most credible empirical work estimates short-run effects, whereas long-run effects are relevant for policy. Innovative research designs that identify longer-run impacts and technological tradeoffs would make important contributions.
 6. The gap between on-road fuel economy and laboratory tests needs to be better understood, and the implications of gaming for policy design should be explored.
 7. Research aimed at understanding the political economy of fuel taxation might help identify scenarios in which such policies would be more politically palatable.

The paper is structured as follows. In section 2, we lay out a simplified model of the car market, which depicts a representative consumer making choices over vehicle fuel economy, other vehicle attributes, and vehicle miles traveled. In section 3, using our theory as a guide, we describe the incentives created by a fuel tax, which we then contrast with the incentives created by other policies, including fuel-economy standards, feebates, and attribute-based standards. Section 4 includes a series of subsections that detail auxiliary issues, including the role of behavioral biases, heterogeneity, supply-side distortions, gaming of emissions ratings, safety externalities, other environmental externalities, technological spillovers, and biofuels policies. We discuss findings from both the theoretical and empirical literature throughout our review. Section 5 concludes.

2. Theoretical framework

This section lays out a stylized model of the automobile sector. It is useful in describing the key distortions and margins of choice that constitute the objects of inquiry for the literature. We use it as a guide to constructing our review.

2.1. Consumer welfare

We assume that the representative consumer has utility that is quasilinear in transportation services and other goods, valued in dollars, such that welfare is also valued in dollars and is given by the following expression:

$$W(g, x, v, n) = \mu(n)\theta(x)u(v) + y - nc(x, g) - n[p + d]gv, \quad (1)$$

where: g is average per-mile fuel consumption; x is a vector of average car attributes (e.g., size and horsepower); v is the average number of miles that each car travels; n is the total number of cars (on a per-capita basis); y is the representative consumer's exogenous income; p is the per-gallon private marginal cost of fuel; and d is the per-gallon marginal social damage from fuel consumption. Note that our quasi-linear utility specification rules out income effects.

The expression $\mu(n)\theta(x)u(v)$ is transportation services generated through the total number of cars, their average attributes (all defined as "goods" that increase utility), and their average utilization, where $\mu'(\cdot), \theta'(\cdot), u'(\cdot) > 0$ and $\mu''(\cdot), \theta''(\cdot), u''(\cdot) \leq 0$. We interpret $\mu(n)\theta(x)u(v)$ as the quality-adjusted utility derived from driving, scaled by a function of market size. Our simplified specification of utility implies that every mile driven in a high-

quality car is more enjoyable than in a low-quality car, while the value of quality is greater if the car is to be driven more often.

The expression $nc(x, g)$ is total costs for n cars with average attributes x and fuel consumption g , with partial derivatives given by $c_x > 0$, $c_{xx} \geq 0$, $c_g < 0$, $c_{gg} \leq 0$, and $c_{xg} \leq 0$. That is, per-car costs are increasing and concave in attributes, decreasing and concave in per-mile fuel consumption, and the marginal cost of attributes is lower for cars with higher per-mile fuel consumption. Meanwhile, the expression $n(p+d)gv$ is total private expenditures on fuel ($npgv$) plus total external damages from fuel consumption ($ndgv$).

The necessary conditions that characterize welfare maximization are as follows:

$$\frac{\partial W}{\partial g} = -nc_g(x, g) - n[p + d]v = 0 \quad (2)$$

$$\frac{\partial W}{\partial x} = \mu(n)\theta'(x)u(v) - nc_x(x, g) = 0 \quad (3)$$

$$\frac{\partial W}{\partial v} = \mu(n)\theta(x)u'(v) - n[p + d]g = 0 \quad (4)$$

$$\frac{\partial W}{\partial n} = \mu'(n)\theta(x)u(v) - c(x, g) - [p + d]gv = 0. \quad (5)$$

Expression (2) says that the marginal cost of a decrease in per-mile fuel consumption should be set equal to the marginal social cost of the implied fuel savings (i.e., including both private expenditures on fuel and external damages from greenhouse gas emissions). Expression (3) says that the marginal utility of better car attributes (e.g., size) should be set equal to the marginal impact on the car's up-front purchase cost. Expression (4) says that the marginal utility of a mile should be set equal to the per-mile marginal social cost of fuel consumption. Finally, expression (5) says that the marginal utility of a car equals the marginal cost, inclusive of the up-front purchase cost, as well as the social cost of the car's fuel consumption. Denote the solution to these conditions by (g^*, x^*, v^*, n^*) .

Much of the theoretical and empirical literature fails to model fuel economy (g), car attributes (x), miles (v), and market size (n) jointly. As we will show below, however, interactions among these different margins of choice can be quite important for predicting the impacts of policies and for interpreting empirical results in the literature.

2.2. Externalities and internalities

Consumer choices will potentially deviate from the welfare-maximizing conditions above for two reasons. First, consumers will ignore external costs when choosing their privately optimal level of fuel consumption, which depends both on fuel consumption per mile and on utilization. Second, consumers may fail even to make privately optimal decisions if they mis-perceive or mis-calculate the costs of future fuel expenditures at the time they initially purchase their cars. In the presence of these market imperfections—the fuel consumption externality, and the fuel economy “internality” (see Allcott, Mullainathan & Taubinsky 2014 and Allcott & Sunstein 2015)—corrective policies will be necessary to restore economic efficiency.

Let $\beta > 0$ be the share of future fuel consumption that consumers perceive at the time they purchase their car. A value of $\beta = 1$ implies that the consumer properly values a dollar of future fuel savings at exactly \$1, while a value of $\beta < 1$ implies that the consumer undervalues such savings. We assume that this “undervaluation factor” infects all decisions involving a comparison of current costs and benefits against future fuel savings, including

both the number of cars to buy and the fuel economy of those cars. In contrast, we assume that consumers properly value the costs and benefits of miles driven, conditional on fuel economy, and that they properly weigh the benefits of vehicle quality against up-front purchase costs.

2.3. Policy options

Following the literature, we consider the effects of fuel taxes, fuel-economy standards, and attribute-based standards—as well as a tax on cars—that could apply individually or simultaneously to affect behavior along various dimensions of consumer choice:

- A **fuel tax** imposes a per-gallon tax on fuel consumption. Let τ_f be this tax, such that the total tax paid is $\tau_f n g v$.
- A **traditional fuel-economy standard** imposes a maximum per-mile fuel consumption rate. When binding, the standard acts as an implicit, revenue-neutral tax on per-mile fuel consumption. Accordingly, let τ_g be this implicit tax rate and let $\sigma > 0$ be the binding standard, such that the revenue-neutral tax imposed is $\tau_g [g - \sigma] n$. In the context of multiple, differentiated car models, revenue neutrality implies that efficient cars are subsidized, while inefficient cars are taxed, which is precisely how a feebate system is typically designed.
- An **attribute-based fuel-economy standard** replaces the fixed standard above with one that depends on car attributes. Let this standard be given by $\sigma(x) > 0$, such that the revenue-neutral tax imposed is now $\tau_g [g - \sigma(x)] n$.
- A **car tax**. Let τ_n be this per-car tax, such that the total tax paid is $\tau_n n$.

These taxes (or subsidies, in the case of a negative tax) create policy incentives or “wedges” that modify consumer behavior, as described immediately below.

2.4. Consumer behavior

In this section, we derive conditions that characterize consumer behavior in the presence of the above policies. We potentially allow for all of the policies to apply simultaneously, but note that any of the policies could be “turned off” by setting the appropriate tax incentives to zero. Further, since the traditional fuel-economy standard is simply a special case of the attribute-based standard (i.e., with $\sigma(x) = \sigma$), we will not separately model the traditional standard. Given these assumptions, the necessary conditions that characterize consumer behavior in the presence of the above policies are as follows:

$$\frac{\partial B}{\partial g} = -n c_g(x, g) - \beta n [p + \tau_f] v - \tau_g n = 0 \quad (6)$$

$$\frac{\partial B}{\partial x} = \mu(n) \theta'(x) u(v) - n c_x(x, g) + \tau_g n \sigma'(x) = 0 \quad (7)$$

$$\frac{\partial B}{\partial v} = \mu(n) \theta(x) u'(v) - n [p + \tau_f] g = 0 \quad (8)$$

$$\frac{\partial B}{\partial n} = \mu'(n) \theta(x) u(v) - c(x, g) - \beta [p + \tau_f] g v - \tau_g [g - \sigma(x)] - \tau_n = 0, \quad (9)$$

where ∂B denotes the consumer’s perceived change in private welfare associated with marginal changes in the respective choice variables. Note in condition (6) that $\beta < 1$ would lead a consumer to under-invest in fuel economy relative to the full-rationality baseline represented by $\beta = 1$. But note also in condition (9) that $\beta < 1$ would lead a consumer

to over-invest in cars in general. Meanwhile, note in conditions (7) and (8) that car attributes and miles traveled are chosen optimally, conditional on the number of cars and their average per-mile fuel consumption.¹ Denote the solution to these conditions by $(\hat{g}, \hat{x}, \hat{v}, \hat{n})$.

