Gasoline Price Uncertainty and the Design of Fuel Economy Standards

Ryan Kellogg *

January 22, 2018

Abstract

What are the implications of gasoline price volatility for the design of fuel economy policies? I show that this problem has a strong parallel to Weitzman's (1974) classic model of using price or quantity controls to regulate an externality. Changes in fuel prices act as shocks to the marginal cost of complying with the standard. Assuming constant marginal damages from fuel consumption, an application of Weitzman (1974) implies that a fixed fuel economy standard reduces expected welfare relative to a "price" policy such as a feebate or, equivalently, a fuel economy standard that is indexed to the price of gasoline. When the regulator is constrained to use a fixed standard, I show that the usual approach to setting the standard—equate expected marginal compliance cost to marginal damage—is likely to be sub-optimal because the standard may not bind if the realized gasoline price is sufficiently high. Instead, the optimal fixed standard will be relatively relaxed and may be non-binding even at the expected gasoline price. Finally, I show that although an attribute-based standard allows vehicle choices to flexibly respond to gasoline price shocks, the resulting distortions imply that the optimal fuel economy standard is not attribute-based.

^{*}Harris School of Public Policy, University of Chicago, and NBER. Email: kelloggr@uchicago.edu. This paper has benefited from comments and suggestions from two anonymous referees and from conference and seminar audiences at the AERE Summer Conference, Energy Institute at Haas Energy Camp, Heartland Environmental and Resource Economics Workshop, Midwest Energy Fest, NBER EEE Summer Institute, PERC Searle Workshop, TE³, UC Santa Barbara, and University of Michigan. I especially thank Lucas Davis, Jing Li, Ian Parry, and James Stock for helpful discussions.

1 Introduction

In October 2012, the U.S. Environmental Protection Agency (EPA) and the National Highway Transportation and Safety Administration (NHTSA) jointly issued new fuel economy standards for U.S. light-duty vehicles. These standards are scheduled to become increasingly stringent over time, reaching a fleet-wide average fuel economy of 55 miles per gallon in 2025. When the standards were issued, the EPA and NHTSA scheduled a "mid-term review" in 2017–2018 to determine whether the standards for 2022–2025 should be revised. Thus, the standards set in 2012 were locked in place through 2021—nine years after the standards were set—and any revisions made in 2017–2018 will be locked in for seven years. The EPA concluded the mid-term review in January, 2017, affirming the original 2022–2025 standards. Two months later, however, the new administration opened a review of that decision. This new review is scheduled to conclude in April, 2018.

Throughout the debate over the 2022–2025 standards, there has been discussion among industry participants and regulators of whether the standards should be loosened given the substantial decrease in gasoline prices since the standards were set. In October 2012, the prevailing U.S. average gasoline price was \$3.45/gallon, but by January 2016 the price had fallen to \$1.51/gallon.¹ Because consumers' willingness to pay for fuel-efficient vehicles depends on fuel prices (Allcott and Wozny 2014; Busse, Knittel, and Zettelmeyer 2013; Sallee, West, and Fan 2016), this decrease in the price of gasoline has increased the cost of complying with the new fuel economy standards and therefore raised concerns as to whether the regulatory path set in 2012 should be sustained. For instance, in April 2015, a spokesperson for the Alliance of Automobile Manufacturers said that "Given the extremely long 15-year lead time for the standards, the government set a mid-term review in 2017 as a reality check for regulatory assumptions. One of the assumptions was that the price of gas would be much higher than today, and that affects what consumers buy. Sales of our most fuel-efficient vehicles go up

¹These prices are pre-tax and real \$2016. See section 2.2 for a discussion of the gasoline price data used in this paper.

and down with the price of gasoline. A mid-term reality check is a good idea, especially since our compliance is based on what consumers buy, not what we offer for sale." (Detroit News 2015).

Spurred by this policy question, this paper examines the welfare economics of fuel economy standards under uncertain future gasoline prices and therefore uncertain future compliance costs. I begin by framing the problem in the context of Weitzman's (1974) "Prices vs Quantities", which facilitates a welfare comparison of a standard that varies with the gasoline price versus a standard that is fixed in place. I present a model in which the private benefits of fuel economy that accrue to vehicle producers and consumers are moderated by the price of gasoline so that, absent regulation, private agents will choose fuel efficient vehicles when gasoline prices are high and inefficient vehicles when gasoline prices are low. Fuel economy policy is motivated in the model by climate change externalities that are associated with vehicles' fuel consumption per mile. In this framework, a fuel economy standard that varies with the price of gasoline—allowing greater fuel use per mile when the gasoline price falls but tightening when the gasoline price rises—can be equivalent to a revenue-neutral "feebate" that taxes inefficient vehicles and subsidizes efficient vehicles (since under a feebate agents' fuel economy choices would also rise and fall with the gasoline price). Thus, in the language of Weitzman (1974), a gasoline price-indexed fuel economy standard acts as a "price" policy. In contrast, a traditional fixed standard acts as a "quantity" policy.

To distill intuition, I simplify the model so that, in the absence of uncertainty over future gasoline prices, a feebate policy can achieve the first-best welfare outcome by acting effectively as a tax on gasoline. Likewise, a fuel economy standard can achieve the firstbest by acting as a standard on the quantity of gasoline consumed. The key assumption underpinning the paper is that the marginal external damage associated with vehicles' fuel use per mile is locally constant and unaffected by gasoline price shocks. Treating marginal damage as constant in U.S. vehicle emissions is a natural consequence of the fact that CO_2 — the primary driver of the externality—is a global stock pollutant (Newell and Pizer 2003).²

The paper's first result has been intuited by others (Anderson and Sallee 2016) and is a direct application of Weitzman (1974): indexing the fuel economy standard to the price of gasoline—or, equivalently, using a revenue-neutral feebate—can achieve the firstbest with uncertain fuel prices because it can equate the marginal cost of abatement with its marginal benefit at any fuel price. This result can also be seen as an application of models of general indexed regulation (Ellerman and Wang 2003; Newell and Pizer 2008), where the index variable (the gasoline price) is perfectly correlated with marginal compliance cost. In contrast, a fixed fuel economy standard will result in an excessive marginal abatement cost (and too much abatement) when fuel prices are lower than what was expected when the policy was set, and too small a marginal abatement cost (and too little abatement) when realized fuel prices are high. In either case, there is a welfare loss relative to the first-best.

I focus the remainder of the paper on the question of how best to set a fuel economy standard when it cannot be indexed to the gasoline price and cannot be changed for many years once it is set. Following a first order condition (FOC) from Weitzman (1974), the usual rule for setting such a fixed standard is to equate the expected marginal cost of compliance, evaluated at the standard, to the marginal external harm. I show, however, that this policy rule ignores the non-trivial probability that if the realized gasoline price is sufficiently high, the standard will not bind. That is, firms and consumers may, if the gasoline price is high enough, voluntarily select a level of fuel economy that exceeds the standard. I derive a new first order condition that accounts for this possibility and show that the optimal fixed standard equates expected marginal compliance cost, *conditional on the standard binding*, to marginal external harm. This optimal standard is more lax than that implied by the usual Weitzman (1974) rule, and I show that given historic gasoline price uncertainty and estimates of marginal damage from the literature, the difference between the optimal standard and the

²Even if marginal damage were not locally constant, setting a tax or feebate schedule would achieve the first-best and welfare dominate a fixed standard in the presence of fuel price uncertainty, per arguments given in Kaplow and Shavell (2002).

usual Weitzman (1974) standard can be economically large. In fact, the optimal standard may be so lax that it is non-binding even at the expected future gasoline price. I also derive an expression for the expected welfare loss (relative to the first-best) under a fixed and possibly non-binding standard, and I show that a standard set using the usual Weitzman (1974) FOC may yield lower welfare than that obtained by not setting any standard at all.

Brozovic, Sunding, and Zilberman (2004) also studies optimal pollution control standards that may not bind, using a model with two discrete firm types. The results presented here expand on this prior work by presenting intuitive formulas for the optimal standard and expected welfare loss under a continuous distribution of abatement cost uncertainty, and by calibrating the model to demonstrate that its implications are economically significant for fuel economy regulation. Another related paper is Costello and Karp (2004), which examines the merits of setting a potentially non-binding pollution standard to learn about a firm's marginal abatement cost.³

Finally, I study whether basing the fuel economy standard on vehicle attributes such as footprint or weight can mitigate the welfare losses associated with a fixed standard. The U.S. fuel economy standards set in 2012 are in fact footprint-based: vehicles with a large wheelbase are assigned a less stringent standard, as shown in figure 1. Ito and Sallee (forthcoming) shows, using a model similar to that presented here but assuming complete information, that attribute-based regulation (ABR) reduces welfare by distorting choices of the attribute: consumers purchase vehicles that are larger than optimal.⁴ When the gasoline price is uncertain, however, ABR may confer benefits by building flexibility into the regulation, as suggested in Anderson and Sallee (2016). For instance, if the gasoline price is lower than expected, agents can shift to vehicles with larger footprints and lower fuel economy, mitigating the increase in the marginal abatement cost. As noted by Lutsey (2015), this flexibility benefit of ABR was highlighted in General Motors' public comments

 $^{^{3}}$ Costello and Karp's (2004) model is structured so that, in the absence of an incentive to learn, the optimal standard always binds.

⁴Whitefoot and Skerlos (2012) and Gillingham (2013) also illustrate the incentive to increase vehicle size under ABR, using models that allow for automakers' market power.

(2009) on the proposed regulation. In addition, the response of vehicle choices to fuel prices was modeled by NHTSA in its Regulatory Impact Analysis (2012), and Leard, Linn, and McConnell (2016) finds evidence that the recent fuel price decrease has indeed modestly affected fleet-average footprint and fuel economy. I show here, however, that even if fuel prices are uncertain, attribute-basing reduces expected welfare relative to an optimallyset non-attribute-based standard, since the distortions caused to the attribute and to fuel economy outweigh the flexibility benefit.

The paper proceeds as follows. Section 2 describes U.S. fuel economy standards and characterizes U.S. gasoline price volatility, and section 3 introduces the paper's model of the vehicle market. Section 4 then applies Weitzman (1974) to the comparison of fixed versus gasoline price-indexed fuel economy standards. Sections 5 and 6 then present the main results of the paper: section 5 discusses the optimal level for a fixed fuel economy standard, and section 6 assesses whether an attribute-based standard can be welfare-improving. Section 7 discusses the potential importance of several non-modeled factors, such as covariance between the gasoline price and marginal damage, banking and borrowing of fuel economy credits, under-valuation of fuel economy by consumers, and firms' investment dynamics. Section 8 concludes.

2 Policy background and gasoline price data

2.1 2012 EPA and NHTSA fuel economy standards

This section discusses the U.S. fuel economy standards set by the EPA and NHTSA in 2012, focusing on three of their characteristics that are important for this paper: attribute-basing, credit trading, and banking and borrowing. For additional information, see EPA (2012) or the complete rule in the Federal Register (2012). Figure 1 shows the planned path of the footprint-based fuel economy standard for passenger cars through 2025 (though the standards for 2022–2025 are subject to revision in the ongoing EPA review). For any given footprint,



Figure 1: U.S. EPA and NHTSA fuel economy standards for passenger cars through 2025

Source: Federal Register, Vol 77, No. 199, p. 62644 (15 Oct, 2012). For reference, the 2012 model year Honda Fit, Ford Fusion, and Chrysler 300 have footprints of 40, 46, and 53 square feet, respectively.

the required level of fuel economy increases over time as technology is projected to advance. In addition (and not shown), light trucks are assigned a separate set of footprint-based standards that are less stringent than those for passenger cars. This separate treatment of passenger cars versus light trucks is essentially a different form of attribute-basing, though on a discrete rather than continuous variable.

Individual vehicle models are not required to meet their fuel economy target shown in figure 1. Instead, each manufacturer must meet a fleet-wide standard based on the aggregated compliance of all vehicles sold, including both passenger cars and light trucks.⁵ Moreover, manufacturers can trade compliance credits with one another so that the entire industry faces an identical marginal cost of compliance. Manufacturers may also engage in limited banking and borrowing of fuel economy credits: unused credits may be banked for up to five years

⁵Any given model may exceed or fall short of its footprint-based standard, yielding the model's manufacturer fuel economy credits or liabilities, respectively, proportional to the difference between the model's fuel use per mile and the model's footprint-based standard.

or used to cover deficits no more than three years in the past. Ultimately, each year every manufacturer must either produce vehicles that in aggregate comply with its fuel economy target or cover any deficit using credits obtained via trading, banking, or borrowing.⁶

2.2 Gasoline price variation

To motivate the problem of setting fuel economy policy under fuel price uncertainty, this section establishes that U.S. gasoline price shocks have historically been both large and persistent. I use data on gasoline prices obtained from the U.S. Energy Information Administration (EIA). Panel (a) of figure 2 illustrates the time path of U.S. gasoline prices before and after the current fuel economy standards were set in October 2012. The prices shown are monthly U.S. volume-weighted average refiner prices of motor gasoline for sale to end users, pre-tax, and deflated to January 2016 U.S. dollars using the Bureau of Labor Statistics' Consumer Price Index for all goods less energy.⁷ The price of gasoline has exhibited substantial volatility, ranging from roughly \$1.00/gallon to \$4.00/gallon since 2000.

I use the historical volatility of gasoline prices to generate estimates of the variance of future gasoline prices that policy-makers face when setting fuel economy policy with a long time horizon. Because the distribution of gasoline prices is right-skewed (panel (b) of figure 2), I assume that monthly price innovations are additive in logs, so that the future price distribution is lognormal. I therefore estimate historic gasoline price volatility by taking the standard deviation of differenced log prices. I use long differences to capture gasoline price uncertainty over time horizons that are similar to fuel economy policy lifetimes. For instance, the 2012 standard is locked-in until 2021—nearly a decade—and may be in effect through 2025 pending the outcome of the current EPA review. The volatility of ten-year differenced

⁶In the original U.S. Corporate Average Fuel Economy(CAFE) standard established in 1975, manufacturers could also comply by paying fines. Fines are no longer a compliance mechanism under the current EPA standards. Also, credit trading across manufacturers was not permitted under the original CAFE standard.

⁷Gasoline price data were accessed from https://www.eia.gov/dnav/pet/pet_pri_refoth_dcu_nus_m.htm on 11 April, 2016. The CPI used covers all urban consumers and is not seasonally adjusted. The CPI series ID is CUUR0000SA0LE.