2.5. First-best corrective policies

Comparing the welfare-maximizing conditions in (2)-(5) to the consumer behavioral conditions in (6)-(9) implies that the first-best policy will include a fuel tax given by $\tau_f = d$, i.e., a Pigouvian tax set equal to marginal damages. In the absence of consumer undervaluation ($\beta = 1$), the fuel tax alone is the first-best policy. Moreover, no combination of alternative policies can achieve the first-best without a gasoline tax, which is necessary to create the appropriate incentive on the miles-driven margin.

However, in the presence of undervaluation ($\beta < 1$), the first-best policy must also include a fuel-economy standard with an implicit tax rate of $\tau_g = (1 - \beta)(p + d)v^*$ and a uniform target of $\sigma(x) = g^*$, as well as a car tax given by $\tau_n = (1 - \beta)(p + d)g^*v^*$. The standard is then equivalent to a revenue-neutral tax on per-mile fuel consumption, with the tax rate set equal to the “undervalued” savings from marginal improvements in fuel economy, while the car tax is set equal to the “undervalued” portion of total lifetime fuel expenditures. Note that if we replaced the revenue-neutral standard with an explicit tax on per-mile fuel consumption, then we could reduce or even eliminate the tax on cars—but a fuel tax by itself is not sufficient to restore full market efficiency (e.g., an “exaggerated” fuel tax equal to d/β could correct the choices of fuel economy and cars but would over-tax the subsequent choice of miles).

2.6. Limitations of our model

While our model is useful for gaining a broad intuition for the effects of different policies, we make a number of important simplifications. Taken literally, our model assumes a single representative consumer (or consumer type) simultaneously choosing cars, fuel economy, attributes, and miles—while facing a continuous car price surface (either a technological cost frontier or hedonic price equilibrium) and an exogenous price for undifferentiated transportation fuel. We do not explicitly model consumer heterogeneity, uncertainty in costs and benefits, the supply of transportation fuels, consumer choice of conventional vs. low-carbon fuels, the dynamics of new car purchases and used car scrappage, or the fact that there are a discrete number of highly differentiated car models manufactured and sold by oligopolistic firms—potentially important details that are featured in many of the papers we cite. Instead, we discuss these issues qualitatively as extensions to our main model.

3. How do alternative policies affect behavior?

We now describe the incentives created by the fuel tax and the behavioral effects it induces. We then describe the incentives created by alternative policies and the behavioral effects they induce, comparing these incentives and effects to those of the fuel tax.

¹We follow the literature in applying the undervaluation factor β only to future fuel costs. An alternative model would allow the undervaluation factor to infect all margins of behavior, rather than just choices for fuel economy (g) and market size (n). This alternative deserves greater consideration in future work.

3.1. The effects of a fuel tax

Raising the price of fuel via a fuel tax causes consumers to adjust their behavior along a host of margins. While it is possible to work out the full, general comparative statics using our model, this work is tedious and the result is too complex to be useful for our present purpose. Thus, we make a number of simplifying assumptions so that we can show graphically the key mechanisms of consumer response. Figure 1 illustrates these responses.




Figure 1

Effects of the fuel tax on (a) fuel economy, (b) car attributes, and (c) miles traveled. These panels correspond respectively to necessary conditions (6), (7), and (8) in the text with a fuel tax but no other policies (i.e., $\tau_f > 0$ with $\tau_g = \tau_n = 0$). The number of cars (n) is implicitly held fixed in all three panels, while panels (a) and (b) implicitly hold miles traveled fixed, as well. In panel (c), we assume that the long-run effect on miles is negative, which we believe to be the empirically relevant case.

3.1.1. Cars get more efficient and shrink in size. When fuel prices are higher, the benefits of fuel economy are greater. Thus, consumers will demand greater fuel economy (lower g), shifting away from other attributes (x). Panels (a) and (b) of figure 1 illustrate the key mechanisms underlying these changes, assuming for ease of exposition that market size (n) and miles traveled (v) are fixed and that the marginal utility of car attributes ($\theta'(x)$) is constant. In panel (a), an increase in the fuel price prices shifts the marginal benefit of fuel economy ($\beta[p + \tau_f]v$) upward, while the marginal cost in terms of higher vehicle price ($c_g(x, g)$) remains fixed. Thus, optimal per-mile fuel consumption decreases (from point A to point B). Simultaneous reductions in car attributes (x), discussed momentarily, lead the marginal cost of fuel economy to shift downward, reinforcing this effect (from point B to point C). Meanwhile, in panel (b), the marginal cost of car attributes ($c_x(x, g)$) increases as a result of the improved fuel economy, while the marginal benefit ($[\mu(n)/n]\theta'(x)u(v)$) remains constant. Thus, the optimal level of attributes also decreases (from point A to B). Intuitively, consumers respond to higher fuel prices both by paying more for cars with additional fuel-saving technologies, conditional on attributes, and by choosing cars with fewer attributes, since both mechanisms save fuel.²

²In appendix A, we reframe our model such that the consumer chooses car attributes x and car technology c , with per-mile fuel consumption $g(x, c)$ given as a direct function of car attributes and technology. This reframed model better highlights the technological tradeoff between per-mile fuel consumption and car attributes by showing that—holding car technology fixed—consumers will

We are not aware of any purely empirical papers that estimate how consumers trade off costs, attributes, and fuel economy in the long run, after allowing for the deployment of new fuel-saving technologies. However, many empirical papers in the literature model consumers' choices of cars from among the current set of model offerings. These papers invariably show that higher fuel prices shift market shares toward smaller, more efficient cars. For example, Busse, Knittel & Zettelmeyer (2013) present reduced-form estimates for 1999–2008, which imply that a \$1 increase in gasoline prices leads to a 3.7% increase in new-car fuel economy.³ Meanwhile, Beresteanu & Li (2011) present a simulation based on structural estimation of demand and supply for 1999–2006, showing that a \$1 increase in gasoline taxes would increase new-car fuel economy by 4.5% (1.01 mpg on a 22.4 mpg base).

All of these studies, however, likely understate the long-run effects of fuel prices on fuel economy, since they implicitly treat as fixed the observed cross-sectional relationship that exists in the new car market between costs, attributes, and fuel economy—in the longer run, we would expect car models to be redesigned and outfitted with new technologies to save more fuel. Thus, studies attempting to model longer-run changes typically augment their empirical analysis with data on engineering costs (see Austin & Dinan 2005, Whitefoot & Skerlos 2012, Jacobsen 2013a, Whitefoot, Fowlie & Skerlos 2013), which can be difficult for economists to validate. Thus, we view the development of new empirical methods for estimating technology-cost tradeoffs as an important avenue for future research.

3.1.2. Consumers drive fewer miles. Conditional on car attributes (x) and fuel economy (g), increasing the price of fuel will increase the per-mile cost of driving while leaving marginal benefits unchanged—and vehicle miles traveled (v) will therefore fall. Meanwhile, the simultaneous reduction in per-mile fuel consumption (g) discussed above will tend to weaken this effect, as consumers demand greater fuel economy, while the simultaneous reduction in car attributes (x) will tend to reinforce this effect by lowering the quality-adjusted utility from driving. Panel (c) of figure 1 illustrates these responses, assuming for ease of exposition that market size (n) is fixed and that the marginal utility of car attributes ($\theta'(x)$) is constant. An increase in fuel prices shifts the per-mile marginal cost of driving ($[p + \tau_f]g$) upward (to the dashed blue line), while the per-mile marginal benefit ($[\mu(n)/n]\theta(x)u'(v)$) remains fixed. Thus, miles traveled decreases (from point A to point B). Note that this partial response, in which car attributes and fuel economy are held fixed, corresponds conceptually to the “short-run” elasticity of miles and fuel demand with respect to changes in fuel prices, as discussed below. The simultaneous reduction in car attributes (x) and per-mile fuel consumption (g) discussed above will, however, lead both the per-mile marginal benefit and per-mile marginal cost of driving to shift downward. Thus, the long-run response of miles traveled to higher fuel prices (from point A to point C) will generally diverge from the short-run response.

A voluminous literature estimates the short-run and long-run response of miles traveled to changes in fuel prices, with a central tendency of around -0.15 in the short run and -0.30 in the long run. For reviews, see Goodwin (1992), Greening, Greene & Difiglio

respond to higher fuel prices by directly trading off car attributes for improved fuel economy.

³They estimate and report changes in market shares by fuel economy quartile. We calculate the implied change in average new-car fuel economy, based on mean fuel economy for each quartile in 2006, as reported in their online appendix.

(2000), Graham & Glaister (2002), and Goodwin, Dargay & Hanly (2004). For more recent examples, see Gillingham, Jenn & Azevedo (forthcoming-a) and Gillingham (2014a), who estimate static models using individual-level micro data, as well as Small & Van Dender (2007), who estimate a dynamic model using state-aggregate panel data. Meanwhile, an even more voluminous literature estimates the short-run price elasticity of fuel demand, with a central tendency of around -0.25 . Short-run responses for miles traveled and fuel demand tend to align closely, since on-road fuel economy is largely fixed in the short run. For reviews, see Dahl & Sterner (1991), Goodwin (1992), Espey (1998), Graham & Glaister (2002), and Goodwin et al. (2004). For recent examples that more credibly identify demand responses using shifts in fuel supply, see Kilian (2010), who use a structural VAR approach, Coglianesi, Davis, Kilian & Stock (2015), who use changes in state fuel taxes as instruments for retail prices, and Levin, Lewis & Wolak (2015), who use city-level weekly panel data and a rich set of fixed effects to control flexibly for potential demand shifters.

3.1.3. Consumers hold fewer cars. An increase in the fuel price makes cars more expensive to own and operate. This implies that the new automobile market can be expected to shrink overall. The mechanisms underlying this effect can be seen in condition (9). Holding miles and car attributes fixed, an increase in fuel prices shifts per-car fueling costs $(\beta[p + \tau_f]gv)$ upward, while the marginal benefit of a car net of up-front purchase costs $(\mu'(n)\theta(x)u(v) - c(x, g))$ remains fixed. Thus, the number of cars declines. Simultaneous changes in miles, fuel economy, and car attributes will lead to further shifts in costs and benefits, but the net effect should be an overall decline in market size. See Goodwin et al. (2004) for a review. Finally, most discrete-choice models of car demand include an “outside good” option and therefore implicitly feature market size effects—although they rarely emphasize such effects in their simulations.

3.1.4. Consumers scrap inefficient cars earlier. A richer understanding of how fuel taxes impact behavior is possible if we move beyond our model and consider dynamic effects and heterogeneity in the car market. A fuel tax raises the optimal long-run fuel economy of the car fleet, which—in addition to shifting market shares toward more efficient new cars—will hasten the scrapping of inefficient used cars relative to efficient cars. Consistent with this intuition, Jacobsen & van Benthem (2015) estimate that when gasoline prices increase by \$1, the annual scrap rate for the least efficient quartile of cars increases by 0.7 percentage points relative to the most efficient quartile, and that this effect is highly concentrated in older cars (see their online appendix for these estimates). Li, Timmins & von Haefen (2009) present qualitatively similar results. Note, however, that effects in the new car market will work against these scrapping effects, because new and used cars are substitutes. A contraction of the new car market will mute the acceleration in scrapping from a gasoline price increase, and likewise a shift towards more efficient new cars can be expected to mitigate the relative acceleration of scrapping among inefficient new cars.

3.2. The effects of a traditional fuel-economy standard

A traditional fuel-economy standard (i.e., with no attribute-basing) also induces a host of responses. Figure 2 illustrates these responses, which differ substantially from those of the fuel tax.




Figure 2

Effects of the traditional fuel economy standard on (a) fuel economy, (b) car attributes, and (c) miles traveled. These panels correspond respectively to necessary conditions (6), (7), and (8) in the text with a binding fuel economy standard but no other policies and no attribute-basing (i.e., $\tau_g > 0$ and $\sigma(x) = \sigma$, such that $\sigma'(x) = 0$, with $\tau_f = \tau_n = 0$). The number of cars (n) is implicitly held fixed in all three panels, while panels (a) and (b) implicitly hold miles traveled fixed, as well.

3.2.1. Cars get more efficient and shrink in size—same as the fuel tax. Like a tax on fuel, a fuel-economy standard directly shifts consumer demand toward greater fuel economy, which will lead to a decrease in per-mile fuel consumption (g) and car performance attributes (x) that compromise fuel economy. These effects are illustrated in panels (a) and (b) of figure 2. In fact, holding market size (n) and miles traveled (v) fixed, a fuel-economy standard can be designed to have precisely the same impact on fuel economy and car attributes as a fuel tax by setting the tax rate on per-mile fuel consumption (τ_g) equal to the perceived per-mile fuel tax that would be paid by each car ($\beta\tau_f v$). But the similarity between a fuel-economy standard and a fuel tax ends there.

3.2.2. As the price of driving falls, consumers drive more—the “rebound effect”. Unlike the fuel tax, a fuel-economy standard provides no direct incentive to reduce miles traveled. Worse, a fuel economy standard can have a perverse effect: by raising fuel economy while leaving fuel prices unchanged, the standard lowers the per-mile cost of driving, and consumers might therefore be expected to drive more. This intuition is illustrated in panel (c) of figure 2 as a downward shift in the marginal per-mile cost of driving (pg), which, holding car attributes and therefore the marginal benefit from driving ($[\mu(n)/n]\theta(x)u'(v)$) fixed, leads to an increase in miles traveled (from point A to point B). This response is known as the “rebound effect” and it is the subject of a voluminous literature (for recent reviews and micro-theoretic analyses, see Gillingham 2014b, Borenstein 2015, Chan & Gillingham 2015, and Gillingham, Rapson & Wagner forthcoming-b). In the presence of various unpriced externalities that scale with miles traveled (e.g., car accidents, congestion, and local pollution), this increase in miles traveled could potentially more than offset the reduction in external costs associated with greenhouse gas emissions. See section 4.6 for detail.

3.2.3. The “rebound effect” may be weakened by shrinking car size. Traditional empirical estimates of the rebound effect have been drawn from studies that identify the derivative of miles traveled (v) with respect to per-mile fuel costs (pg) using time-series variation in fuel prices, while implicitly holding vehicle attributes (x) constant—corresponding to the

behavioral response described immediately above. But note that consumers will respond to fuel economy standards, in part, by choosing smaller cars—which, in the context of our model, generate fewer benefits per mile. Thus, adjustments in car attributes will tend to weaken the rebound effect. This intuition is illustrated in panel (c) of figure 2 as a downward shift in the marginal benefits from driving ($[\mu(n)/n]\theta(x)u'(v)$), such that the net effect of fuel economy standards on miles traveled (now, from point A to C) is weaker than when attributes are held constant—and the net response of miles could potentially be zero or even negative.

Consistent with this theoretical prediction, West, Hoekstra, Meer & Puller (2015) use a compelling regression discontinuity design to show that households that were “just barely eligible” for the U.S. cash-for-clunkers program of 2009—based on the age and fuel economy of their used, trade-in vehicles—purchased new cars with 3.9% higher fuel economy than those purchased by very similar ineligible households. Meanwhile, these eligible households did not travel any more (or fewer) miles. These results are relevant for understanding the potential impacts of fuel economy standards since the cash-for-clunkers program effectively operated as a direct subsidy to the fuel economy of new vehicles, closely approximating the incentives generated by a fuel-economy standard.

But two important notes of caution are warranted. First, the cash-for-clunkers subsidy was short-lived—only eight weeks. Thus, West et al. (2015) implicitly estimate the behavior of consumers constrained to choose between the car models offered for sale during summer 2009. Miles traveled could very well increase in the long run, as carmakers deploy fuel-saving technologies that relax the tradeoff between fuel economy and car attributes. Second, even if the rebound effect net of attribute changes is precisely zero, a fuel tax is still more efficient than a fuel-economy standard, because the fuel tax directly increases the cost of driving, reducing miles traveled among both new and used cars (note that the corresponding decrease in car attributes will tend to reinforce this effect). This response can be qualitatively important (see Austin & Dinan 2005).

3.2.4. Consumers have zero direct incentive to hold fewer cars. A fuel tax raises the operating cost of all automobiles, but a revenue-neutral fuel-economy standard implicitly taxes some products while subsidizing others. This fact is clear in our definition of the policy in section 2.3, which shows that a car is taxed based on the difference between its per-mile fuel consumption (g) and the standard (σ). Intuitively, this means that a fuel-economy standard could plausibly tilt relative car prices to reflect external damages, but it cannot directly change average prices. Thus, a fuel-economy standard fails to optimally shrink market size—and could even increase it. Holland, Hughes & Knittel (2009) make these points about performance standards more generally.

3.2.5. Consumers scrap inefficient cars later—the “Gruenspecht effect”. As discussed above, a fuel tax will shrink the market share of inefficient new cars, while hastening the scrapping of inefficient used cars. A fuel-economy standard will also shrink the market share of inefficient new cars. However, unlike a fuel tax, a fuel-economy standard does not increase the operating costs of inefficient used cars. Meanwhile, under the intuitive assumption that a large, inefficient used car will often be the closest substitute to a large, inefficient new car, the contraction in the new-car market will, via substitution, tend to increase the resale values of inefficient used cars. Thus, on net, the resale values of inefficient used cars increase, perversely extending their lives—the so-called “Gruenspecht effect”

(see Gruenspecht 1982). Jacobsen & van Benthem (2015) explore this dynamic in detail, using variation in used-car resale values induced by changes in gasoline prices to estimate the “scrap elasticity”—the effect of resale value on the probability of scrappage. They combine these estimates with a rich model of car supply and demand to simulate the effects of fuel-economy standards on overall fuel consumption, showing that 13%–16% of the savings anticipated for new cars is lost via “leakage” to the used car market—comparable to the “leakage” due to the rebound effect.