Figure 2: U.S. gasoline price variation

Note: Gasoline price data are monthly U.S. volume-weighted average refiner prices of motor gasoline for sale to end users, obtained from the U.S. Energy Information Administration. These prices are pre-tax and are deflated to January 2016 U.S. dollars using the Bureau of Labor Statistic's Consumer Price Index for all goods less energy (CPI series CUUR0000SA0LE).

log gasoline prices, estimated for January 2000 to January 2016,⁸ equals 0.42, equivalent to 52%.⁹ The 95% confidence interval for gasoline prices ten years in the future therefore encompasses prices more than twice or less than half of the expected future price. Volatilities estimated using shorter differences naturally result in less uncertainty. The volatility of one-month differences, for instance, is 0.085, and the volatility of one-year differences is 0.272.¹⁰

Finally, the data indicate that gasoline price shocks are highly persistent. The monthly AR1 coefficient on the logged gasoline price since 2000 is 0.97, and a unit root cannot be

⁸Gasoline price data are available since January 1990, so January 2000 is the first month for which ten-year lagged prices are available.

⁹Specifically, 0.42 is obtained by taking the standard deviation of the ten-year difference of the logged gasoline price over January 2000 to January 2016. 52% is obtained by $e^{0.42} - 1$. The 95% confidence interval for future prices is obtained by $e^{0.42 \cdot 1.96}$ and its reciprocal, yielding a range from 0.44 to 2.28 times the expected price.

¹⁰Whenever I estimate volatility for differences of less than 10 years, I use the same gasoline price sample used to estimate the ten-year volatility; i.e., a sample from January 2000 to January 2016. Note that as the lag length increases, volatility increases at a rate less than the \sqrt{t} rule that would hold for a geometric random walk (e.g., $0.085 \cdot \sqrt{120} = 0.93 > 0.42$), suggesting that price innovations in levels do not scale fully proportionally with the current price level.

rejected.¹¹ This persistence is consistent with results on the relative accuracy of no-change forecasts of the long-run real price of oil (Alquist, Kilian, and Vigfusson 2013) and relates to evidence that consumers hold no-change beliefs about future gasoline prices (Anderson, Kellogg, and Sallee 2013).

3 Model set up

This section introduces the model of the vehicle market that I use throughout the paper. Like the model used in Ito and Sallee (forthcoming), I impose a number of simplifying assumptions such that, under complete information, either a fuel economy standard or a feebate can attain the first-best because they can mimic a quantity standard or a Pigouvian tax on gasoline. In particular, I assume that consumers have unit demand for vehicles, fully internalize future fuel costs when purchasing a vehicle, and do not respond to per-mile fuel costs when selecting how many miles to drive (i.e., there is no "rebound effect"). I treat vehicle manufacturers as perfectly competitive, and I abstract away from interactions with the used car market and from interactions with other taxes or un-priced externalities.

These assumptions are all strong, and their implications for the efficiency of fuel economy standards are discussed in Ito and Sallee (forthcoming) and many other works. However, abstracting away from these issues conveys the substantial benefit of being able to obtain analytical results and thereby distill intuition for the first-order implications of fuel price uncertainty for fuel economy policy. Moreover, the fact that the model is not precisely tailored to the features of the vehicle market facilitates linking the model to Weitzman (1974) and enhances the generalizability of the paper's findings to other pollution control settings in the presence of uncertain marginal abatement costs. In section 7, I return to several of the above assumptions, as well as other factors that are not explicitly modeled,

¹¹Using the augmented Dickey-Fuller test of Elliot, Rothenberg, and Stock (1996), I obtain a test statistic of -1.24 at the Ng and Perron (2001) optimal lag length of 10, relative to a 10% critical value of -2.567. Similar results are obtained if the AR1 regression and unit root tests are carried out in levels rather than logs, or if the full dataset back to 1990 is used.

and discuss how relaxing them may affect the paper's conclusions. And in appendix C, I show that the "prices versus quantities" result holds when miles driven respond to per-mile fuel costs, so long as the gasoline price elasticity of driving is not too large.

3.1 Private demand and supply of vehicles

I begin with the consumer side of the market. Consider a consumer *i* purchasing a vehicle with fuel consumption per mile e_i and attribute a_i .¹² a_i can represent a single attribute that can ultimately be targeted by the regulation (i.e., the vehicle's footprint), or more abstractly a scalar composite of multiple attributes. The price of gasoline at the time of purchase is given by *G*. I then denote consumer *i*'s indirect utility from this vehicle purchase by $U^i(e_i, a_i, G) - P_i$, where P_i is the price of the vehicle and $U^i(e_i, a_i, G)$ is a twice continuously differentiable function. This specification implicitly assumes that consumer utility is quasilinear, with all other goods as a numeraire.

Increases in e_i reduce utility because increased fuel consumption per mile implies increased expenditure on fuel; thus, the derivative $U_e^i < 0$. Because each consumer's vehicle miles travelled is exogenous and utility is quasi-linear, the second derivative U_{ee}^i equals zero.¹³ The disutility from a vehicle's fuel use increases in magnitude with the price of gasoline at the time of purchase (since gasoline prices are persistent), so that $U_{eG}^i < 0$. U_{eG}^i , moreover, represents the consumer's expected discounted miles travelled over the lifetime of the vehicle. The attribute *a* is assumed to be desirable but with diminishing returns, so that $U_a^i > 0$ and $U_{aa}^i < 0$. Quasi-linearity of utility implies that $U_{aG}^i = 0$.

On the supply side of the market, vehicle manufacturers competitively produce vehicles with a cost function $C(e_i, a_i)$. Reducing the fuel use of any particular vehicle is costly, so that $C_e < 0$, and there are diminishing returns to fuel economy investments, so that $C_{ee} > 0$.

¹²It is convenient to model e_i as fuel use rather than efficiency $(= 1/e_i)$ because environmental damage is (locally) linear in fuel use but not efficiency.

 $^{{}^{13}}U_{ee}^i = 0$ also implies that consumers only value fuel efficiency for its monetary benefits, not because of green preferences. To the extent that consumers have green preferences, fuel efficiency could also be viewed as part of the characteristic *a* that consumers value with diminishing returns.

Increasing the attribute *a* involves convex costs, so $C_a > 0$ and $C_{aa} > 0$. Finally, letting *a* represent attributes such as footprint, weight, or horsepower, reductions in *e* should be weakly more costly for vehicles that have a large value of *a*. Thus, I assume $C_{ea} \leq 0$.

As in Weitzman (1974) and Ito and Sallee (forthcoming), I assume that the $U^i(e_i, a_i, G)$ and $C(e_i, a_i)$ functions can be approximated by second-order Taylor expansions and that all of the U^i have identical second derivatives, which substantially improves the model's tractability.¹⁴ Under these assumptions, appendix A shows that a sufficient statistic for measuring impacts of fuel economy policy on aggregate (utilitarian) consumer welfare is the representative consumer's utility U(e, a, G) obtained from the average vehicle purchased, where e is the average of the e_i , a is the average of the a_i , and the first derivatives of U equal the average of the U_e^i and U_a^i . Likewise, the cost C(e, a) of the average vehicle purchased is a sufficient statistic for aggregate production cost.¹⁵

Combining the two sides of the market, define $B(e, a, G) \equiv U(e, a, G) - C(e, a)$ as the private welfare benefit when the average vehicle purchased has fuel use per mile e and attribute a, and the gasoline price is G. B(e, a, G) inherits the property that it can be expressed as a second-order Taylor expansion, and its second derivatives have the properties $B_{ee} < 0, B_{aa} < 0, B_{eG} < 0, B_{ea} \ge 0$, and $B_{aG} = 0$. I also naturally assume that $B_e(0, a, G) >$ 0 and $B_a(e, 0, G) > 0$ to allow for an interior private optimum.

Absent regulation, private agents will choose vehicles so that e and a satisfy the private first order conditions (FOCs): $B_e(e, a, G) = 0$ and $B_a(e, a, G) = 0$. I assume that the second order condition (SOC), $B_{ee}B_{aa} - B_{ea}^2 > 0$, is satisfied, guaranteeing a unique, interior solution.

¹⁴A necessary condition for the assumption that the U_{eG}^i are identical is that all consumers share the same annual vehicle miles travelled.

¹⁵The sufficient statistic claims hold for the policies considered in this paper: fuel economy feebates and fuel economy standards that are implemented with cross-manufacturer credit trading, so that the marginal cost of abatement is equated across all vehicle models. In the absence of trading, so that a standard can bind against individual firms or even individual models, changes to the average vehicle sold are no longer sufficient for understanding welfare impacts.

3.2 Externalities from fuel consumption and the social optimum

Let ϕ_g denote the constant marginal damage resulting from consuming one gallon of gasoline.¹⁶ Then define $\phi \equiv -B_{eG}\phi_g$ as the marginal damage caused by an increase in a vehicle's fuel use per mile, discounted over the lifetime of the vehicle. The full-information social planner's problem is then given by:

$$\max_{e,a} B(e,a,G) - \phi e \tag{1}$$

The planner's FOCs are then given by $B_e(e, a, G) = \phi$ and $B_a(e, a, G) = 0$,¹⁷ and the SOC is identical to that in the private agents' problem. Intuitively, the planner's solution involves a lower e than the private optimum, but conditional on the choice of e the planner chooses the same a as private agents, since the attribute does not create an externality on its own.¹⁸

For the results in sections 4 through 6, it will be useful to discuss the externalities that are included in ϕ_g as well as estimates of their magnitudes. The median social cost of CO₂ used by the U.S. Government (Interagency Working Group 2013) for 2015 is \$38 per metric ton, with a range from \$12/ton to \$58/ton under alternative discount rate assumptions, and a value of \$109/ton under a "right-tail" damage assumption. These values translate to \$0.30/gallon, \$0.10/gallon, \$0.47/gallon, and \$0.87/gallon of gasoline, respectively.¹⁹ In addition to the climate externality, a survey by Parry, Wells, and Harrington (2007) considers externalities from U.S. dependency on foreign oil, concluding that "a corrective tax might be anything from roughly 8–50 cents (excluding geo-political costs); a recent panel of experts (National

¹⁶Per Weitzman (1974), the results of the paper are unchanged if ϕ_g is defined as the expected value of uncertain marginal damages, so long as the stochastic component of marginal damage is uncorrelated with G. In this case, what the paper describes as "first-best" is first-best subject to incomplete information about marginal damage. Section 7.1 discusses implications of a potentially non-zero correlation between marginal damage and G.

¹⁷I assume that $B_e(0, a, G) > \phi$, so that the planner's problem has an interior solution.

¹⁸I am abstracting away from effects of attributes such as footprint, size, and type (car vs. truck) on accident risk, which are examined empirically in Jacobsen (2013) and Anderson and Auffhammer (2014), and discussed in Ito and Sallee (forthcoming).

¹⁹Per EIA (2011), one gallon of E10 gasoline produces 17.68 pounds of CO_2 .

Research Council 2002) recommended a value of 12 cents per gallon". Combining these costs with the climate externality yields a range for ϕ_g from \$0.18/gallon to \$1.37/gallon; I focus below on a "median" estimate of \$0.42/gallon.²⁰ Finally, Parry et al. (2014) discusses externalities from gasoline consumption associated with local pollutants, congestion, and accident risk.²¹ Because these externalities are associated with miles driven and not fuel use per mile, I do not include them in ϕ_g for the purpose of regulating fuel economy.²²

4 Prices versus quantities

This section applies Weitzman (1974) to evaluate the welfare implications of a fixed fuel economy standard versus a feebate or, equivalently, a gasoline price-indexed standard. Because I do not consider an attribute-based standard until section 6, this section and section 5 will suppress the attribute a so that the private benefit function is simply given by B(e, G).²³ The model proceeds in two stages. In the first stage, the regulator must commit to a particular level for the standard or the feebate, without knowing what the realized gasoline price will be in the second stage. Then in stage 2 agents choose vehicles conditional on the regulation and the realized gasoline price. I model stage 2 as a single compliance period, though it is straightforward to extend the model so that stage 2 involves multiple compliance periods with different price realizations.

Letting G denote the realized gasoline price in stage 2, let F(G) denote the policy-maker's stage 1 rational belief about the distribution of G. F(G) has support on $[G_L, G_U]$, where $G_L \ge 0$ and G_U is possibly infinite. The expected stage 2 price is \overline{G} , and the variance of Gis σ^2 .

 $^{^{20}}$ \$0.42/gallon is the sum of \$0.30/gallon for climate and \$0.12/gallon for foreign oil dependency.

²¹Parry et al.'s (2014) estimates for the United States are \$0.09/gallon for local pollutants, \$0.85/gallon for congestion, and \$0.37/gallon for accidents, all in 2010 U.S. dollars.

²²Fuel economy improvements will not affect emissions of local pollutants in the U.S. because these pollutants are regulated under separate per-mile emission standards. Fuel economy regulation will only affect congestion and accidents indirectly via the rebound effect and impacts on vehicles' size and weight.

²³Under non-attribute based regulation, agents will always choose the socially optimal (= privately optimal) level of a conditional on the regulated choice of e.

Suppose that a tax is imposed on each vehicle in proportion to its fuel use per mile. Letting τ denote the tax rate, the tax on the representative vehicle is then given by τe . If the tax revenues are redistributed lump-sum, this tax is equivalent to a revenue-neutral feebate that taxes inefficient vehicles and subsidizes efficient vehicles.²⁴ It is straightforward to see that by setting $\tau = \phi$, the regulator can achieve the full-information first-best: the feebate is a Pigouvian tax, so it causes agents to perfectly internalize the externalities of fuel consumption per mile at any gasoline price level. This feebate is equivalent to a fuel economy standard that is indexed to gasoline prices so that the targeted fleet-wide average fuel consumption rises (falls) when gasoline prices fall (rise). The feebate, however, involves a substantially lower informational requirement for the regulator than the indexed fuel economy standard: the latter requires information on ϕ and the parameters of B(e, G), while the former only requires information on ϕ .

Now suppose that the regulator instead sets a fixed fuel economy standard such that fleet-wide average fuel use per mile cannot exceed the value μ . The "usual" FOC for the solution to the regulator's problem is to equate the expected marginal compliance cost at μ to marginal damage, as discussed in Weitzman (1974):

$$E[B_e(\mu, G)] - \phi = 0,$$
 (2)

where $E[B_e(e,G)] = B_e(e,\bar{G})$ due to the linearity of B_e . I show in section 5 below that equation (2) does not yield the optimal standard if there is a possibility that the standard may not bind. Putting that possibility aside for now, the standard implied by equation (2) will only yield the first-best in the event that the realized gasoline price is \bar{G} . Generically, however, the standard $\tilde{\mu}$ (in gallons per mile) that solves (2) will either be too low (if $G < \bar{G}$) or too high (if $G > \bar{G}$) relative to the first best, resulting in a marginal abatement cost that is, respectively, either too high or too low. The resulting welfare loss is illustrated in figure

 $^{^{24}}$ The tax or subsidy on any given vehicle is proportional to the difference between that vehicle's fuel use and that of the average vehicle.

Figure 3: Welfare from a fixed fuel economy standard that always binds



Note: e denotes fuel consumed per mile, and $B_e(e, G)$ is its marginal benefit. ϕ denotes marginal damage. Gasoline prices are distributed on $\{G_L, G_U\}$. $\tilde{\mu}$ is the optimal fixed standard. See text for details.

3 for the simple case where f(G) has two point masses at G_L and G_U . Given a standard $\tilde{\mu}$ set to satisfy FOC (2), and under the Taylor series approximation for B(e, G), the expected welfare loss per vehicle Δ is equal to a Harberger triangle with height $-\sigma B_{eG}$ (the standard deviation of marginal abatement cost) and width $\sigma B_{eG}/B_{ee}$. Thus, $\Delta = \frac{-\sigma^2 B_{eG}^2}{2B_{ee}}$. This result is a direct analogue of the Δ from Weitzman (1974) in the case where marginal damages are constant, and its formal derivation is given in appendix B.1.

In contrast, a feebate with $\tau = \phi$ allows the chosen level of fuel use e to shift with the gasoline price, so that the marginal abatement cost $B_e(e, G)$ is constant and always equal to ϕ , and the first-best is realized. An appropriately-set gasoline price-indexed standard can achieve the same end.

The price of gasoline during the compliance period is the only source of uncertainty in this paper's model. In principle, it is straightforward to extend the model to also allow for uncertainty about fuel economy technology, since shocks to the intercept of the marginal cost function $C_e(e)$ are isomorphic to gasoline price shocks.²⁵ Including such shocks in the model would increase the variance of the marginal cost of complying with a fixed fuel economy standard beyond the variance caused by gasoline price volatility alone. Moreover, conditioning the fuel economy standard solely on the price of gasoline would no longer be sufficient to achieve the full-information first-best. Instead, the standard would also have to be conditioned on the state of fuel economy technology, which may present practical challenges since technological evolution is more challenging to objectively measure than the price of gasoline.

5 The optimal level of a fixed standard

Suppose that a feebate or gasoline price-indexed fuel economy standard is not a tool in the regulator's choice set. Instead, the regulator must set a fixed fuel economy standard, knowing that whatever standard it sets will be locked in place for many years. These restrictions are strong, but they are reflective of the way fuel economy policy has been implemented in the United States. As discussed above, the usual, intuitive way to set the standard in this situation, following Weitzman (1974) and expressed in equation (2), is to equate the expected marginal compliance cost at the standard to marginal damage. Doing so is equivalent to adding marginal damage per gallon ϕ_g to the expected future gasoline price \bar{G} and setting the standard at the level of fuel consumption per mile that private agents would voluntarily choose at that price. Denote this standard, which solves (2), by $\tilde{\mu}$.

This approach to setting the standard is incorrect, however, when uncertainty about marginal compliance cost is sufficiently large that a standard set at $\tilde{\mu}$ may not bind. The

²⁵To see this isomorphism, observe that shocks to the intercept of $C_e(e)$ (e.g., via a stochastic technology variable η that vertically shifts the affine $C_e(e)$ function) result in an identical but oppositely-signed vertical shift to the net marginal benefit function $B_e(e)$ (holding G constant), since $B_e(e, G) = U_e(e, G) - C_e(e)$. Gasoline price shocks are isomorphic to these technology shocks because they also shift $B_e(e)$ vertically (scaled by B_{eG}). More generalized shocks to fuel economy technology may change the slope of $C_e(e)$: such shocks would not be isomorphic to gasoline price shocks (because they would change the value of B_{ee}) and would take the model outside of the Weitzman (1974) framework.

problem lies in that fact that, for a sufficiently negative shock to marginal compliance cost, private agents will on their own choose to consume less fuel than what is prescribed in the standard. To see that this possibility is non-trivial in the case of fuel economy regulation, it is useful to compare the variance of future gasoline prices, discussed in section 2.2, to estimates of ϕ_g , discussed in section 3.2. Looking ten years forward, the 95% confidence interval for the future price of gasoline encapsulates a doubling or a halving of the current price (under a no-change forecast for the expected future price). There exists, therefore, a substantial probability that the price of gasoline may increase by more than \$1.00/gallon over a ten-year policy horizon. Such an increase exceeds most estimates of ϕ_g . Thus, the realized future price of gasoline may exceed the sum of \bar{G} and ϕ_g , so that a policy set at $\tilde{\mu}$ using equation (2) may not bind. Given this possibility, what then is the ex-ante welfare maximizing level of the standard?

5.1 First order condition for an optimal fixed standard

To optimally set the level of the fuel economy standard, it is useful to write the regulator's problem in a way that makes clear that there is a critical gasoline price, $\hat{G}(\mu)$, such that for $G > \hat{G}(\mu)$ a given standard μ will not bind:

$$\max_{\mu} \int_{G_L}^{\hat{G}(\mu)} \left(B(\mu, G) - \phi \mu \right) f(G) dG + \int_{\hat{G}(\mu)}^{G_U} \left(B(e(G), G) - \phi e(G) \right) f(G) dG \tag{3}$$

where e(G) denotes agents' choice of e given a gasoline price realization G. Taking the derivative yields the first order necessary condition for an optimal standard:²⁶

$$\int_{G_L}^{\hat{G}(\mu)} (B_e(\mu, G) - \phi) f(G) dG = 0.$$
(4)

²⁶The terms in the derivative of (3) from Leibniz's Rule involving $d\hat{G}/d\mu$ cancel out. Under the assumption that $B_e(e, G)$ is continuously differentiable, this derivative is continuous if f(G) does not have any point masses. If f(G) has point masses (as in the case presented in figures 3 and 4, for example), then the derivative will discontinuously jump up as μ increases at each value of μ for which $\hat{G}(\mu)$ falls on a point mass. Because the derivative can only jump up but not down, the optimal standard μ^* will never have $\hat{G}(\mu^*)$ fall on a point mass, so FOC(5) is still a necessary condition for an optimal standard even when f(G) has point masses.

Dividing through by $F(\hat{G}(\mu))$ yields

$$\frac{1}{F(\hat{G}(\mu))} \int_{G_L}^{\hat{G}(\mu)} B_e(\mu, G) f(G) dG - \phi = 0.$$
(5)

The first term of FOC (5) is the expected marginal compliance cost conditional on the gasoline price being sufficiently low that the regulation binds. For any standard μ such that $\hat{G}(\mu) < G_U$, this conditional expected marginal compliance cost will be strictly greater than $E[B_e(\mu, G)]$ from equation (2), since $B_{eG} < 0$. Thus, if the $\tilde{\mu}$ standard that solves (2) may be non-binding, its conditional expected marginal compliance cost will strictly exceed ϕ , implying that $\tilde{\mu}$ is not optimal. Instead, because $B_{ee} < 0$, the optimal standard must be more lax—and therefore involve greater fuel use—than the $\tilde{\mu}$ standard.

This principle is illustrated in figure 4, which is similar to figure 3 in that the distribution of fuel prices is concentrated on two point masses at G_L and G_U . In figure 4, however, the standard $\tilde{\mu}$ that solves the Weitzman FOC (2) does not bind if $G = G_U$. If the standard is increased—i.e., shifted to the right—beyond $\tilde{\mu}$, observe that the welfare loss under the G_L price realization decreases while the welfare loss under the G_U price realization is unchanged. Thus, expected welfare can be improved by increasing the standard (in gallons per mile). The optimal standard is μ^* , where expected marginal compliance cost, conditional on the standard binding, equals marginal external harm ϕ , as prescribed by equation (5). For this simple two-point distribution of G, the welfare loss from the optimal standard is actually zero whenever $G = G_L$.

Under more general distributions of G, progress on characterizing the difference between $\tilde{\mu}$ and the optimal standard can be made by making use of the second order Taylor expansion for B(e, G). Because $B_e(\tilde{\mu}, \bar{G}) = \phi$, we may write:

$$B_{e}(\mu, G) = \phi + B_{eG}(G - \bar{G}) + B_{ee}(\mu - \tilde{\mu}).$$
(6)

Next, note that by the definition of \hat{G} , $B_e(\mu, \hat{G}(\mu)) = 0$. Combining this fact with

Figure 4: Welfare from a fixed fuel economy standard that does not bind for a high fuel price



Note: e denotes fuel consumed per mile, and $B_e(e, G)$ is its marginal benefit. ϕ denotes marginal damage. Gasoline prices are distributed on $\{G_L, G_U\}$. $\tilde{\mu}$ is the fixed standard that solves equation (2). μ^* is the optimal fixed standard that solves equation (4). See text for details.

equation (6) yields $B_{ee}(\mu - \tilde{\mu}) = -\phi - B_{eG}(\hat{G}(\mu) - \bar{G})$, so that $B_e(\mu, G) = B_{eG}(G - \hat{G}(\mu))$. We can use this result to replace $B_e(\mu, G)$ in equation (5) and then divide through by $-B_{eG}$ to obtain:

$$-\frac{1}{F(\hat{G}(\mu))} \int_{G_L}^{\hat{G}(\mu)} \left(\phi_g - (\hat{G}(\mu) - G)\right) f(G) dG = 0$$
(7)

$$\leftrightarrow \hat{G}(\mu) = E[G|G \le \hat{G}(\mu)] + \phi_g \tag{8}$$

Let μ^* denote the solution to FOC (8). The optimal standard μ^* has the property that the gasoline price at which the standard binds is equal to the expected gasoline price, conditional on the standard binding, plus the per-gallon marginal damage ϕ_g . In the event that the standard always binds (i.e., $\hat{G}(\mu^*) \geq G_U$), FOC (8) reduces to the usual condition that the optimal standard compels agents to choose vehicles that they would have purchased at a gasoline price equal to the unconditional expected gasoline price \bar{G} plus ϕ_q .

The value of equation (8) is that it allows a calculation of $\hat{G}(\mu^*)$ —the gasoline price at which the optimal fuel economy standard will bind—given only a value for ϕ_g and a distribution F(G).²⁷ Table 1 shows the results of solving equation (8) for $\hat{G}(\mu^*)$, using several values for both ϕ_g and the time horizon over which the standard will be enforced. Throughout, I assume that the expected price of gasoline, \bar{G} , is \$1.50/gallon (roughly the same price as that prevailing in January 2016)²⁸ and that F(G) is lognormal. For a given policy horizon, I obtain the standard deviation of F(G) by averaging gasoline price volatilities (calculated per the discussion in section 2.2) obtained from gasoline prices differenced over one month (capturing low uncertainty immediately after the policy is implemented) through prices differenced over the specified horizon length (capturing greater uncertainty at the end of the policy's compliance period).²⁹

The middle three rows of table 1 use a value for ϕ_g of \$0.42/gallon, the "median" externality estimate discussed in section 3.2. Under this estimate of ϕ_g , a fuel economy policy that will be committed to for one year—over which the average gasoline price volatility is 25%—has an optimal binding price $\hat{G}(\mu^*)$ of \$1.80/gallon rather than the \$1.92/gallon implied by the "usual" policy of setting $\hat{G}(\tilde{\mu}) = \bar{G} + \phi_g$. Policies implemented over five-year and ten-year horizons have optimal binding prices of only \$1.65 and \$1.60/gallon, respectively. Intuitively, because uncertainty about G increases with the policy horizon, the optimal stan-

²⁷Obtaining μ^* itself requires knowledge of B_{eG} and B_{ee} .

²⁸For higher values of \bar{G} , the variance of the future gasoline price increases, holding volatility (based on log prices) fixed. Because price variance enters equation (8) in levels, the difference between $\hat{G}(\mu^*)$ and $\hat{G}(\tilde{\mu})$ will increase with \bar{G} .

²⁹For instance, to obtain the one-year horizon volatility of 25% reported in table 1, I calculate the volatility of one-month, two-month, three-month, ..., twelve-month gasoline price differences, and I then take the average of these volatilities. In principle, the exact optimal gas price at which the standard binds should minimize expected welfare losses across the full time horizon over which the standard is applied, rather than be calculated from the average volatility over this time. I adopt the average volatility approach here both because it is more straightforward to implement and because it directly connects to the paper's analytical model, especially equation 8.

			Gas price \hat{G} at which		Expected welfare loss as %	
			standard binds (\$/gal)		of loss from no regulation	
Externality	Policy	Gas price	"Usual"	Optimal	"Usual"	Optimal
$\phi_g \; (\$/\text{gal})$	horizon	volatility	standard $\tilde{\mu}$	standard μ^*	standard $\tilde{\mu}$	standard μ^*
0.18	1 year	25%	1.68	1.34	181%	84%
	5 years	38%	1.68	1.13	312%	90%
	10 years	42%	1.68	1.06	367%	91%
0.42	1 year	25%	1.92	1.80	47%	43%
	5 years	38%	1.92	1.65	76%	56%
	10 years	42%	1.92	1.60	87%	59%
1.37	1 year	25%	2.87	2.87	6%	6%
	5 years	38%	2.87	2.84	12%	12%
	10 years	42%	2.87	2.83	14%	14%

Table 1: Effects of fixed fuel economy standards for various externalities and gasoline price volatilities. Expected future gasoline price \overline{G} is \$1.50 per gallon.

Note: Calculations assume that F(G) is lognormal. Volatilities of 25%, 38%, and 42% are estimated by averaging gasoline price volatilities over horizons from one month to 12 months, 60 months, and 120 months, respectively. The gasoline price $\hat{G}(\mu^*)$ at which the optimal fuel economy standard binds is calculated using equation (8). See text for details.

dard becomes progressively more relaxed as the horizon length increases.