3.2.6. Cost-effectiveness is undermined by firm, fleet, and year-specific standards. Our discussion above implicitly considers a second-best fuel-economy standard whose structure is chosen optimally. But fuel-economy standards in practice fall far short of this second-best ideal. First, standards historically have regulated automaker and fleet-specific averages, leading—in the absence of efficient trading—to different marginal incentives to improve fuel economy for different firms and products. Although trading between firms has been permitted under U.S. standards since 2012, very little trading has actually taken place so far, and the future success of this new market remains uncertain (Leard & McConnell 2015). Second, because light trucks face a less stringent standard, there is a potentially strong incentive for automakers to “re-classify” their large cars (e.g., station wagons) to be small light trucks (e.g., minivans and SUVs) through superficial design changes. Re-classifying a large car to be a small truck in this way raises the average fuel economy of the remaining car models, while simultaneously raising the average fuel economy of the now-expanded set of truck models, which undermines the overall effectiveness of the standards (see Henderson 1985 and Sallee 2011). Finally, shifts in consumer demand due to volatile fuel prices imply that the marginal incentive to improve fuel economy is also quite volatile over time, in spite of limited provisions for credit banking and borrowing. This volatility implies potentially large departures from cost-effectiveness, given that the marginal benefits of carbon abatement are relatively constant over time. We return to this issue below.

3.3. Is a feebate the same as a fuel-economy standard?

A “feebate” system is a series of “fees” on inefficient models and “rebates” for efficient models. A major theme in environmental regulation is the close symmetry between price and quantity instruments—and these same ideas and intuitions carry over to the comparison of a feebate (a “price” mechanism) with a fuel-economy standard (a “quantity” mechanism). A pollution tax makes explicit the implicit tax imposed by a cap on pollution. Likewise, a feebate simply makes explicit the implicit system of taxes and subsidies associated with a fuel-economy standard—thus, the effects of these policies are very similar.⁴ See Gillingham (2013) for analysis and extended discussion of feebates and how they compare to fuel-economy standards, echoing many of our points, and see Sallee (2011) for a review of tax and subsidy policies related to fuel economy more generally.

In principle, a feebate allows for a positive tax on all vehicles, more closely approximating the effects of a first-best tax on fuel and raising revenue to offset distortionary taxes. In

⁴As with other price and quantity mechanisms, this symmetry can break down when considering policy interactions. For example, in the presence of a binding fuel-economy standard, raising the fuel tax would have no impact on fuel economy—it would only lower the shadow price on the fuel-economy constraint. In the presence of a feebate, however, raising the fuel tax would further increase fuel economy.

practice, feebates have typically been designed to be revenue-neutral, largely for political reasons, obviating these purported advantages. Moreover, a fuel-economy standard could be combined with a cap (or a tax) on the number of cars to achieve the same effects. Thus, the perceived differences of a feebate are more about political and administrative convenience. Likewise, it is sometimes asserted that a feebate allows for greater flexibility in pricing—for example, by allowing the tax schedule to be a nonlinear function of fuel consumption rates. But a fuel-economy standard could achieve the same effects by re-weighting vehicles based on a nonlinear function of fuel consumption when calculating fleet-average fuel economy, again reflecting the close symmetry between these price and quantity instruments. Moreover, departures from linear pricing come with costs. Indeed, feebates everywhere have featured highly nonlinear discrete jumps or “notches” in their tax and subsidy schedules. Sallee & Slemrod (2012) demonstrate that carmakers reacted to notches in the U.S. Gas Guzzler Tax by bunching vehicles around favorable fuel consumption rates. They show theoretically that such a response is inefficient and argue for the use of smooth schedules instead of notched ones.

One real benefit of feebates is that they harmonize incentives to improve fuel economy across firms. Although this same benefit can, in principle, be achieved in a standard by allowing compliance trading across firms, efficient trading could be limited for a variety of reasons, including the small number of firms (Leard & McConnell 2015).

Another real benefit of feebates in practice is that they provide certainty about the marginal incentive to improve fuel economy. The marginal benefits from carbon abatement in the U.S. transportation sector, arguably, are approximately constant within and across model years. Thus, standard reasoning implies that setting a certain “price” via the feebate and being uncertain about the level of fuel economy is better (i.e., has lower expected deadweight loss) than setting a certain “quantity” via the standard and being uncertain about marginal compliance costs (see Weitzman 1974). This logic is especially compelling when one considers the difficulty of estimating the marginal cost of a given standard, even *ex post*—indeed, several studies have attempted to estimate such costs and have come to widely divergent conclusions (see Anderson & Sallee 2011 for a discussion). Moreover, fuel-economy standards in the United States are typically set several years in advance, while demand for fuel economy fluctuates continuously and unpredictably with oil prices.

While prior discussions comparing feebates to fuel-economy standards have emphasized revenue-neutrality and flexible pricing as benefits of feebates, we argue that the more important benefits are likely related to uncertainty. However, we know of no efforts to quantify these benefits in the literature. This would be a fruitful area for future research.

3.4. What are the effects of an attribute-based fuel-economy standard?

In 2012, the U.S. fuel economy program became attribute-based. This means that, rather than simply requiring all firms to meet a common fuel-economy target, a firm’s target depends on some other attribute. In the United States that attribute is a vehicle’s footprint—the rectangular area bounded by a vehicle’s tires. Larger vehicles are credited with a lower fuel economy target. Canada and Mexico also have footprint-based standards, while other countries with attribute-based standards all have weight-based standards. Thus, the world’s largest auto markets—the United States, Europe, China, Japan, Mexico, Canada, Brazil, South Korea, and India—all have size-based standards (U.S. Energy Information Administration 2015).

3.4.1. Attribute-based standards implicitly subsidize car size, potentially leading to up-sizing. Attribute-basing changes the economics of standards by introducing an additional margin of compliance—changes in the attribute (e.g., car size). In our conceptual model, this effect is captured by showing that an attribute-based standard creates an implicit subsidy to the attribute (the $\tau_g n \sigma'(x)$ term in equation 7, where we have implicitly assumed that $\sigma'(x) > 0$), which accompanies the implicit tax on fuel economy (the $\tau_g n$ term in equation 6). Holding the implicit tax on per-mile fuel consumption (τ_g) fixed, this extra incentive would manifest in figure 2 as a smaller downward shift in the marginal cost of fuel economy improvements ($-c_g(x, g)$) in panel (a), and a smaller upward shift in the marginal cost of the attribute ($c_g(x, g)$) in panel (b)—implying smaller reductions in both per-mile fuel consumption and car attributes. Thus, relative to a traditional fuel-economy standard with the same incentive rate (τ_g), the attribute-based standard leads to up-sizing and implies lower fuel economy.

Motivated by these theoretical predictions, Ito & Sallee (2015) consider the case of Japan’s weight-based standards. Though broadly set to match the pre-regulation correlation between fuel economy and weight, these standards nevertheless feature large “notches” below which the incentive to increase weight is extremely powerful (note that the slope of the weight-based standard at the notch is $\sigma'(x) = \infty$). Accordingly, they find that these notches are associated with significant “excess bunching” in the distribution of weight on the policy-favored side of these notches. They show that this bunching can be used to infer the increase in average vehicle weight in Japan that is caused by the attribute-based policy, as compared to a flat fuel-economy standard featuring the same shadow price, and conclude that the weight increase in Japan is economically significant.

We are not aware of any similar studies that have used ex-post data to examine the effects of footprint-based standards in the United States—where regulators explicitly chose footprint, in part, because they believed it would be more difficult for carmakers to manipulate than weight. Nevertheless, Whitefoot & Skerlos (2012) use an engineering cost model to conclude that distortions to footprint will be substantial. Meanwhile, the same perverse incentive to up-size cars potentially exists in the European Union, China, and India—all of which feature weight-based regulations—although Reynaert (2015) finds no strong evidence of weight manipulation in Europe. Additional research on the actual effects of attribute-based standards in other markets would be valuable, given the growing prevalence of such standards.

3.4.2. When calibrated to match pre-regulation market data, attribute-based standards neutralize the incentive to up-size and instead force technology adoption. How do regulators choose the slope of an attribute-based policy? In practice, most regulators have chosen the slope of the fuel-economy target ($\sigma'(x)$) by fitting a function to historical or projected data on fuel economy vs. car size across the range of car models offered in the market. How should we interpret this slope? Intuitively, we can think of this slope as deriving from heterogeneity in consumer preferences for car size. This heterogeneity then traces out an equilibrium relationship between fuel economy and car size, as consumers with stronger preferences for size—facing the same cost and technological tradeoffs—optimally choose bigger, less efficient cars.