Table 1 also presents optimal values for $\hat{G}(\mu^*)$ when the externality ϕ_g equals \$0.18/gallon or \$1.37/gallon, the lowest and highest plausible values from the estimates discussed in section 3.2. For $\phi_g =$ \$0.18/gallon, the reduction in $\hat{G}(\mu^*)$ at long policy horizons is pronounced: a ten-year policy has an optimal binding price of only \$1.06/gallon. Even for a one-year policy, the optimal standard is actually non-binding at at the expected gasoline price of \$1.50/gallon. In other words, the optimal standard specifies more gasoline consumption per mile than what private agents are expected to choose in the absence of regulation. In contrast, at $\phi_g =$ \$1.37/gallon the optimal policy is not substantially different from setting a standard that binds at $\bar{G} + \phi_g =$ \$2.87/gallon, even if the policy will be applied for ten years. This result stems from the fact that for such a large externality, it is unlikely that gasoline prices will be so high that the standard will not bind. Thus, table 1 overall illustrates how, given a distributional assumption on F(G), the difference between the optimal standard and the usual Weitzman (1974) calculation depends entirely on how large the volatility of future prices is relative to the externality, since these magnitudes govern the probability with which the standard will not bind and the size of the corresponding expected welfare loss.

It is important to emphasize that the optimal fixed standard μ^* is still inferior from a welfare perspective to a gasoline price-indexed standard or a feebate, as discussed in section 4. The μ^* standard derived here is optimal only subject to the constraint that it is not feasible for the regulator to make the standard contingent on the fuel price. The expected welfare loss Δ per vehicle for an optimal fixed standard μ^* that may not bind is derived in appendix B.1 and is given by

$$\Delta = -\frac{B_{eG}^2}{2B_{ee}} \left(F(\hat{G}(\mu^*))\sigma_c^2 + (1 - F(\hat{G}(\mu^*)))\phi_g^2 \right), \tag{9}$$

where σ_c denotes the standard deviation of fuel prices conditional on the standard binding.³⁰ Expression (9) indicates that the expected welfare loss is given by a weighted average of two Harberger triangles. First, with probability $F(\hat{G}(\mu^*))$ the standard will bind, and the expected height and width of the triangle will be proportional to σ_c . Second, with probability $1 - F(\hat{G}(\mu^*))$ the standard will not bind, and the height and width of the triangle will be proportional to the externality ϕ_g . Note that if the optimal standard binds for all realizations of G, then equation (9) reduces to the standard Weitzman (1974) formula, $\Delta = \frac{-\sigma^2 B_{cG}^2}{2B_{ce}}$.

One way to quantify the expected welfare losses Δ for the parameters used in table 1 is to express them as a percentage of the losses that would be realized in the absence of any regulation: $\frac{-\phi_g^2 B_{eG}^2}{2B_{ee}}$. This quantification does not require values for B_{eG} or B_{ee} . These percentage losses are displayed in the right-most column of table 1. For optimally-set standards, an intuitive pattern emerges that as gasoline price volatility increases relative to the externality ϕ_g , the percentage improvement in welfare generated by the standard is reduced. At the extremes, an optimal policy with a one-year horizon that addressed a \$1.37/gallon externality would eliminate 94% of the welfare loss from no regulation, while a policy with a ten-year horizon addressing a \$0.18/gallon externality would eliminate only 9% of the no-regulation welfare loss.

³⁰That is, $\sigma_c^2 \equiv \operatorname{Var}(\overline{G|G \leq \hat{G}(\mu)}).$

Table 1 also displays the welfare implications of setting a standard via the "usual" policy of $\hat{G}(\tilde{\mu}) = \bar{G} + \phi_g$. For the median externality estimate of \$0.42/gallon and a policy with a ten-year horizon, this expected welfare loss is 1.5 times that from the optimal standard.³¹ The relative welfare loss from adopting the $\tilde{\mu}$ policy rather than the optimal standard increases as fuel price volatility becomes large relative to the externality. In the low externality case of $\phi_g =$ \$0.18/gallon, the $\tilde{\mu}$ policy results in more than twice the welfare loss than that from the optimal policy—and a greater loss than that from not imposing a standard at all—even for a policy horizon of just one year.

Obtaining expected welfare losses in dollars per vehicle requires values for B_{eG} and B_{ee} . B_{eG} simply represents expected miles driven, discounted over the life of the vehicle. I follow the assumptions of Ito and Sallee (forthcoming) that new vehicles are driven 12,000 miles per year over a lifetime of 13 years, with a 5% discount rate, yielding $B_{eG} = 118,359$ miles. B_{ee} , the rate at which the marginal cost of reducing e (fuel consumption per mile) increases as edecreases, is more difficult to pin down. I use engineering estimates of the costs of fuel-saving technologies from National Research Council (2015), which estimates "pathways" by which fuel economy can be improved via sequential addition of fuel-saving technology to a baseline vehicle. Fitting an affine marginal cost curve to these estimates, I obtain a value for B_{ee} of \$1,756 per (gallon/100 miles)².³² Using these estimates, implementing the optimal standard for a ten year policy horizon results in expected welfare losses, relative to the first best, of \$12 per vehicle for $\phi_g =$ \$0.18/gallon, \$42 per vehicle for $\phi_g =$ \$0.42/gallon, and \$105 per vehicle for $\phi_g =$ \$1.37/gallon. Under the sub-optimal $\tilde{\mu}$ standard, the losses are \$48 per

 $^{^{31}\}text{The }1.5$ figure is the ratio of 87% versus 59% of the loss from no policy.

 $^{^{32}}$ I use the pathways presented in tables 8.4a and 8.4b of National Research Council (2015), which correspond to a midsize gasoline-fueled car under the NRC's low-cost and high-cost estimates, respectively. These pathways account for interactions between technologies as they are added to a vehicle. Using the cost estimates for 2017, I fit a line to the marginal cost curve for the technologies presented in each table. In each case, I drop the highest-cost technology (cooled exhaust gas recirculation level 2) because it is an extreme outlier (its MC is three times that of the next most costly technology) and because it is not required to meet the 2025 standards. I obtain estimates of B_{ee} equal to 1,474 using the low-cost estimates and 2,038 using the high-cost estimates. The average of these two values equals the \$1,756 per (gallon/100 miles)² used in the text. This value is similar to the \$1,941 estimated in Parry, Fischer, and Harrington (2004) using data from National Research Council (2002).

vehicle, \$61 per vehicle, and \$105 per vehicle, respectively.

5.2 Second order condition for an optimal fixed standard

The results shown in table 1 assumed that satisfaction of FOC (4) was sufficient for optimality.³³ It is important, however, to consider the second order condition (SOC), which is given by:

$$\frac{dG(\mu)}{d\mu} (B_e(\mu, \hat{G}(\mu)) - \phi) f(\hat{G}(\mu)) + B_{ee} F(\hat{G}(\mu)) < 0.$$
(10)

We can simplify this condition by using the facts that $d\hat{G}(\mu)/d\mu = -B_{ee}/B_{eG}$ and $B_e(\mu, \hat{G}(\mu)) = 0$, obtaining the SOC

$$B_{ee}F(\hat{G}(\mu))\left(1-\phi_g\frac{f(\hat{G}(\mu))}{F(\hat{G}(\mu))}\right)<0.$$
(11)

The term to the left of the parentheses in equation (11) is negative; thus, a solution μ^* to FOC (4) is a local maximum iff $\phi_g f(\hat{G}(\mu^*))/F(\hat{G}(\mu^*)) < 1$. When does this inequality hold? First, observe that a standard that never binds $(\hat{G}(\mu) \leq G_L)$ is always a solution to FOC (4). At this degenerate standard, however, $F(\hat{G}) = 0$, so the SOC clearly fails to hold. A standard that never binds is, therefore, never optimal (unless $\phi_g = 0$). Second, for a lognormal distribution, as well as many other distributions including the normal and uniform, f(G)/F(G) is decreasing in G (and strictly decreasing within the support of G), with a limit of zero as G goes to infinity. Thus, as μ decreases and \hat{G} increases, the SOC must eventually become negative and stay negative forever.³⁴ There can consequently be only one non-degenerate solution to FOC (4)—the solution obtained in section 5.1—and this solution will be the globally optimal standard.

For other distributions—particularly cases in which f(G) is left-tailed and marginal com-

³³In this section, I focus on FOC (4) rather than FOC (5) because the latter was obtained by dividing through by $F(\hat{G})$, which eliminates the fact that a degenerate standard satisfies the FOC.

³⁴This property continues to hold as \hat{G} increases beyond the support of G; i.e., when the standard always binds (as may be the case, for example, under a uniform distribution for G). Thus, the unique optimum could be a sometimes-binding or an always-binding standard.

pliance cost is therefore right-tailed—f(G)/F(G) can be non-monotonic so that the uniqueness of the non-degenerate solution to FOC (4) is not guaranteed. Because gasoline prices are right-tailed empirically (figure 2), this issue is likely not relevant for fuel economy regulation, but it may be important for other settings in which regulators set fixed standards and the distribution of uncertain marginal abatement cost is right-tailed. In such cases, solutions to (4) may be local minima or local maxima that are not globally optimal. Furthermore, appendix B.2 shows that a solution to (4) that is not the optimal standard will actually yield strictly lower expected welfare than the degenerate standard that never binds iff $\phi_g < \sigma_c$.³⁵ This result follows from the intuition that an "interior" standard generates welfare losses that increase with the conditional variance of gasoline prices, while the losses from a degenerate standard increase with the size of the unpriced externality, ϕ_q .

6 Uncertainty and attribute-based regulation

A fixed fuel economy standard, even if set optimally per section 5, will be too stringent relative to the first-best when the realized gasoline price is low and will be too lax (and may not bind) when the gasoline price is high. This section explores whether the consequent welfare loss can be mitigated by implementing an attribute-based regulation (ABR) that allows vehicle choices to flexibly respond to gasoline price shocks.

Under an affine ABR, the fuel economy standard is given by a function $\mu(a) = \mu_0 + \gamma a$. Before formally deriving the optimal standard—which involves optimal values for μ_0 and the slope parameter γ —it is useful to build intuition for agents' vehicle choices under ABR. Figure 5 depicts choices of a and e both without ($\gamma = 0$) and with ($\gamma > 0$) attribute-basing. In both panels of the figure, the steep dashed line depicts vehicle choices in the absence of regulation, when a higher (lower) gasoline price G decreases (increases) the choices of both

³⁵Appendix B.2 shows that this possibility extends to cases in which a standard satisfying (2) always binds. Such a standard will always be a local optimum (since the SOC for an always-binding standard is simply equal to B_{ee} , which is strictly negative) but will nonetheless be inferior to the degenerate standard iff $\phi_q < \sigma$, where σ is the unconditional standard deviation of gasoline prices (in levels).



Figure 5: Attribute-based fuel economy regulation

Note: e denotes fuel consumption per mile, and a denotes a vehicle attribute (e.g., footprint). When unconstrained, agents' choices lie on the dashed upward-sloping line, where a lower realized gasoline price G leads to higher choices of e and a. The fuel economy standard $\mu(a) = \mu_0 + \gamma a$ binds only for gasoline prices less than $\hat{G}(\mu_0, \gamma)$. See text for details.

a and e and therefore shifts the location on the line. The slope of this line equals $-B_{aa}/B_{ea}$. A first-best regulatory instrument such as a gasoline tax or feebate shifts agents' vehicle choices along this line towards lower values of a and e, for any given realization of G.

Panel (a) of figure 5 imposes a "flat" non-attribute-based standard of $\mu_0 = \mu^*$ that binds whenever G is less than $\hat{G}(\mu^*)$ and is optimal conditional on $\gamma = 0$, following the discussion in section 5. When the standard binds, the agents' choice of a always equals a^* as shown in the figure, since the optimal a for the agents, conditional on the choice of e, does not depend on the realized gasoline price $G.^{36}$ Equivalently, whenever $G \leq \hat{G}(\mu^*)$, the level set of agents' private welfare B(e, a, G) that corresponds to the private optimum, subject to the standard, always has a point of tangency at a^* .

Panel (b) of figure 5 sets $\gamma > 0$ so that the regulation is attribute-based. The standard is rotated so that the gasoline price $\hat{G}(\mu_0, \gamma)$ at which the standard binds is the same as that from panel (a), and therefore the "pivot point" (a^*, μ^*) is the same as (a^*, μ^*) from panel

³⁶This result follows from the assumption that $B_{aG} = 0$.

(a). Now, as the gasoline price decreases below $\hat{G}(\mu_0, \gamma)$, agents' choices (corresponding to the point of tangency of the level set) move along the standard to progressively higher levels of a and e.³⁷ It is in this way that ABR creates flexibility with which vehicle choices can respond to gasoline price shocks. The welfare impact of this flexibility is second-order in γ , since the sensitivity of e to the gasoline price G is second-order in γ in the neighborhood of $\gamma = 0.^{38}$

Panel (b) also illustrates that ABR distorts the choice of a, conditional on e, away from its optimum (which lies on the dashed line of unconstrained choices) whenever the standard is binding. This distortion generates a welfare loss, per Ito and Sallee (forthcoming), that is second-order in γ in the neighborhood of $\gamma = 0$ (since the welfare loss is a standard Harberger triangle).³⁹ Panel (b) also shows that, given the rotation of the standard about the point (a^*, μ^*) , ABR increases the choice of e, relative to the case when $\gamma = 0$, whenever the standard binds.

I now derive the optimal μ_0 and γ . The regulator's problem is given by:

$$\max_{\mu_{0},\gamma} \int_{G_{L}}^{\hat{G}(\mu_{0},\gamma)} \left(B(e(\mu_{0},\gamma,G),a(\mu_{0},\gamma,G),G) - \phi e(\mu_{0},\gamma,G) \right) f(G) dG + \int_{\hat{G}(\mu_{0},\gamma)}^{G_{U}} \left(B(e(G),a(G),G) - \phi e(G) \right) f(G) dG$$
(12)

The first order conditions are then given by (13) and (14) below, where the notation

 $^{3^{37}}$ I assume that private agents' SOC when faced with a binding attribute-based standard, $B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa} < 0$, holds, so that vehicle choices lie on an interior optimum.

 $^{{}^{38}}de/dG$, which drives the flexibility benefit, is a function of the product of the slope γ with the magnitude of the distortion in *a* away from the pivot point attribute a^* . Both of these effects are first-order in γ as γ increases away from zero, so de/dG is therefore second-order in γ in the neighborhood of $\gamma = 0$.

³⁹If the attribute *a* itself drives an externality, then the welfare effects of changing γ will be first-order at $\gamma = 0$, as discussed in Ito and Sallee (forthcoming). If increasing vehicle size and weight generates external benefits, then optimal ABR will involve $\gamma > 0$ rather than $\gamma = 0$ as derived below. Alternatively, if vehicle size and weight drive external harm, such as the accident damages found in Jacobsen (2013) and Anderson and Auffhammer (2014), then optimal ABR should actually discourage larger vehicles by setting $\gamma < 0$.

suppresses the dependence of B_e , B_a , $de/d\mu_0$, $de/d\gamma$, $da/d\mu_0$, and $da/d\gamma$ on μ_0 , γ , and G:

$$\operatorname{FOC}_{\mu_0} : \int_{G_L}^{\hat{G}(\mu_0,\gamma)} \left(B_e \frac{de}{d\mu_0} + B_a \frac{da}{d\mu_0} - \phi \frac{de}{d\mu_0} \right) f(G) dG = 0$$
(13)

$$FOC_{\gamma} : \int_{G_L}^{\tilde{G}(\mu_0,\gamma)} \left(B_e \frac{de}{d\gamma} + B_a \frac{da}{d\gamma} - \phi \frac{de}{d\gamma} \right) f(G) dG = 0$$
(14)

Addressing first FOC_{μ_0} , appendix B.3 shows that equation (13) simplifies to:

$$\frac{B_{eG}(B_{ea}\gamma + B_{aa})}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}} \int_{G_L}^{\hat{G}(\mu_0,\gamma)} \left(\phi_g - (\hat{G}(\mu_0,\gamma) - G)\right) f(G) dG = 0$$
(15)

Equation (15) is similar to FOC (7) for the level of the fuel economy standard under no attribute-basing, except that (15) is pre-multiplied by a term that is constant given γ . Intuitively, when $\gamma = 0$ this first term simply equals B_{eG} , so that equation (7) is a special case of equation (15).⁴⁰ Equation (15) therefore implies that, conditional on γ , μ_0 should be set so that the price \hat{G} at which the standard binds equals $\hat{G}(\mu^*, 0)$, the same price as for the optimal flat standard. In terms of figure 5, this result implies that the optimal pivot point (a^*, μ^*) for an attribute-based standard is invariant to the slope γ of the standard. Thus, as shown in figure 5, when γ increases, the expected fuel use e rises, so that the overall stringency of the standard effectively falls.

Why is the pivot point invariant to γ ? The intuition flows from the fact that, when $\gamma > 0$, the distortion to the attribute *a* increases with the stringency of the standard.⁴¹ Thus, as γ increases, the permitted fuel consumption per mile under the optimal standard must rise in order to mitigate the distortion to *a*. This outcome is accomplished by holding fixed the optimal pivot point (a^*, μ^*) .

⁴⁰Recall that equation (7) was derived by dividing the FOC through by $-B_{eG}F(\hat{G}(\mu))$, so the two FOCs are identical when $\gamma = 0$.

⁴¹To see the intuition, observe in figure 5(b) that reducing μ_0 shifts the choice of *a* away from the dashed line denoting agents' unconstrained (and optimal) attribute choices, conditional on *e*. Formally, the effect of changing μ_0 on the distortion to *a* is given by $da/d\mu_0 - (-B_{ea}/B_{aa})(1 + \gamma da/d\mu_0)$. Applying results from appendix B.3, this expression reduces to $\gamma(B_{ea}^2 - B_{aa}B_{ee})/(B_{aa}(B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}))$, which must be strictly negative (assuming the SOCs for the private agents' unconstrained and constrained problems are satisfied).

The fixed pivot point also has implications for the optimal level of γ . One way to view the tradeoff in setting γ is that ABR causes a distortion in a, per Ito and Sallee (forthcoming), but also creates flexibility for vehicle choices to respond to gasoline price shocks. Also important, however, is the fact that ABR increases fuel use, relative to a flat standard, for any fuel price at which the standard binds. ABR therefore does not, as one might have initially expected, cause fuel use to decrease when the gasoline price is relatively high and increase when the gasoline price is relatively low.

This intuition leads to this section's main result that the optimal fuel economy standard has $\gamma = 0$ and is therefore not attribute-based. Formally, appendix B.3 shows that the FOC for γ given by equation (14) simplifies to equation (16) below, where σ_c is the standard deviation of gasoline prices conditional on the standard binding, and it is assumed that the pivot point is set at (a^*, μ^*) so that equation (15) is satisfied:

$$\frac{\gamma F(\hat{G})(B_{aa} + B_{ea}\gamma)}{(B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa})^2} B_{eG}^2(\phi_g^2 - \sigma_c^2) = 0.$$
(16)

There are only two solutions to equation (16): $\gamma = 0$ and $\gamma = -B_{aa}/B_{ea}$. The former corresponds to a non-attribute-based standard, while the latter is an attribute-based standard such that the slope of the standard is the same as the rate at which unconstrained private agents trade off a and e. Following figure 5, if the pivot point is fixed at (a^*, μ^*) , setting $\gamma = -B_{aa}/B_{ea}$ allows agents to choose the same vehicles that they would have chosen in the absence of regulation. Optimal ABR therefore involves one of two possible corner solutions for γ : either the standard is flat ($\gamma = 0$), or it is so heavily attribute-based ($\gamma = -B_{aa}/B_{ea}$) that the standard is non-binding for any realization of G. These corner outcomes result from the fact that the welfare impacts of both the flexibility benefits and attribute distortions of ABR are second rather than first-order.

At $\gamma = 0$, the SOC for γ reduces to $B_{aa}^{-1}F(\hat{G})B_{eG}^2(\phi_g^2 - \sigma_c^2) < 0$. Recall from section 5.2 that if the level of the standard is set so that it not only satisfies FOC (8) (equivalent

to FOC (15) when $\gamma = 0$) but is also globally optimal conditional on $\gamma = 0$, we must have $\phi_g \geq \sigma_c$. Since $B_{aa} < 0$, it then follows that if μ_0 is globally optimal within the set of flat standards, the SOC for γ is satisfied at $\gamma = 0.4^2$ The globally optimal fixed fuel economy standard therefore has $\gamma = 0$ and is not attribute-based. Thus, the central insight of Ito and Sallee (forthcoming) that ABR reduces welfare extends to the case when uncertainty over future gasoline prices drives uncertainty over future marginal compliance costs. The intuition for this result stems from the fact that the flexibility benefit of ABR is countered both by the distortion to the attribute and by the increase in expected fuel consumption when the standard binds.

Finally, the $\gamma = -B_{aa}/B_{ea}$ solution to FOC (16) only satisfies the SOC when $\phi_g \leq \sigma_c$,⁴³ which per section 5.2 can only happen when μ_0 is set so stringently that a flat standard at μ_0 yields lower welfare than no standard at all. In this case, setting $\gamma = -B_{aa}/B_{ea}$, which is equivalent to a non-binding standard at any gasoline price, will intuitively dominate the flat $\gamma = 0$ standard. Thus, ABR can be valuable when it is used to mitigate the effects of setting a standard with a level μ_0 that is far too stringent. If the level of the standard is set optimally, however, the optimal fuel economy standard is flat.

⁴²FOC (16) and the SOC for γ at $\gamma = 0$ (given by $B_{aa}^{-1}F(\hat{G})B_{eG}^2(\phi_g^2 - \sigma_c^2) < 0$) both assume that FOC (15) for μ_0 is satisfied. Thus, if SOC (11) for μ_0 , conditional on $\gamma = 0$, is satisfied, then satisfaction of the SOC for γ at $\gamma = 0$ is sufficient for a local optimum at $(\mu_0, 0)$ (since the SOC for γ already accounts for the cross-terms in the Hessian). Moreover, since satisfaction of the SOC for γ at $\gamma = 0$ implies that the SOC is violated at $\gamma = -B_{aa}/B_{ea}$, a $\gamma = 0$ local optimum is globally optimal if the choice of μ_0 , conditional on $\gamma = 0$, is globally optimal.

⁴³At $\gamma = -B_{aa}/B_{ea}$, the SOC is given by $-B_{aa}F(\hat{G})Z^{-2}B_{eG}^2(\phi_g^2 - \sigma_c^2) < 0$, where Z is the denominator in equation (16) evaluated at $\gamma = -B_{aa}/B_{ea}$.

7 Additional factors affecting fuel economy feebates or standards under uncertainty

7.1 Covariance between gasoline prices and marginal damage

The results in sections 4 through 6 assume that marginal damage ϕ_g is locally constant and unaffected by changes to the gasoline price G. However, if gasoline price changes are driven by shocks to the global economy or global crude oil supply, the path of global CO₂ emissions may shift. This shift will then cause ϕ_g to change to the extent that the global marginal damage curve for CO₂ is upward-sloping. If changes in ϕ_g and G are positively (negatively) correlated, then the first-best price-indexed fuel economy standard should be more (less) responsive to gasoline price shocks than the price-indexed standard discussed in section 4. For instance, if a decrease in G is associated with a decrease in ϕ_g , the optimal feebate must decrease when G decreases, and therefore the optimal fuel use per mile increases by more than it would have were ϕ_g constant.

The sign of the correlation between shocks to ϕ_g and G is not obvious *a priori*, since it depends on whether global demand shocks or supply shocks are dominant in determining crude oil price movements. For instance, a persistent slowdown in the global economy will decrease the path of global CO₂ emissions and therefore decrease marginal damage ϕ_g . The reduction in economic output will also reduce demand for crude oil, thereby reducing the price of gasoline. ϕ_g and G are therefore positively correlated in this scenario. Shocks to global crude oil supply, however, will drive a negative correlation.

There exists considerable work in macroeconomics that attempts to disentangle supply and demand shocks in the global crude oil market. See Hamilton (2013) and Kilian (2014) for recent surveys and discussion; it is clear from this literature that the extent to which oil price shocks are supply versus demand-driven has varied over time. Baumeister and Kilian (2016) attribute the recent 2014–2015 oil and gasoline price decrease primarily to demand shocks, though positive shocks to oil production played a role as well.

7.2 Banking and borrowing of fuel economy credits

The U.S. fuel economy standards established in 2012 allow manufacturers to bank unused fuel economy credits or borrow credits against future years. This banking and borrowing is a flexibility that in principle can help firms smooth shocks to marginal compliance cost that are driven by gasoline prices. Application of results from Pizer and Prest (2016)—which studies binding dynamic quantity standards under abatement cost uncertainty—implies that even if a fuel economy standard is fixed for many years, the first-best can be achieved if banking and borrowing is permitted (with specific trading ratios) over time horizons that exceed the time elapsed between the regulator's updates to the standard.⁴⁴

In practice, however, banking and borrowing are unlikely to help vehicle manufacturers respond to gasoline price shocks under the U.S. fuel economy standards. First, credits may only be banked for up to five years, and they may be borrowed only to cover deficits from no more than three years in the past. These time periods are substantially shorter than the time between policy updates (the standards set in 2012 are locked-in through at least 2021). Second, gasoline price shocks are typically persistent, as shown in section 2.2, so that a negative shock this year does not imply that the price will rebound next year. Given this persistence, banking and borrowing cannot help firms smooth out compliance cost shocks driven by gasoline price volatility. This point is illustrated in appendix D using a simple two-period model that examines how the incentive to bank or borrow emissions credits varies with the persistence of gasoline price shocks.

7.3 Consumer valuation of fuel economy

What if consumers do not fully value future fuel costs when making a vehicle purchase? If policy-makers have the objective of correcting for not just external damage ϕ but also

⁴⁴Pizer and Prest (2016) shows that a "quantity" policy with banking and borrowing can actually improve welfare over a "price" policy, even if marginal damage is locally constant, if: (1) there are not just shocks to compliance costs but also to marginal damage, (2) these shocks are observable to firms, and (3) the price policy cannot be applied retroactively.

consumer undervaluation, then the optimal tax on fuel consumption per mile must account for both effects. This paper's model can be extended to accommodate standard models of undervaluation (such as Allcott and Wozny (2014); Busse, Knittel, and Zettelmeyer (2013); and Sallee, West, and Fan (2016)) by replacing the B_{eG} term in the private benefit function with λB_{eG} , where $\lambda \in [0, 1]$, with $\lambda = 1$ connoting full valuation. The optimal tax on fuel consumption per mile is then equal to $\phi + (1 - \lambda)B_{eG}G$. This tax increases with the price of gasoline whenever $\lambda < 1$, since the magnitude of undervaluation of future fuel costs increases with the fuel price.⁴⁵ To mimic this tax with a gasoline price-indexed fuel economy standard, the standard must be more responsive to gasoline price shocks the greater is consumer undervaluation. Thus, when consumers undervalue fuel economy in this way, the welfare benefit (relative to a fixed standard) of optimally indexing the fuel economy standard to gasoline prices is enhanced relative to the full valuation baseline case presented in section 4.

7.4 Technology, investment, and adjustment costs

Compliance with fuel economy standards requires not just shifting of the fleet-wide sales mix but also design changes and, ultimately, technology innovation. Indeed, the "ramping up" of U.S. fuel economy standards shown in figure 1 was designed with technological progress in mind. Compliance costs in practice are therefore dynamic rather than static as modeled here. How might dynamic compliance costs affect the economics of a fixed fuel economy standard versus a feebate?

Under perfect competition, a feebate policy will still attain the first-best with dynamic compliance costs, while the fixed standard will not. The intuition is given in Williams (2012): while the optimal quantity of emissions abated may phase in over time in the presence of adjustment costs or technology investment, a Pigouvian tax that causes private actors to equate marginal compliance cost to marginal damage at all times is still optimal. Dynamic

⁴⁵In an alternative model in which fuel costs become more salient to consumers when the gasoline price is high, so that λ is an increasing function of G, the optimal tax may instead be constant or decreasing in G.

compliance costs do, however, complicate the equivalence between a feebate policy and a gasoline price-indexed fuel economy standard. In the static model presented in the paper, the price-indexed standard responds immediately to fuel price shocks. In the presence of dynamic compliance costs, however, the optimal price-indexed standard must respond gradually, mimicking how firms would respond to fuel price shocks under the feebate.