What then are the consequences of choosing the slope of an attribute-based standard to match the slope of this equilibrium relationship? Intuitively, the equilibrium relationship between fuel economy and size defines the “baseline” car choices made by consumers prior to

regulation. Thus, in setting the target fuel economy function to be parallel to this baseline, an attribute-based standard effectively neutralizes the incentive to comply via moving left or right—that is, by up-sizing or down-sizing individual cars or by up-sizing or down-sizing the “average” car via mix-shifting. Instead, the only compliance benefit is in moving up—that is, adding fuel-saving technologies.

Consistent with this theoretical prediction, Reynaert (2015) studies the roll-out of weight-based standards in the European Union, finding that compliance there is achieved entirely through technology gains—which are found to accelerate sharply when the weight-based standard is introduced. Note that this finding does not rule out the counterfactual possibility that, had the standard been flat, the fleet would have been downsized as part of the compliance strategy. As such, the findings of Reynaert (2015) are consistent with our theory and with the findings of Ito & Sallee (2015).

3.4.3. Is there any efficiency or other justification for attribute-based standards? As we note above, there is no efficiency justification for attribute-based standards in our conceptual model—the optimal policy is a Pigouvian tax on fuel or (if v is exogenous) a tax on per-mile fuel consumption with no attribute-basing. These same points are made by Ito & Sallee (2015), who emphasize that the practical rationale for attribute-based standards is likely distributional. Attribute-based standards spread compliance burdens more equally and tend to favor makers of bigger and heavier cars—domestic carmakers in the United States and German carmakers in the European Union.

However, these distributional aims are achieved at the cost of introducing several possible inefficiencies. First, the implicit subsidy to attributes under attribute-based standards distorts consumers’ choices. As Ito & Sallee (2015) show in the context of Japan’s weight-based standards, the implicit subsidy to weight creates a Harberger-like deadweight loss triangle between the marginal benefits and costs of weight—the more elastic is weight, the larger is this deadweight loss. Second, attribute-based standards substantially mute the incentive to down-size cars, which must weakly increase the cost of achieving a given fuel-economy target. Ito & Sallee (2015) note that forcing technology adoption could potentially be justified by technology spillovers, but they are skeptical that these spillovers are large, given extensive patenting and licensing in the car market. Third, weight-based standards exacerbate—or at least fail to mitigate—the safety-related externality that heavier cars impose on occupants of other cars (see section 4.5 below). Finally, by neutralizing the incentive to down-size and forcing technology adoption, attribute-based standards likely generate a larger rebound effect than traditional standards. Thus, while we argue above that the “holding attributes constant” empirical approaches common in the literature likely overstate the size of the rebound effect for conventional fuel-economy standards, these concerns are rightly placed on an attribute-based standard.

Our discussion above ignores uncertainty. In a world with variable demand for fuel economy due to volatile fuel prices, attribute-based standards—like feebates—have the potential benefit of stabilizing the marginal cost of regulation. Under attribute-based standards, target fuel economy would rise and fall in tandem with fuel prices, as demand shifted between big and small cars. Thus, these shifts would have little effect on the regulatory shadow price. Instead, the shadow price would be driven mainly by the costs less benefits of fuel-saving technology, which would vary some—but much less—with fuel prices. As noted above, however, a feebate could also stabilize incentives but without the concomitant distortions.

3.4.4. Attribute-biasing undermines the logic of using standards to counteract myopia.

Recent fuel-economy rules in the United States have, in their supporting documents, emphasized consumer myopia (or other sources of consumer undervaluation) as a key justification for regulation. However, attribute-based standards partially undermine this logic by neutralizing the incentive to down-size cars and instead forcing the adoption of fuel-saving technologies. There is no reason to believe that myopia, if it exists, applies only to the valuation of fuel-saving technologies for cars with identical attributes and not to comparisons of cars with different attributes—indeed, empirical evidence for myopia is based largely on observing how consumers trade off car size for greater fuel economy when fuel prices rise, as we discuss in detail below. Thus, if myopia is the rationale for regulation, the appropriate policy is a fuel-economy standard that preserves incentives for down-sizing.

4. Additional considerations

In this section, we cover a number of auxiliary topics that do follow directly from the discussion of our conceptual model.

4.1. Behavioral biases do little to justify fuel-economy standards

A substantial literature in energy economics studies the possibility that consumers might fail to make rational investments in energy efficiency, due to present bias or some other behavioral bias or market friction. Might behavioral biases overturn the traditional wisdom that a fuel tax is more efficient than standards?

As we discuss above, consumer undervaluation of fuel economy ($\beta < 1$) creates a second market failure, such that a fuel tax must be complemented by a fuel-economy standard (or direct subsidy) to restore full market efficiency. These results match those in Allcott et al. (2014) and Heutel (2015)—although a tax on cars is also required in our setting. Note, however, that unpriced miles-related externalities (e.g., congestion) weaken the case for a complementary fuel-economy standard, since the standard induces a rebound effect. Indeed, given these un-priced externalities, and assuming an optimal fuel tax, Parry, Evans & Oates (2014b) argue that an extreme value of undervaluation (with $\beta < 0.3$) would be required before a complementary fuel-economy standard becomes welfare-improving.

Is such a large degree of undervaluation consistent with the data? In short: no. We focus on what we believe to be the most credible empirical papers in the literature—those that use panel identification strategies, regressing the market prices of various car models in different time periods on a constructed measure of these cars' future fueling costs (see Helfand & Wolverton 2011 and Greene 2010 for recent reviews of the broader literature). This approach enables the inclusion of both time and car model fixed effects—thereby controlling for unobserved car attributes and macroeconomic shocks—while still identifying the effect of future fuel costs on vehicle prices. Kahn (1986) pioneered this approach using a handful of car models, while more recent iterations—notably, Allcott & Wozny (2014), Busse et al. (2013), and Sallee, West & Fan (forthcoming)—use much more precise and comprehensive data. These three papers are united in finding substantial valuation (estimates of β are closer to 1 than 0.5). However, all three papers emphasize that their results are somewhat sensitive to chosen samples, specifications, and assumptions about consumers' fuel price forecasts and other objects used to construct future fuel costs (see these papers and Anderson, Kellogg & Sallee 2013). Thus, while their preferred specifications do not

consistently reject full valuation, alternative assumptions can lead to results consistent with modest undervaluation.

What about consumer heterogeneity? Bento, Li & Roth (2012) show in a simulation that falsely assuming a homogeneous valuation parameter can lead to an erroneous finding of mean undervaluation. Meanwhile, Grigolon, Reynaert & Verboven (2014) directly estimate heterogeneity in preferences for fuel economy in a structural model applied to the European car market, finding that the average car buyer fully values fuel economy. Again, however, the literature cannot rule out that at least some consumers undervalue fuel economy.

Such heterogeneity complicates policy analysis since a fuel-economy standard that corrects the biases of “behavioral” agents will simultaneously distort the choices of “rational” agents (see Allcott et al. 2014). Instead, behavioral nudges, information provision, and fuel consumption labels might be better targeted to behavioral agents and could potentially have large effects. Indeed, Allcott (2013) finds that there is wide dispersion in the accuracy of beliefs about fuel costs and that consumers report spending very little time considering such costs. Meanwhile, Davis & Metcalf (2014) and Newell & Siikamäki (2014) suggest that information design can have important impacts on appliance choice. More such work is needed, in particular for automobiles.

4.2. Heterogeneity in vehicle longevity and on-road fuel economy strengthen the case for a fuel tax

Fuel-economy standards (and feebates) only affect up-front purchase decisions, treating all cars with the same fuel economy identically. Thus, such policies fail to account for substantial heterogeneity in actual emissions across both cars and drivers, implying wide dispersion in marginal abatement costs. One source of heterogeneity is vehicle longevity (durability). Jacobsen, Knittel, Sallee & van Benthem (2016) show that different types of cars have vastly different values for expected lifetime mileage. They develop a sufficient statistics approach for characterizing the welfare implications of a fuel-economy standard in light of this heterogeneity and conclude that the second-best standard achieves only one-quarter to one-third of the welfare gains from a fuel tax. Another source of heterogeneity is driver behavior. Langer & McRae (2014) observe wide dispersion in on-road fuel economy across drivers operating a literally identical car in the same metro area, while Reynaert & Sallee (2016) documents substantial variation using on-road data from the Netherlands, and Sallee (2011) argues that the disparity between city and highway fuel economy in government tests implies substantial heterogeneity across drivers based on where and when they drive. In contrast to the standard, the fuel tax scales directly with total emissions, equalizing marginal abatement costs across different cars and drivers—and without requiring any additional information on the part of the policy maker.

4.3. The implications of imperfect competition for energy policy are not well-understood

Our model focuses on the consumer problem and posits a competitive supply side. In reality, the automobile market is composed of a handful of major automakers who sell differentiated products and thus enjoy some degree of market power. How might this influence our conclusions about policy design?