When firms that invest in fuel efficiency technology compete imperfectly, the choice between a price or quantity policy may affect the extent to which the level of innovation approaches the social optimum. Montero (2011) and Scotchmer (2011), for instance, develop models of innovation under price or quantity regulation when innovators have the ability to extract monopoly rents by licensing the technology they develop. The extent to which outcomes from a price or quantity policy deviate from the first-best depends on the size of the innovation, the elasticity of demand for abatement, and the ability of the regulator to commit to its policy.

8 Conclusion

Fuel economy policy in the United States is beset by the difficult problem that standards enacted remain in place for long time horizons over which the price of gasoline can vary substantially. Gasoline price swings, such as the fall in prices experienced between 2014 and 2015, affect the marginal cost of compliance with a fixed fuel economy standard because consumers' valuation of fuel economy co-varies with the price of gasoline. Following Weitzman (1974), the first-best welfare outcome can be achieved by allowing the standard, in miles per gallon, to rise and fall with the price of gasoline (and possibly also with shocks to fuel economy technology, to the extent it can be objectively measured), mimicking what would happen under a feebate. In practice, however, U.S. fuel economy standards are updated only infrequently. The standards set in 2012 are locked in place until 2021, and an ongoing EPA review is expected to finalize standards for 2022–2025. During these long regulatory commitment periods, inevitable swings in the price of gasoline will open up a gap between the regulation's marginal compliance cost and its marginal benefit, reducing expected welfare relative to the first-best.

When fuel economy standards are periodically revised, these revisions can take into account fuel prices at the time of the revision. But what then is the optimal standard to set? The usual rule when committing to a fixed pollution control standard is to equate the expected marginal cost of abatement at the standard to marginal external harm. This rule ignores, however, the possibility that fuel price uncertainty (and therefore marginal cost uncertainty) is sufficiently large, and the policy horizon sufficiently long, that the standard set will not actually bind. I show that given historic gasoline price volatility and estimates of marginal damage from the literature, this possibility is sufficiently likely that an optimallyset fuel economy standard is substantially relaxed relative to that prescribed by the usual rule. The optimal fixed standard may in fact be so lax that it is non-binding at the expected gasoline price, since doing so substantially mitigates the welfare losses associated with low gasoline price realizations. Furthermore, I show that while an attribute-based standard does introduce some flexibility by which consumers can vary their vehicle choice with gasoline prices, the distortions caused to the attribute and to the expected level of fuel consumption are so costly that attribute-basing reduces expected welfare.

While this paper is focused on fuel economy regulation, the concepts and conclusions presented are likely to be relevant for other pollution control settings in which marginal abatement costs are uncertain and marginal damage is locally constant. For instance, the CO_2 standards under the Clean Power Plan relied in part on anticipated reductions in emissions from switching from coal-fired to gas-fired electricity generation. However, the cost of fuel switching, and therefore the cost of abatement, depends on the relative prices of coal and natural gas. If natural gas becomes very cheap relative to coal, the regulation may not bind, much like fuel economy regulations may not bind if the price of gasoline rises. Because the price of natural gas has a volatility even larger than that of gasoline, the probability of such an outcome is far from trivial, so that the optimal standard will deviate substantially from the usual Weitzman (1974) calculation.⁴⁶ As another example, Borenstein *et al.* (2016) estimates substantial uncertainty in business-as-usual emissions, and therefore marginal abatement costs, for California's 2013–2020 greenhouse gas cap-and-trade market. In that market, Borenstein *et al.* (2016) demonstrates that the regulation that is likely to bind is not the quantity standard itself but rather the floor or ceiling on emission permit prices. A potential extension of the results presented here could consider optimal fixed standards in the presence of a ceiling or a non-zero floor on the price of emission permits and therefore marginal abatement costs.

 $^{^{46}}$ Using five-year differences on EIA natural gas prices for the electric power sector (obtained from https://www.eia.gov/dnav/ng/hist/n3045us3m.htm on 19 January, 2018), natural gas price volatility Given a January, 2016 natural gas price of \$3.17 per million Btu (mmBtu), the stanis 64.8%. dard deviation of natural gas prices five years in the future is therefore \$2.05/mmBtu. This volatility is large relative to the carbon externality in terms of \$ per mmBtu of natural gas. For a social cost of carbon of \$38/mton and a January, 2016 electric power sector average coal price of \$2.11/mmBtu (coal price data in \$ per short ton obtained from EIA's Annual Coal Report (released November 2017), table 34, at https://www.eia.gov/coal/annual/pdf/table34.pdf, and coal heat content data (mmBtu per short ton) obtained from EIA's Monthly Energy Review, table A5, at https://www.eia.gov/totalenergy/data/monthly/pdf/sec13_5.pdf), calculations from Cullen and Mansur (2017) indicate that a natural gas price decrease of \$1.27/mmBtu would have an equivalent effect on power sector emissions as would imposition of a carbon tax at the SCC (I use equations (2) and (3), and footnote 6, from Cullen and Mansur (2017) to find the natural gas price decrease that would have the same impact on the coal-to-gas cost ratio as would the carbon tax, locally to the cost ratio at the fuels' social cost). Because this \$1.27/mmBtu externality is less than the natural gas price volatility (in levels) of \$2.05/mmBtu, there is a significant probability that optimally-set electric power CO₂ regulation may not bind.

References

- Allcott, Hunt and Nathan Wozny (2014), "Gasoline Prices, Fuel Economy, and the Energy Paradox", *Review of Economics and Statistics* 96(5), 779-795.
- [2] Alquist, Ron, Lutz Kilian, and Robert J. Vigfusson (2013), "Forecasting the Price of Oil", in Graham, Elliot and Allan Timmerman (Eds.), Handbook of Economic Forecasting vol.2. Amsterdam: North Holland.
- [3] Anderson, Michael L. and Maximilian Auffhammer (2014), "Pounds that Kill: The External Costs of Vehicle Weight", *Review of Economic Studies* 81(2), 535-571.
- [4] 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, 383-403.
- [5] Anderson, Soren T. and James M. Sallee (2016), "Designing Policies to Make Cars Greener: A Review of the Literature", Annual Review of Resource Economics, forthcoming.
- [6] Baumeister, Christiane and Lutz Kilian (2016), "Understanding the Decline in the Price of Oil since June 2014", Journal of the Association of Environmental and Resource Economists 3(1), 131-158.
- [7] Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins (2016), "Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design", Energy Institute at Haas working paper #274.
- [8] Brozovic, Nicholas, David L. Sunding, and David Zilberman (2004), "Prices Versus Quantities Reconsidered", Urbana 51, 61801.

- [9] Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer (2013), "Are Consumers Myopic? Evidence from New and Used Car Purchases", American Economic Review 103(1), 220-256.
- [10] Costello, Christopher and Larry Karp (2004), "Dynamic Taxes and Quotas with Learning", Journal of Economic Dynamics and Control 28, 1661-1680.
- [11] Cullen, Joseph A. and Erin T. Mansur (2017), "Inferring Carbon Abatement Costs in Electricity Markets: A Revealed Preference Approach Using the Shale Revolution", *American Economic Journal: Economic Policy* 9, 106-133.
- [12] Detroit News (7 Apr., 2015), "54.5-mpg Rules Expected to Stick".
- [13] Ellermann, A. Denny and Ian Sue Wing (2003), "Absolute vs. Intensity-Based Emission Caps", MIT Joint Program on the Science and Policy of Global Change report #100.
- [14] Elliot, Graham, Thomas J. Rothenberg, and James H. Stock (1996), "Efficient tests for an autoregressive unit root", *Econometrica* 64(4), 813-836.
- [15] Energy Information Administration (2011), "Voluntary Reporting of Greenhouse Gases Program: Fuel Emission Coefficients". Accessed from http://www.eia.gov/oiaf/1605/coefficients.html#tbl2 on 11 April, 2016.
- [16] Environmental Protection Agency (Aug., 2012), "EPA and NHTSA Set Standards to Reduce Greenhouse Gases and Improve Fuel Economy for Model Years 2017–2025 Cars and Light Trucks", EPA-420-F-12-051.
- [17] Federal Register (15 Oct., 2012), Vol 77, No. 199, p. 62644.
- [18] General Motors (20 Nov., 2009), "General Motors Comments On Proposed Rulemaking to Establish Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards for 2012–2016 model years". Accessed from

https://www.regulations.gov/#!documentDetail D=EPA-HQ-OAR-2009-0472-6953 on 15 May, 2016.

- [19] Gillingham, Kenneth (2013), "The Economics of Fuel Economy Standards versus Feebates", paper prepared for the National Energy Policy Institute.
- [20] Hamilton, James D. (2013), "Historical Oil Shocks", in Randall E. Parker and Robert Whaples, Eds., Routledge Handbook of Major Events in Economic History. New York: Routledge Taylor and Francis Group.
- [21] Interagency Working Group on Social Cost of Carbon (2013), "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866". United States Government. Accessed from https://www.whitehouse.gov/sites/default/files/omb/inforeg/social_cost_of_carbon_for _ria_2013_update.pdf on 19 May, 2016.
- [22] Ito, Koichiro and James M. Sallee (forthcoming), "The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards", *Review of Economics* and Statistics.
- [23] Jacobsen, Mark (2013), "Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior", American Economic Journal: Applied Economics 5(3), 1-26.
- [24] Kaplow, Louis and Steven Shavell (2002), "On the Superiority of Corrective Taxes to Quantity Regulation", American Law and Economics Review 4(1), 1-17.
- [25] Kilian, Lutz (2014), "Oil Price Shocks: Causes and Consequences", Annual Review of Resource Economics 6, 133-154.
- [26] Leard, Benjamin, Joshua Linn, and Virginia McConnell (2016), "Fuel Prices, New Vehicle Fuel Economy, and Implications for Attribute-Based Standards", Resources for the Future Discussion Paper #16-04.

- [27] Levin, Laurence, Matthew S. Lewis, and Frank A. Wolak (forthcoming), "High Frequency Evidence on the Demand for Gasoline", American Economic Journal: Economic Policy.
- [28] Lutsey, Nic (23 Jan., 2015), "Do the automakers really need help with the U.S. efficiency standards?" International Council on Clean Transportation. Blog post. Accessed from http://www.theicct.org/blogs/staff/do-automakers-really-need-help-us-efficiencystandards on 29 Dec., 2015.
- [29] Montero, Juan Pablo (2011), "A Note on Environmental Policy and Innovation When Governments Cannot Commit", *Energy Economics* 33, 513-519.
- [30] National Highway Transportation and Safety Administration (Aug., 2012), "Final Regulatory Impact Analysis, Corporate Average Fuel Economy for MY 2017–MY 2025 Passenger Cars and Light Trucks".
- [31] National Research Council (2002), Effectiveness and Impact of Corporate Average Fuel Economy (CAFE) Standards. Washington, D.C.: National Academy Press.
- [32] National Research Council (2015), Cost, Effectiveness, and Deployment of Fuel Economy Technologies for Light-Duty Vehicles. Washington, D.C.: National Academy Press.
- [33] Newell, Richard G. and William A. Pizer (2003), "Regulating Stock Externalities Under Uncertainty", Journal of Environmental Economics and Management 45, 416-432.
- [34] Newell, Richard G. and William A. Pizer (2008), "Indexed Regulation", Journal of Environmental Economics and Management 56, 221-233.
- [35] Ng, Serena and Pierre Perron (2001), "Lag, Length, Selection and the Construction of Unit Root Tests With Good Size and Power", *Econometrica* 69(6), 1519-1554.

- [36] Parry, Ian W.H., Carolyn Fischer, and Winston Harrington (2004), "Should Corporate Average Fuel Economy (CAFE) Standards be Tightened?" Resources for the Future Discussion Paper #04-53.
- [37] Parry, Ian W.H., Dirk Heine, Eliza Lis, and Shanjun Li (2014), Getting Energy Prices Right: From Principle to Practice. Washington, D.C.: International Monetary Fund.
- [38] Parry, Ian W.H., Margaret Wells, and Winston Harrington (2007), "Automobile Externalities and Policies", *Journal of Economic Literature* 45(2), 373-399.
- [39] Pizer, William A. and Brian Prest (2016), "Prices versus Quantities with Policy Updating", NBER working paper #22379.
- [40] Sallee, James M., Sarah E. West, and Wei Fan (2016), "Do Consumers Recognize the Value of Fuel Economy? Evidence from Used Car Prices and Gasoline Price Fluctuations", *Journal of Public Economics* 135, 61-73.
- [41] Scotchmer, Suzanne (2011), "Cap-and-Trade, Emissions Taxes, and Innovation", in J. Lerner and S. Stern, Eds. Innovation Policy and the Economy vol. 11. University of Chicago Press.
- [42] Weitzman, Martin L. (1974), "Prices vs. Quantities", Review of Economic Studies 41(4), 477-491.
- [43] Whitefoot, Kate S. and Steven J. Skerlos (2012), "Design Incentives to Increase Vehicle Size Created From the U.S. Footprint-Based Fuel Economy Standards", *Energy Policy* 41, 402-411.
- [44] Williams, Roberton C., III (2012), "Setting the Initial Time-Profile of Climate Policy: The Economics of Environmental Policy Phase-Ins", in D. Fullerton and C. Wolfram, Eds. The Design and Implementation of U.S. Climate Policy. University of Chicago Press.

Appendix

A Representative consumer model

Consider a unit mass of consumer types *i*, each of whom purchases a vehicle with fuel use per mile e_i and attribute a_i , yielding private benefit $B^i(e_i, a_i, G) = U^i(e_i, a_i, G) - C(e_i, a_i, G)$. This section proves that a sufficient statistic for the aggregate private welfare effects of fuel economy policies is the private benefit B(e, a, G) obtained from the average vehicle purchased, where *e* is the average of the e_i , *a* is the average of the a_i , and the first derivatives of *B* equal the average of the B_e^i and B_a^i .

I begin by showing that all consumer types make identical changes to their fuel use and attribute choices when there is a change to either G or a tax τ on fuel use per mile. Given G and τ , consumer *i*'s choice of e_i and a_i is given by the solution to the system of FOCs

$$B_e^i(e_i, a_i, G) = \tau \tag{17}$$

$$B_a^i(e_i, a_i, G) = 0. (18)$$

Application of the implicit function theorem yields:

$$\frac{de_i}{dG} = \frac{-B_{aa}B_{eG}}{B_{ee}B_{aa} - B_{ea}^2} < 0; \ \frac{de_i}{d\tau} = \frac{B_{aa}}{B_{ee}B_{aa} - B_{ea}^2} < 0 \tag{19}$$

$$\frac{da_i}{dG} = \frac{B_{ea}B_{eG}}{B_{ee}B_{aa} - B_{ea}^2} \le 0; \ \frac{da_i}{d\tau} = \frac{-B_{ea}}{B_{ee}B_{aa} - B_{ea}^2} \le 0,$$
(20)

where it is assumed that the SOC $B_{ee}B_{aa} - B_{ea}^2 > 0$ holds so that the private problem has an interior optimum. Under the second order Taylor expansion assumptions of the text, each of the derivatives above is a constant scalar that is identical across all types *i*. Thus, changes to *G* or to the fuel economy tax τ affect all types identically. Moreover, changes to τ , given *G*, are one-to-one with changes in *e* and *a*, the average fuel use and attribute. Thus, if a fuel economy standard (with compliance credit trading) is used in lieu of a tax, shifts in the standard are isomorphic to shifts in the implied tax τ . Thus, it is sufficient to show that B(e, a, G) is a sufficient statistic for welfare under a tax policy; I do not need to explicitly model a fixed fuel economy standard.

Using the second order Taylor expansion, let B_e^i and B_a^i be written as:

$$B_e^i(e_i, a_i, G) = B_{e0}^i + B_{ea}(a_i - a^*) + B_{ee}(e_i - e^*) + B_{eG}(G - \bar{G})$$
(22)

$$B_a^i(e_i, a_i, G) = B_{a0}^i + B_{ea}(e_i - e^*) + B_{aa}(a_i - a^*),$$
(23)

where e^* and a^* are the socially optimal fuel use and attribute, when the gasoline price equals its expected price \bar{G} , for the representative consumer, who has marginal utility intercepts given by $B_{e0} = \int_i B_{e0}^i di$ and $B_{a0} = \int_i B_{a0}^i di$. Aggregate private welfare is then given by:

$$\int_{i} B^{i}(e_{i}, a_{i}, G) di = \int_{i} B^{i}_{e0}(e_{i} - e^{*}) di + \int_{i} B^{i}_{a0}(a_{i} - a^{*}) di + \frac{1}{2} B_{ee} \int_{i} (e_{i} - e^{*})^{2} di + \frac{1}{2} B_{aa} \int_{i} (a_{i} - a^{*})^{2} di + B_{ea} \int_{i} (e_{i} - e^{*})(a_{i} - a^{*}) di + B_{eG} \int_{i} (e_{i} - e^{*})(G - \bar{G}) di.$$
(24)

Define $e \equiv \int_i e_i di$ and $a \equiv \int_i a_i di$. Replace the $(e_i - e^*)$ terms in equation (24) with $(e_i - e + e - e^*)$, and do likewise with the $(a_i - a^*)$ terms. After removing terms for which the expectation is zero, and after grouping related terms together, we obtain:

$$\int_{i} B^{i}(e_{i}, a_{i}, G) di = [B_{e0}(e - e^{*}) + B_{a0}(a - a^{*}) + \frac{1}{2}B_{ee}(e - e^{*})^{2} + \frac{1}{2}B_{aa}(a - a^{*})^{2}
+ B_{ea}(e - e^{*})(a - a^{*}) + B_{eG}(e - e^{*})(G - \bar{G})]
+ [\int_{i} B^{i}_{e0}(e_{i} - e)di + \int_{i} B^{i}_{a0}(a_{i} - a)di + \frac{1}{2}B_{ee}\int_{i}(e_{i} - e)^{2}di
+ \frac{1}{2}B_{aa}\int_{i}(a_{i} - a)^{2}di + B_{ea}\int_{i}(e_{i} - e)(a_{i} - a)di]$$
(25)

The first bracketed term in equation (25) is the second order Taylor expansion of B(e, a, G). The second bracketed term is invariant to G or to fuel economy policy, since $e_i - e$ and $a_i - a$ are constants for each *i*. Thus, B(e, a, G) is a sufficient statistic for the private welfare effects of fuel economy policy.

B Proofs and derivations

B.1 Expected welfare loss from an optimal standard

This section derives equation (9) from the main text, which gives the expected welfare loss, relative to the first-best, for an optimal flat fuel economy standard μ^* that may not bind.

Define \hat{G} implicitly by $B_e(\mu^*, \hat{G}) = 0$. As in the main text, define $\tilde{\mu}$ implicitly by $E[B_e(\tilde{\mu}, G)] = \phi$. Per the second order Taylor expansion, marginal utility is given by

$$B_e(e,G) = B_{eG}(G - \hat{G}) + B_{ee}(e - \mu^*).$$
(26)

When the standard does not bind, rearranging (26) yields the function e(G) that gives agents' fuel economy choice as a function of G:

$$e(G) = \mu^* - \frac{B_{eG}}{B_{ee}}(G - \hat{G}).$$
(27)

If the first-best tax $\tau = \phi$ is imposed instead of the fixed standard, agents' choices are

instead given by the function $\tilde{e}(G)$:

$$\tilde{e}(G) = \mu^* + \frac{\phi}{B_{ee}} - \frac{B_{eG}}{B_{ee}}(G - \hat{G}).$$
(28)

Using the Taylor expansion $B(e,G) = B_{eG}(e - \mu^*)(G - \hat{G}) + \frac{B_{ee}}{2}(e - \mu^*)^2$, we may then obtain an expression for B(e(G), G):

$$B(e(G), G) = B_{eG} \frac{-B_{eG}}{B_{ee}} (G - \hat{G})^2 + \frac{B_{ee}}{2} \left(\frac{-B_{eG}}{B_{ee}}\right)^2 (G - \hat{G})^2$$
$$= \frac{-B_{eG}^2}{2B_{ee}} (G - \hat{G})^2,$$
(29)

and an expression for $B(\tilde{e}(G), G)$:

$$B(\tilde{e}(G), G) = B_{eG} \left(\frac{\phi}{B_{ee}} - \frac{B_{eG}}{B_{ee}} (G - \hat{G}) \right) (G - \hat{G}) + \frac{B_{ee}}{2} \left(\frac{\phi}{B_{ee}} - \frac{B_{eG}}{B_{ee}} (G - \hat{G}) \right)^2$$

$$= \frac{-B_{eG}^2}{2B_{ee}} (G - \hat{G})^2 + \frac{\phi^2}{2B_{ee}}.$$
 (30)

Define Δ to be the expected welfare loss under μ^* relative to the first best. Δ is given by:

$$\Delta \equiv \int_{G_L}^{\hat{G}} \left[B(\tilde{e}(G), G) - \phi \tilde{e}(G) - B(\mu^*, G) + \phi \mu^* \right] f(G) dG + \int_{\hat{G}}^{G_U} \left[B(\tilde{e}(G), G) - \phi \tilde{e}(G) - B(e(G), G) + \phi e(G) \right] f(G) dG.$$
(31)

Let Δ_1 denote the first integral in (31). Noting that $B(\mu^*, G) = 0$, substituting using (28) and (30), and cancelling terms, we obtain

$$\Delta_1 = \frac{-1}{2B_{ee}} \int_{G_L}^{\hat{G}} \left[B_{eG}^2 (G - \hat{G})^2 - 2\phi B_{eG} (G - \hat{G}) + \phi^2 \right] f(G) dG.$$
(32)

Next, note that if we multiply FOC (7) by B_{eG} , we obtain

$$\int_{G_L}^{\hat{G}} \left(B_{eG}(G - \hat{G}) - \phi \right) f(G) dG = 0.$$
(33)

Substituting (33) into (32) yields:

$$\Delta_1 = \frac{-1}{2B_{ee}} \int_{G_L}^{\hat{G}} \left[B_{eG}^2 (G - \hat{G})^2 - \phi^2 \right] f(G) dG.$$
(34)

Define $\dot{G} \equiv \frac{1}{F(\hat{G})} \int_{G_L}^{\hat{G}} Gf(G) dG$ (i.e., \dot{G} is the expected value of G conditional on the

standard binding). Then define $\sigma_c^2 \equiv \frac{1}{F(\hat{G})} \int_{G_L}^{\hat{G}} (G - \dot{G})^2 f(G) dG$. We then have:

$$B_{eG}^{2} \int_{G_{L}}^{\hat{G}} (G - \hat{G})^{2} f(G) dG = B_{eG}^{2} \int_{G_{L}}^{\hat{G}} (G - \dot{G} + \dot{G} - \hat{G})^{2} f(G) dG$$

$$= F(\hat{G}) B_{eG}^{2} (\sigma_{c}^{2} + (\dot{G} - \hat{G})^{2})$$

$$= F(\hat{G}) (B_{eG}^{2} \sigma_{c}^{2} + \phi^{2}), \qquad (35)$$

where the last line makes use of equation (33) and the definition of G. Substituting (35) into (34) yields:

$$\Delta_1 = \frac{-F(\hat{G})B_{eG}^2 \sigma_c^2}{2B_{ee}}.$$
(36)

Returning to equation (31), let Δ_2 denote the second integral. Making substitutions using equations (27)-(30), we obtain:

$$\Delta_2 = \int_{\hat{G}}^{G_U} \frac{-\phi^2}{2B_{ee}} f(G) dG = \frac{-\phi^2 (1 - F(\hat{G}))}{2B_{ee}}.$$
(37)

Combining terms yields expression (9) from the main text, which simplifies to the standard Weitzman (1974) formula, $\Delta = \frac{-\sigma^2 B_{eG}^2}{2B_{ee}}$, when the standard always binds:

$$\Delta = -\frac{B_{eG}^2}{2B_{ee}} \left(F(\hat{G}(\mu^*))\sigma_c^2 + (1 - F(\hat{G}(\mu^*)))\phi_g^2 \right).$$

B.2 Welfare from a fixed standard versus a degenerate standard

This section derives the expected welfare benefit δ of a standard μ^* that satisfies FOC (4) relative to imposing a degenerate standard that never binds. Reusing notation from appendix B.1, δ is then defined by:

$$\delta \equiv \int_{G_L}^{\hat{G}} \left[B(\mu^*, G) - \phi \mu^* - B(e(G), G) + \phi e(G) \right] f(G) dG.$$
(38)

Applying equations (27) and (29), we can simplify expression (38) to:

$$\delta \equiv \int_{G_L}^{\hat{G}} \left[\frac{-B_{eG}}{B_{ee}} (G - \hat{G})\phi + \frac{B_{eG}^2}{2B_{ee}} (G - \hat{G})^2 \right] f(G) dG.$$
(39)

Next, apply equations (33) and (35) to obtain

$$\delta = \frac{F(\hat{G})}{2B_{ee}} \left(B_{eG}^2 \sigma_c^2 + \phi^2 - 2\phi^2 \right) = \frac{-F(\hat{G}) B_{eG}^2}{2B_{ee}} (\phi_g^2 - \sigma_c^2).$$
(40)

Thus, a standard μ^* that satisfies FOC (4) improves welfare relative to a degenerate standard iff $\phi_g^2 - \sigma_c^2 \ge 0$. This result can be extended to the case in which μ^* always binds (so that $\sigma_c^2 = \sigma^2$) by defining f(G) = 0 between G_U and $\hat{G}(\mu^*)$.

B.3 Optimal attribute-based standard

This section derives expressions (15) and (16)—the FOC for μ_0 and the FOC for γ (conditional on satisfaction of (15)) from the main text. To begin, it is useful to first write down the private agents' vehicle choice problem subject to a binding attribute-based standard and derive how the choices of a and e vary with μ_0 , γ , and G.

When e is constrained to lie on $e = \mu_0 + \gamma a$, the agents' problem can be written as:

$$\max_{a} B(\mu_0 + \gamma a, a, G) \tag{41}$$

The agents' FOC and SOC, using the notation $e(a) = \mu_0 + \gamma a$ are then:

$$B_e(e(a), a, G)\gamma + B_a(e(a), a, G) = 0$$
(42)

$$B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa} < 0 \tag{43}$$

The solution to FOC (42), combined with the constraint $e = \mu_0 + \gamma a$, yields agents' choices of e and a as functions $e(\mu_0, \gamma, G)$ and $a(\mu_0, \gamma, G)$. The implicit function theorem can then be used to establish the following useful derivatives when the standard binds (where the notation suppresses the dependence of B_a , B_e , da/dG, de/dG, $da/d\mu_0$, $de/d\mu_0$, $da/d\gamma$, and $de/d\gamma$ on μ_0 , γ , and G):

$$\frac{da}{dG} = \frac{-B_{eG}\gamma}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}} \le 0 \tag{44}$$

$$\frac{de}{dG} = \gamma \frac{da}{dG} = \frac{-B_{eG}\gamma^2}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}} \le 0$$
(45)

$$\frac{da}{d\mu_0} = \frac{-(B_{ee}\gamma + B_{ea})}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}}$$
(46)

$$\frac{de}{d\mu_0} = 1 + \gamma \frac{da}{d\mu_0} = \frac{B_{ea}\gamma + B_{aa}}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}} \tag{47}$$

$$\frac{da}{d\gamma} = \frac{-(B_e(e(a), a, G) + B_{ee}a\gamma + B_{ea}a)}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}}$$
(48)

$$\frac{de}{d\gamma} = a + \gamma \frac{da}{d\gamma} = \frac{B_{ea}a\gamma + B_{aa}a - B_e(e(a), a, G)\gamma}{B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}}$$
(49)

We next develop simple expressions for B_e and B_a while constrained. Let (\hat{a}, \hat{e}) denote the pivot point for the standard and \hat{G} denote the gasoline price at which the standard just binds $(\hat{e}, \hat{a}, \text{ and } \hat{G} \text{ are all functions of } \mu_0 \text{ and } \gamma)$. Because $B_a(\hat{e}, \hat{a}, G) = 0$ (recall that $B_{aG} = 0$), we may write the Taylor expansion for B_a as:

$$B_a(e, a, G) = B_{ea}(e - \hat{e}) + B_{aa}(a - \hat{a}).$$
(50)

Along the standard, $e - \hat{e}$ and $a - \hat{a}$ are linear functions of $G - \hat{G}$, per the derivatives given by equations (45) and (44), respectively. Substituting, we obtain:

$$B_{a}(e, a, G) = \frac{-\gamma (B_{ea}\gamma + B_{aa})}{B_{ee}\gamma^{2} + 2B_{ea}\gamma + B_{aa}} B_{eG}(G - \hat{G}).$$
(51)

We can then use FOC (42) to obtain $B_e(e, a, G)$:

$$B_{e}(e, a, G) = \frac{B_{ea}\gamma + B_{aa}}{B_{ee}\gamma^{2} + 2B_{ea}\gamma + B_{aa}}B_{eG}(G - \hat{G}).$$
(52)

We are now in position to begin working with the FOC for μ_0 given by equation (13) in the text. To simplify notation, let $S = B_{ee}\gamma^2 + 2B_{ea}\gamma + B_{aa}$. Sequentially applying the agents' FOC (42) and then expressions (46), (47), and (52), we obtain:

$$FOC_{\mu_{0}} : \int_{G_{L}}^{\hat{G}} \left(B_{e} \left(\frac{de}{d\mu_{0}} - \gamma \frac{da}{d\mu_{0}} \right) - \phi \frac{de}{d\mu_{0}} \right) f(G) dG = 0$$

$$= \frac{1}{S} \int_{G_{L}}^{\hat{G}} \left(\frac{(B_{ea}\gamma + B_{aa})B_{eG}(G - \hat{G})S}{S} - \phi(B_{ea}\gamma + B_{aa}) \right) f(G) dG = 0$$

$$= \frac{B_{ea}\gamma + B_{aa}}{S} \int_{G_{L}}^{\hat{G}} (B_{eG}(G - \hat{G}) - \phi) f(G) dG = 0$$
(53)

$$=\frac{B_{eG}(B_{ea}\gamma+B_{aa})}{S}\int_{G_L}^{\hat{G}}(\phi_g-(\hat{G}-G))f(G)dG=0.$$
(54)

Equation (54) matches equation (15) from the text.