In general, welfare analysis is complicated by the presence of market power so that it is impossible to know *ex ante* whether any particular attribute (e.g., fuel economy) is over or

under-provided in the private equilibrium. First, firms will price to the marginal consumer for a product, rather than the average (Spence 1975). Next, firms in the automobile market sell multiple differentiated products, so the welfare analysis is further complicated by the potential incentives of firms to attempt to price discriminate. As argued in Plourde & Bardis (1999) and Fischer (2005), firms might use quality (e.g., fuel economy) as a way of segmenting the market, which can lead to inefficient provision of attributes (e.g., fuel economy) to some market segments as a way of extracting surplus from other segments.

Thus, in general, imperfect competition could lead to either over or under-provision of fuel economy. Testing these alternatives is very difficult and the literature provides little guidance or evidence on this issue. Structural models typically estimate markups and separate producer and consumer surplus, but those models generally take product attributes as fixed—allowing only prices to change—so they are unable to tell us how policy design might affect fuel economy in the longer run (a notable exception is Whitefoot et al. 2013).

There are, however, some recent intriguing results regarding energy-efficiency policies and supply-side behavior in the appliance market. Houde (2014) argues that imperfectly competitive firms respond to Energy Star certification of appliances by increasing markups on certified products, while Houde & Spurlock (2015) demonstrate that energy-efficiency standards are correlated with technological innovation in the appliance market, which they argue is rationalized by imperfect competition. Clearly, greater consideration of the supply side in the car market is an area where future research could have a major impact—although it is equally apparent that the empirical challenges are substantial.

4.4. Cheating erodes the efficiency of standards

Fuel-economy standards and feebates must rely on official fuel-economy ratings based on laboratory tests, which may differ from on-road performance. Indeed, mounting evidence gathered by industry sources shows a remarkable divergence between official ratings and on-road performance. This gap is modest in the United States (The International Council on Clean Transportation 2013), but it is very large (around 40% in 2014) in Europe (The International Council on Clean Transportation 2015). Such gaps may, in fact, be due to the intentional gaming of test procedures, as evidenced by the 2015 Volkswagen scandal involving its NOx emissions control system or controversy in 2012 that forced Hyundai/Kia to downgrade fuel economy ratings for most of its cars.

Reynaert & Sallee (2016) analyze the welfare implications of such cheating. When regulators are aware of evasion and the cost of evasion is homogeneous across firms, regulators can still achieve the intended level of fuel economy by tightening standards until actual fuel economy reaches the desired level. But the gaming itself is socially wasteful. Moreover, when the cost of gaming is heterogeneous across cars, consumers will generally be unable to infer the precise fuel economy of each car, leading to inefficient car choices and further lowering welfare. Reynaert & Sallee (2016) use data on heterogeneity in gaming and a discrete choice model to argue that these welfare costs are substantial. Carmakers cannot cheat the gas tax in the same way—although consumers could still be misled into making inefficient car choices. Thus, gaming provides yet another potential benefit of a fuel tax over fuel-economy standards, though the size of this benefit will depend on the ability of consumers (perhaps in conjunction with third-party testing) to see through such gaming and observe true fuel economy.

4.5. Safety-related externalities are large

A politically and scientifically controversial issue in regulating fuel economy has been the potential adverse impact on car safety, as it is intuitive that larger cars should provide their occupants more protection in accidents, but potentially increase the risk to occupants of other cars—a risk that may be exacerbated if cars and light trucks have “mismatched” body styles (see Transportation Research Board & National Research Council 2002). A major rationale in moving to attribute-based standards in the United States has therefore been the assumption that such standards would substantially mute the incentive to down-size the fleet and thereby neutralize this controversial issue (see U.S. Environmental Protection Agency 2010).

This is a complex issue, as overall fatality risks depend not only on the overall mix of cars on the road, but on where, when, how much, and by whom these cars are driven. To address these issues, Jacobsen (2013b) estimates fatality risks for accidents involving different combinations of car classes, while controlling for driver behavior and other factors. His findings help explain why, for example, minivans cause far fewer fatalities than their large size would imply—these cars tend to be driven by safer drivers, or at safer times and locations. Jacobsen (2013b) then combines these estimates with a model of the car market to simulate the impacts of fuel-economy standards on fatalities, finding that 1 mpg increase would cause 149 extra deaths per year—but that a unified car and truck standard or a footprint-based standard would have negligible impacts on fatalities. In contrast, Anderson & Auffhammer (2014) focus specifically on the externality related to safety, which is what matters for assessing the relative efficiency of competing policies. Using a rich dataset on car accidents, they estimate that a 100-pound increase in weight is associated with a 4%–5% increase in the probability that the other car’s occupants are killed in a serious accident. They argue that these results would justify a gasoline tax of approximately \$1, simply to address the weight externality. We discuss implications for policy immediately below.

4.6. Accounting for other un-priced externalities reinforces the efficiency advantage of a fuel tax

Thus far, we have focused on greenhouse warming externalities, which scale with total fuel consumption. However, Parry, Walls & Harrington (2007) survey evidence on all automobile-related externalities, showing that un-priced externalities that scale with miles—namely, local pollution, congestion, and car accidents—are potentially just as or even more important. Meanwhile, Anderson & Auffhammer (2014) recently quantify a large externality associated specifically with vehicle weight, as discussed above.

Thus, accounting for other un-priced externalities tends to reinforce our efficiency ranking of policies: fuel tax > fuel-economy standard > attribute-based standard. Why? A fuel tax provides a direct incentive to decrease miles—reducing externalities from pollution, congestion, and accidents—whereas traditional fuel-economy standards either fail to mitigate or exacerbate these externalities. Likewise, traditional fuel-economy standards give potentially strong incentives to reduce vehicle size—reducing externalities from accidents—whereas size-based standards either fail to mitigate or exacerbate these externalities. Of course, these second-best considerations would be irrelevant in the presence of a first-best tax on local pollution emissions (or a combination of instruments that replicate such a tax, as in Fullerton & West 2002, 2010), time-of-use congestion pricing, and a weight-differentiated tax on miles to internalize accident risks (see Anderson & Auffhammer 2014).

4.7. Knowledge spillovers may justify complementary policies

If firms are unable to appropriate the full benefits of their own technological advances (i.e., there are knowledge spillovers), then there is another market failure. This market failure could justify policies to subsidize technology (see Anderson, Fischer & Egorenkov 2015). As noted above, attribute-based standards subsidize technology, but they do so in a way that distorts other margins. Thus, to justify the use of such policies, technology spillovers must be substantial. Moreover, where subsidies for technology are deployed in addition to a binding fuel-economy regulation, the subsidy may reduce the cost of complying with the regulation, but it will not lower emissions beyond what is achieved by the regulation alone (due to the logic of overlapping regulation, as discussed in Goulder, Jacobsen & van Benthem 2012).

4.8. Barriers limiting the adoption of biofuels may justify complementary policies

Our discussion above focuses on reducing overall fuel consumption, but an alternative pathway is to switch to low-carbon fuels. Generalizing to multiple fuel types, the first-best policy is a carbon tax, which would raise the prices of all fuels in proportion to their carbon contents—inducing the same behavioral responses as the fuel tax above—while optimally tilting relative prices in favor of low-carbon fuels, such as ethanol.

How do fuel markets respond to such incentives? We do not have a carbon tax, but the federal Renewable Fuel Standard (RFS) and California’s low-carbon fuel standard (LCFS) also tilt relative prices in favor of ethanol and other low-carbon fuels—albeit without directly raising average fuel prices. Thus far, these policies have been accommodated by blending ethanol into regular gasoline for use in conventional cars. However, to blend ethanol in ratios higher than 10%—the so-called blend-wall—requires flexible-fuel vehicles and new retail fueling infrastructure (e.g., pumps and storage tanks). Thus, as we approach the 10% blend wall, potential barriers to the adoption of high-ethanol blends become important—and there are several such barriers.

First, micro-empirical studies on fuel-switching behavior among drivers of flexible-fuel vehicles in the United States (Anderson 2012) and Brazil (Salvo & Huse 2013) indicate that many consumers demand large price discounts before switching to ethanol. These consumer-level barriers are exacerbated by a lack of flexible-fuel vehicles and ethanol refueling infrastructure. Pouliot & Babcock (2014a,b) show in a simulation model that these consumer-level barriers and infrastructure constraints can severely limit ethanol adoption in the short-run. Second, network externalities and coordination failures in adoption—i.e., drivers depending on station owners to install pumps, and station owners depending on drivers to buy flexible-fuel vehicles—could hinder investment in new infrastructure. Indeed, Shriver (2015) estimates that new retailer entry increases flexible-fuel adoption in a county by 6%, but that such entry first depends on the adoption of 300 or more new flexible-fuel vehicles. Thus, current policies that implicitly subsidize the production of flexible-fuel vehicles (see Anderson & Sallee 2011) or that directly subsidize new retail fueling infrastructure (recently and massively expanded under the USDA Biofuels Infrastructure Partnership initiative) could potentially be justified on efficiency grounds. Similar logic would apply to completely new technologies (e.g., hydrogen cars). Finally, there is suggestive evidence that a lack of competition throughout the supply chain may be limiting adoption. For example, Anderson (2012) and Knittel, Meiselman & Stock (2015) show that wholesale fuel prices and ethanol subsidies do not pass through fully to retail prices for high-ethanol

blends—consistent with retail market power—while it is rumored that vertically integrated oil companies are loath to supply alternative fuels at competitive prices, as such fuels would then compete directly with their upstream operations. Additional work should explore these issues and the potential role of policy in addressing them.

5. Conclusion

How can we make cars greener? The chief aim of our review has been to assess the efficiency of fuel-economy standards relative to the benchmark policy of a fuel tax. Traditional economic wisdom advocates the use of a fuel tax (or carbon price) to achieve cost-effective mitigation. Policy makers demonstrate a consistent preference for efficiency standards, however, and several decades of research have demonstrated that this preference is costly—a fuel tax could reduce greenhouse gas emissions by the same amount at substantially lower social cost.

Might feebates or attribute-based standards overturn this conclusion? Or could behavioral biases or other market failures imply an important role for efficiency standards? Our reading of the literature suggests that the case for standards remains weak. As detailed above and enumerated in the introduction, a number of important questions remain, so it is possible that our conclusions could change. In the meantime, however, we believe that policy advocates should focus on pushing carbon pricing (or fuel taxation) around the world as a way of mitigating carbon emissions and other harmful pollutants from automobiles.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Hunt Allcott for helpful suggestions and Jonathan Siegle for valuable research assistance. All errors are our own.

LITERATURE CITED

- Allcott H. 2013. The welfare effects of misperceived product costs: Data and calibrations from the automobile market. *American Economic Journal: Economic Policy* 5:30–66
- Allcott H, Mullainathan S, Taubinsky D. 2014. Energy policy with externalities and internalities. *Journal of Public Economics* 112:72–88
- Allcott H, Sunstein CR. 2015. Regulating internalities. *Journal of Policy Analysis and Management* 34:698–705
- Allcott H, Wozny N. 2014. Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics* 96:779–795
- Anderson M, Auffhammer M. 2014. Pounds that kill: The external costs of vehicle weight. *Review of Economic Studies* 82:535–571
- Anderson S, Fischer C, Egorenkov A. 2015. Overlapping strategies for reducing carbon emissions from the personal transportation sector. Working Paper
- Anderson ST. 2012. The demand for ethanol as a gasoline substitute. *Journal of Environmental Economics and Management* 63:151–168

- Anderson ST, Kellogg R, Sallee JM. 2013. What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management* 66:383–403
- Anderson ST, Parry IW, Sallee JM, Fischer C. 2011. Automobile fuel economy standards: Impacts, efficiency and alternatives. *Review of Environmental Economics and Policy* 5:89–108
- Anderson ST, Sallee JM. 2011. Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. *American Economic Review* 101:1375–1409
- Austin D, Dinan T. 2005. Clearing the air: The costs and consequences of higher CAFE standards and increased gasoline taxes. *Journal of Environmental Economics and Management* 50:562–582
- Bento AM, Li S, Roth K. 2012. Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias. *Economics Letters* 115:44–48
- Beresteanu A, Li S. 2011. Gasoline prices, government support, and the demand for hybrid vehicles in the U.S. *International Economic Review* 52:161–182
- Borenstein S. 2015. A microeconomic framework for evaluating energy efficiency rebound and some implications. *Energy Journal* 36:1–21
- Busse MR, Knittel CR, Zettelmeyer F. 2013. Are consumers myopic? Evidence from new and used car purchases. *American Economic Review* 103:220–256
- Chan N, Gillingham K. 2015. The microeconomic theory of the rebound effect and its welfare implications. *Journal of the Association of Environmental and Resource Economists* 2:133–159
- Coglianese J, Davis LW, Kilian L, Stock JH. 2015. Anticipation, tax avoidance, and the price elasticity of gasoline demand. NBER Working Paper 20980
- Dahl C, Sterner T. 1991. Analyzing gasoline demand elasticities: A survey. *Energy Economics* 13:203–210
- Davis LW. 2014. The cost of global fuel subsidies. *Milken Institute Review*
- Davis LW, Metcalf GE. 2014. Does better information lead to better choices? Evidence from energy-efficiency labels. NBER Working Paper 20720
- Espey M. 1998. Gasoline demand revisited: An international meta-analysis of elasticities. *Energy Economics* 20:273–295
- Fischer C. 2005. On the importance of the supply side in demand-side management. *Energy Economics* 27:165–180
- Fullerton D, West SE. 2002. Can taxes on cars and gasoline mimic an unavailable tax on emissions? *Journal of Environmental Economics and Management* 42:135–157
- Fullerton D, West SE. 2010. Tax and subsidy combinations for the control of car pollution. *The B.E. Journal of Economic Analysis & Policy* 10:1–33
- Gillingham K. 2013. The economics of fuel economy standards versus feebates. Working Paper: National Energy Policy Institute
- Gillingham K. 2014a. Identifying the elasticity of driving: Evidence from a gasoline price shock in California. *Regional Science and Urban Economics* 47:13–24
- Gillingham K. 2014b. Rebound effects. In *The New Palgrave Dictionary of Economics*, eds. SN Durlauf, LE Blume. Basingstoke: Palgrave Macmillan
- Gillingham K, Jenn A, Azevedo IML. forthcoming-a. Heterogeneity in the response to gasoline prices: Evidence from Pennsylvania and implications for the rebound effect. *Energy Economics*
- Gillingham K, Rapson D, Wagner G. forthcoming-b. The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*
- Goodwin P. 1992. A review of new demand elasticities with special reference to short and long run effects of price changes. *Journal of Transport Economics and Policy* 26:155–169
- Goodwin P, Dargay J, Hanly M. 2004. Elasticities of road traffic and fuel consumption with respect to price and income: A review. *Transport Reviews* 24
- Goulder LH, Jacobsen MR, van Benthem AA. 2012. Unintended consequences from nested state and federal regulations: The case of Pavley greenhouse-gas-per-mile limits. *Journal of Environmental Economics and Management* 63:187–207
- Graham DJ, Glaister S. 2002. The demand for automobile fuel: A survey of elasticities. *Journal of*

- Transportation Economics and Policy* 36:1–26
- Greene DL. 2010. How consumers value fuel economy: A literature review. U.S. Environmental Protection Agency report EPA-420-R-10-008
- Greening LA, Greene DL, Difiglio C. 2000. Energy efficiency and consumption—the rebound effect—a survey. *Energy Policy* 28:389–401
- Grigolon L, Reynaert M, Verboven F. 2014. Consumer valuation of fuel costs and the effectiveness of tax policy: Evidence from the European car market. Manuscript: University of Leuven
- Gruenspecht HK. 1982. Differentiated regulation: The case of auto emissions standards. *American Economic Review* 72:328–331
- Helfand G, Wolverton A. 2011. Evaluating the consumer response to fuel economy: A review of the literature. *International Review of Environmental and Resource Economics* 5:103–146
- Henderson DR. 1985. The economics of fuel economy standards. *Regulation* 9:45–48
- Heutel G. 2015. Optimal policy instruments for externality-producing durable goods under present bias. *Journal of Environmental Economics and Management* 72:54–70
- Holland SP, Hughes JE, Knittel CR. 2009. Greenhouse gas reductions under low carbon fuel standards? *American Economic Journal: Economic Policy* 1:106–146
- Houde S. 2014. Bunching with the stars: How firms respond to environmental certification. Manuscript: University of Maryland
- Houde S, Spurlock CA. 2015. Do energy efficiency standards improve quality? Evidence from a revealed preference approach. Manuscript: University of Maryland
- Ito K, Sallee JM. 2015. The economics of attribute-based regulation: Theory and evidence from fuel-economy standards. Manuscript: University of Chicago
- Jacobsen MR. 2013a. Evaluating U.S. fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy* 5:148–187
- Jacobsen MR. 2013b. Fuel economy and safety: The influences of vehicle class and driver behavior. *American Economic Journal: Applied Economics* 5
- Jacobsen MR, Knittel CR, Sallee JM, van Benthem AA. 2016. Sufficient statistics for imperfect externality-correcting policies. Manuscript: University of California at Berkeley
- Jacobsen MR, van Benthem AA. 2015. Vehicle scrappage and gasoline policy. *American Economic Review* 105:1312–1338
- Kahn JA. 1986. Gasoline prices and the used automobile market: A rational expectations asset price approach. *The Quarterly Journal of Economics* 101:323–339
- Kilian L. 2010. Explaining fluctuations in gasoline prices: A joint model of the global crude oil market and the U.S. retail gasoline market. *The Energy Journal* 31:87–112
- Knittel CR, Meiselman BS, Stock JH. 2015. The pass-through of RIN prices to wholesale and retail fuels under the renewable fuel standard. NBER Working Paper 21343
- Langer A, McRae S. 2014. Step on it: Evidence on the variation in on-road fuel economy. Manuscript: University of Arizona
- Leard B, McConnell V. 2015. New markets for pollution and energy efficiency. Resources for the Future Discussion Paper 15-16
- Levin L, Lewis M, Wolak F. 2015. High frequency evidence on the demand for gasoline. Working Paper
- Li S, Timmins C, von Haefen RH. 2009. How do gasoline prices affect fleet fuel economy? *American Economic Journal: Economic Policy* 1:113–137
- Newell RG, Siikamäki J. 2014. Nudging energy efficiency behavior: The role of information labels. *Journal of the Association of Environmental and Resource Economists* 1:555–598
- Parry I, Veung C, Heine D. 2014a. IMF Working Paper WP/14/174
- Parry I, Walls M, Harrington W. 2007. Automobile externalities and policies. *Journal of Economic Literature* 45:373–399
- Parry IW, Evans D, Oates WE. 2014b. Are energy efficiency standards justified? *Journal of Environmental Economics and Management* 67:104–125

- Parry IWH, Small KA. 2005. Does Britain or the United States have the right gasoline tax? *American Economic Review* 95:1276–1289
- Plourde C, Bardis V. 1999. Fuel economy standards in a model of automobile quality. *Energy Economics* 21:309–319
- Pouliot S, Babcock BA. 2014a. The demand for E85: Geographical location and retail capacity constraints. *Energy Economics* 45:134–143
- Pouliot S, Babcock BA. 2014b. Impact of ethanol mandates on fuel prices when ethanol and gasoline are imperfect substitutes. CARD Working Paper 14-WP 551
- Proost S, Van Dender K. 2011. What long-term road transport future? Trends and policy options. *Review of Environmental Economics and Policy* 5:44–65
- Reynaert M. 2015. Abatement strategies and the cost of environmental regulation: Emission standards on the European car market. Manuscript: KU Leuven
- Reynaert M, Sallee JM. 2016. Strategic gaming of environmental ratings: Evidence and welfare implications from automobiles. Manuscript: University of California at Berkeley
- Sallee JM. 2011. The taxation of fuel economy. *Tax Policy and The Economy* 25:1–37
- Sallee JM, Slemrod J. 2012. Car notches: Strategic automaker responses to fuel economy policy. *Journal of Public Economics* 96:981–999
- Sallee JM, West SE, Fan W. forthcoming. Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics*
- Salvo A, Huse C. 2013. Build it, but will they come? evidence from consumer choice between gasoline and sugarcane ethanol. *Journal of Environmental Economics and Management* 66:251–279
- Shriver SK. 2015. Network effects in alternative fuel adoption: Empirical analysis of the market for ethanol. *Marketing Science* 34:78–97
- Small KA, Van Dender K. 2007. Fuel efficiency and motor vehicle travel: The declining rebound effect. *Energy Journal* 28:25–51
- Spence MA. 1975. Monopoly, quality, and regulation. *The Bell Journal of Economics* 6:417–429
- The International Council on Clean Transportation. 2013. From laboratory to road: A comparison of official and “real-world” fuel consumption and CO₂ values for cars in Europe and the united states. ICCT White Paper
- The International Council on Clean Transportation. 2015. From laboratory to road: A 2015 update of official and “real-world” fuel consumption and CO₂ values for passenger cars in Europe. ICCT White Paper
- Transportation Research Board, National Research Council. 2002. Effectiveness and impact of corporate average fuel economy (CAFE) standards. Washington, DC: The National Academies Press
- U.S. Energy Information Administration. 2015. Vehicle standards around the world aim to improve fuel economy and reduce emissions. Today In Energy, available at: <http://www.eia.gov/todayinenergy/detail.cfm?id=23572src=email>
- U.S. Environmental Protection Agency. 2010. Final rulemaking to establish light-duty vehicle greenhouse gas emission standards and corporate average fuel economy standards, regulatory impact analysis. EPA-420-R-10-009
- Weitzman ML. 1974. Prices vs. quantities. *The Review of Economic Studies* 41:477–491
- West J, Hoekstra M, Meer J, Puller SL. 2015. Vehicle miles (not) traveled: Why fuel economy requirements don’t increase household driving. NBER Working Paper 21194
- Whitefoot K, Fowle M, Skerlos S. 2013. Evaluating U.S. reformed corporate average fuel economy standards using an engineering attribute selection model. Manuscript: University of California, Berkeley
- Whitefoot K, Skerlos S. 2012. Design incentives to increase vehicle size created from the U.S. footprint-based fuel economy standards. *Energy Policy* 41:402–411

A. Alternative model with fuel economy given as a function of car attributes and technology

In section 2 above, we model a consumer's choice of car attributes and fuel economy, with the car's cost given as a function of car attributes and fuel economy. In this section, we model a consumer's choice of car attributes and car cost (i.e., technology), with the car's fuel economy given as a function of car attributes and costs. By reframing the model in this way, we are able to show that consumer undervaluation of fuel economy will contaminate the tradeoff between car attributes and fuel economy (in addition to the tradeoff between car technology and fuel economy that is clear in the original formulation of our model).

A.1. Consumer welfare

Following the setup of our model in the main text above, we assume that the representative consumer has utility that is quasilinear in transportation services and other goods, valued in dollars, such that welfare is also valued in dollars and is given by the following expression:

$$W(c, x, v, n) = \mu(n)\theta(x)u(v) + y - nc - n[p + d]g(x, c)v, \quad (10)$$

where: c is the cost of the car, which indexes the car's technology level; $g(x, c)$ is the average car's per-mile fuel consumption as a function of attributes x and cost c ; and all other variables are as above. Thus, $g(x, c)$ traces out an iso-cost curve in $g - x$ space. We assume partial derivatives given by $g_x > 0$, $g_{xx} \geq 0$, $g_c < 0$, $g_{cc} \geq 0$, and $g_{xc} \geq 0$. That is, the per-mile fuel consumption rate is increasing and convex in attributes, decreasing and concave in costs, and a dollar of technology is less effective in reducing per-mile fuel consumption in cars with greater attributes.

The necessary conditions that characterize welfare maximization are as follows:

$$\frac{\partial W}{\partial c} = -n - n[p + d]g_c(x, c)v = 0 \quad (11)$$

$$\frac{\partial W}{\partial x} = \mu(n)\theta'(x)u(v) - n[p + d]g_x(x, c)v = 0 \quad (12)$$

$$\frac{\partial W}{\partial v} = \mu(n)\theta(x)u'(v) - n[p + d]g(x, c) = 0 \quad (13)$$

$$\frac{\partial W}{\partial n} = \mu'(n)\theta(x)u(v) - c - [p + d]g(x, c)v = 0 \quad (14)$$

Expression (11) says that, holding car attributes fixed, fuel-saving technology should be added to cars until the last dollar spent on technology reduces the social marginal costs of fuel consumption by exactly one dollar. Expression (12) says that, holding car technology fixed, the marginal utility of better car attributes (e.g., size) should be set equal to the corresponding marginal impact on the social cost of fuel consumption that results from less efficient cars. The other expressions have similar interpretations as above. Denote the solution to these conditions by (c^*, x^*, v^*, n^*) .

A.2. Externalities and internalities

Consumer choices in general will differ from the welfare-maximizing conditions for the same reasons as above. The marginal external damage from fuel consumption is again d , while the degree of consumer undervaluation of future fuel expenditures is again given by β .

A.3. Policy options

We consider the same set of policies as outlined above.

A.4. Consumer behavior

Given these assumptions, the necessary conditions that characterize decentralized consumer behavior are given as follows:

$$\frac{\partial B}{\partial c} = -n - \beta n[p + \tau_f]g_c(x, c)v - \tau_g n g_c(x, c) = 0 \quad (15)$$

$$\frac{\partial B}{\partial x} = \mu(n)\theta'(x)u(v) - \beta n[p + \tau_f]g_x(x, c)v - \tau_g n [g_x(x, c) - \sigma'(x)] = 0 \quad (16)$$

$$\frac{\partial B}{\partial v} = \mu(n)\theta(x)u'(v) - n[p + \tau_f]g(x, c) = 0 \quad (17)$$

$$\frac{\partial B}{\partial n} = \mu'(n)\theta(x)u(v) - c - \beta[p + \tau_f]g(x, c)v - \tau_g [g(x, c) - \sigma(x)] - \tau_n = 0, \quad (18)$$

where ∂B denotes the consumer's perceived change in private welfare associated with marginal changes in the respective choice variables. Note that the presence of the consumer undervaluation factor β directly impacts choices related to car technology, car attributes, and the number of cars purchased in conditions (15), (16), and (18). Denote the solution to these conditions by $(\hat{c}, \hat{x}, \hat{v}, \hat{n})$.

From here, it is clear in conditions (15) and (16) that both the fuel tax and the fuel-economy standard will generate incentives for consumers to add fuel-saving technologies (c) and, holding technology fixed, scale back car attributes (x) to save fuel.