We may next turn to the FOC for γ given by equation (14). Sequentially applying the agents' FOC (42) and then expressions (48), (49), and (52), we obtain:

$$FOC_{\gamma} : \int_{G_L}^{\hat{G}} \left(B_e(\frac{de}{d\gamma} - \gamma \frac{da}{d\gamma}) - \phi \frac{de}{d\gamma} \right) f(G)dG = 0$$

$$= \frac{1}{S} \int_{G_L}^{\hat{G}} \left(\frac{(B_{ea}\gamma + B_{aa})B_{eG}(G - \hat{G})Sa}{S} - \phi(B_{ea}a\gamma + B_{aa}a - B_e\gamma) \right) f(G)dG = 0$$

$$= \frac{B_{ea}\gamma + B_{aa}}{S} \int_{G_L}^{\hat{G}} \left(B_{eG}(G - \hat{G})a - \phi a + \frac{\phi B_{eG}(G - \hat{G})\gamma}{S} \right) f(G)d(G) = 0$$

$$= \frac{B_{ea}\gamma + B_{aa}}{S^2} \int_{G_L}^{\hat{G}} ((B_{eG}(G - \hat{G})a - \phi a)S + \phi\gamma B_{eG}(G - \hat{G}))f(G)dG = 0$$
(55)

Next, we can use the FOC for μ_0 , equation (54), to simplify further:

$$= \frac{B_{ea}\gamma + B_{aa}}{S^2} \int_{G_L}^{\hat{G}} ((a-\hat{a})S(B_{eG}(G-\hat{G}) - \phi) + \phi^2\gamma)f(G)dG = 0$$
(56)

Next apply equation (44) and then (54) again to obtain:

$$= \frac{B_{ea}\gamma + B_{aa}}{S^2} \int_{G_L}^{\hat{G}} \left(-B_{eG}\gamma(G - \hat{G})(B_{eG}(G - \hat{G}) - \phi) + \phi^2\gamma \right) f(G)dG = 0$$

$$= \frac{\gamma(B_{ea}\gamma + B_{aa})}{S^2} \int_{G_L}^{\hat{G}} \left(-B_{eG}^2(G - \hat{G})^2 + 2\phi^2 \right) f(G)dG = 0$$
(57)

Finally, apply equation (35) to obtain:

$$FOC_{\gamma} : \frac{\gamma F(\hat{G})(B_{ea}\gamma + B_{aa})}{S^2} (\phi^2 - B_{eG}^2 \sigma_c^2) = 0$$
$$= \frac{\gamma F(\hat{G})(B_{ea}\gamma + B_{aa})}{S^2} B_{eG}^2 (\phi_g^2 - \sigma_c^2) = 0$$
(58)

Equation (58) matches equation (16) from the main text.

C Prices versus quantities with a rebound effect

This appendix shows that the main result from section 4—a feebate delivers greater expected welfare than a fuel economy standard—holds in the presence of a rebound effect, so long as the gasoline price elasticity of driving is not too large. In this setting, neither instrument achieves the first-best (a gasoline tax is required), and under complete information the feebate and fuel economy standard are equivalent second-best policies. Throughout this appendix, I abstract away from vehicle attributes other than fuel economy, and I assume that all consumers are homogenous.⁴⁷

Allowing a rebound effect requires a substantial modification of the model. Let e and C(e) denote the vehicle's fuel use per mile and production cost, as in the main text, and let G and ϕ_g continue to represent that gasoline price and per-gallon externality. Let Q denote the consumer's miles driven, and let V(Q) denote the utility from driving, where V'(Q) > 0, and V''(Q) < 0 and is locally constant. Private indirect utility is then given by:

$$U(Q, e) = V(Q) - C(e) - eGQ.$$
(59)

Note that as $V'' \to -\infty$ so that miles driven are perfectly inelastic, the model becomes equivalent to that in the main text, where $B_{ee} = -C''$ and $B_{eG} = -Q$. Because the feebate strictly welfare dominates the fuel economy standard with perfectly inelastic demand for miles driven (so long as f(G) is non-degenerate), it is straightforward to show that, given f(G), ϕ_g , and C'', there must exist some arbitrarily small but non-zero value of -1/V'' such that the feebate continues to dominate the standard. That is, the result in section 4 is robust to a small rebound effect.

⁴⁷Allowing for heterogeneity in the intercepts of consumers' marginal utility functions, similar to appendix A, will here generate heterogeneity in the responses of miles driven and fuel economy to gasoline price shocks. Welfare can then no longer be modeled using a representative consumer.

To make further progress, consider welfare under the fuel economy standard. For simplicity, assume that the standard always binds (allowing for the possibility that the standard may not bind will further reduce the expected welfare from the standard relative to the feebate). Given a standard μ , the consumer's FOCs are given by

$$V'(Q) - \mu G = 0 \tag{60}$$

$$-C'(\mu) - GQ - \lambda = 0, \tag{61}$$

where λ is the shadow value of the imposed constraint. With *e* fixed at μ , we have $dQ/dG = \mu/V'' \leq 0$ and $dQ/d\mu = G/V'' \leq 0$. The optimal standard then solves:

$$\max_{\mu} \int_{G_L}^{G_H} \left(V(Q) - C(\mu) - \mu Q(G + \phi_g) \right) f(G) dG.$$
(62)

The FOC is given by:

$$\int_{G_L}^{G_H} \left(V'(Q) \frac{dQ}{d\mu} - C'(\mu) - (Q + \mu \frac{dQ}{d\mu})(G + \phi_g) \right) f(G) dG = 0.$$
(63)

Applying equation (60) and $dQ/d\mu = G/V''$ then yields that the optimal fuel economy standard μ^* satisfies:

$$-C'(\mu^*) = \int_{G_L}^{G_H} Q(G + \phi_g) f(G) dG + \frac{\mu^* \phi_g \bar{G}}{V''}.$$
(64)

Note that as $V'' \to -\infty$, equation (64) simplifies to $-C'(\mu^*) = Q(\bar{G} + \phi_g)$, consistent with equation (2) from the main text. Let Q^* denote miles driven when the fuel economy standard is μ^* and the gasoline price is \bar{G} ; i.e, $V'(Q^*) = \mu^* \bar{G}$.

Expected welfare under the fuel economy standard can then be obtained using the Taylor expansion for V(Q) and the fact that $dQ/dG = \mu^*/V''$. Given the standard μ^* and a gasoline price G, we have:

$$V(Q(\mu^*, G)) = V(Q^*) + \mu^* \bar{G}(Q - Q^*) + \frac{V''}{2} (Q - Q^*)^2$$

= $V(Q^*) + \frac{\mu^{*2} \bar{G}}{V''} (G - \bar{G}) + \frac{\mu^{*2}}{2V''} (G - \bar{G})^2$
= $V(Q^*) + \frac{\mu^{*2}}{V''} \left(\bar{G}(G - \bar{G}) + \frac{1}{2} (G - \bar{G})^2 \right).$ (65)

Expected social welfare $E[W_{\mu}]$ under the optimal standard μ^* is then:

$$E[W_{\mu}] = V(Q^*) + \frac{\mu^{*2}}{V''} \int_{G_L}^{G_H} \left(\bar{G}(G - \bar{G}) + \frac{1}{2}(G - \bar{G})^2 \right) f(G) dG - C(\mu^*) - \mu^* \int_{G_L}^{G_H} (G + \phi_g) (Q^* + \frac{\mu^*}{V''}(G - \bar{G})) f(G) dG.$$
(66)

Next consider a vehicle tax of t dollars per gallon per mile. Under this tax, the consumer's FOCs are given by:

$$V'(Q) - eG = 0 \tag{67}$$

$$-C'(e) - GQ - t = 0. (68)$$

Application of the implicit function theorem yields the following derivatives:

$$\frac{dQ}{dG} = \frac{eC'' - GQ}{V''C'' + G^2}$$
(69)

$$\frac{de}{dG} = \frac{-eG - QV''}{V''C'' + G^2}$$
(70)

$$\frac{dQ}{dt} = \frac{-G}{V''C'' + G^2} \tag{71}$$

$$\frac{de}{dt} = \frac{-V^*}{V''C'' + G^2}.$$
(72)

The consumer's SOC requires that the denominator of each of the above derivatives be strictly negative. As $V'' \to -\infty$, the derivatives de/dG and de/dt collapse to -Q/C'' and -1/C'', consistent with the main text. Using FOCs (67) and (68), a necessary and sufficient condition for the derivatives dQ/dG and de/dG to both be negative (consistent with empirical observations) is that -eC''/(C'(e) + t) > 1 and -QV''/V'(Q) > 1, respectively. In other words, the response of fuel economy to the marginal cost of fuel economy, holding Q fixed, must be inelastic, as must be the response of miles driven to the marginal cost of driving, holding e fixed. I adopt these assumptions on these elasticities for the remainder of the appendix.

The optimal fuel economy tax solves:

$$\max_{t} \int_{G_L}^{G_H} \left(V(Q) - C(e) - eQ(G + \phi_g) \right) f(G) dG, \tag{73}$$

where Q and e are both functions of t and G. The FOC is given by:

$$\int_{G_L}^{G_H} \left(V'(Q) \frac{dQ}{dt} - C'(e) \frac{de}{dt} - \left(Q \frac{de}{dt} + e \frac{dQ}{dt}\right) (G + \phi_g) \right) f(G) dG = 0.$$
(74)

Using the consumer's FOCs (67) and (68) and the derivatives (71) and (72), we obtain:

$$\int_{G_L}^{G_H} \frac{1}{V''C'' + G^2} \left(e\phi_g G - tV'' + \phi_g QV'' \right) f(G) dG = 0$$

$$\leftrightarrow t \int_{G_L}^{G_H} \frac{1}{V''C'' + G^2} f(G) dG = \int_{G_L}^{G_H} \frac{1}{V''C'' + G^2} \left(\phi_g Q + \frac{e\phi_g G}{V''} \right) f(G) dG.$$
(75)

Note that as $V'' \to -\infty$, equation (75) simplifies to $t = \phi_g Q = \phi$, consistent with the main text. More generally, when V'' is real valued the level of fuel economy under the optimal tax and at the average gasoline price \bar{G} will differ from the optimal fuel economy standard

 μ^* given by equation (64), since the response of Q to shocks to G differs between these two instruments.

The goal of this section is to compare expected welfare under the optimal tax, $E[W_t]$, to expected welfare under the optimal standard, $E[W_{\mu}]$. A tractable comparison requires the use of Taylor expansions for V(Q) and C(e) under the tax locally to μ^* and Q^* as defined by equation (64) and $V'(Q^*) = \mu^* \bar{G}$. Because the difference between μ^* and $e(\bar{G})$ under the optimal tax is not tractable owing to the integrals in equation (75), I instead calculate expected welfare for the tax assuming that the tax is set so that, at \bar{G} , the realized fuel economy is μ^* and the realized miles driven are Q^* . The difference in expected welfare between the tax and standard policies that I derive will then be a lower bound on the actual difference between the optimal tax and the optimal standard.

For tractability, I assume that dQ/dG and de/dG are locally constant; i.e., I approximate these derivatives at \bar{G} . The Taylor expansions for V(Q) and C(e) around Q^* and μ^* can then be expressed as a quadratic functions of $G - \bar{G}$:

$$V(Q) = V(Q^*) + \mu^* \bar{G} \frac{dQ}{dG} (G - \bar{G}) + \frac{V''}{2} (\frac{dQ}{dG})^2 (G - \bar{G})^2$$

$$C(e) = C(\mu^*) + C'(\mu^*) \frac{de}{dG} (G - \bar{G}) + \frac{C''}{2} (\frac{de}{dG})^2 (G - \bar{G})^2.$$

Define Δ as the difference in expected welfare between the optimal tax and the optimal standard. Given a non-optimal tax, the Taylor approximations, and the approximations on dQ/dG and de/dG, we have:⁴⁸

$$\Delta \ge \int_{G_L}^{G_H} [(\frac{V''}{2} (\frac{dQ}{dG})^2 - \frac{\mu^{*2}}{2V''})(G - \bar{G})^2 - \frac{C''}{2} (\frac{de}{dG})^2 (G - \bar{G})^2 - \mu^* \frac{dQ}{dG} (G - \bar{G})(G + \phi_g) - Q^* \frac{de}{dG} (G - \bar{G})(G + \phi_g) - \frac{dQ}{dG} \frac{de}{dG} (G - \bar{G})^2 (G + \phi_g) + \frac{\mu^{*2}}{V''} (G - \bar{G})(G + \phi_g)] f(G) dG,$$
(76)

where all dQ/dG and de/dG terms pertain to these derivatives under the tax. Simplifying further yields:

$$\Delta \geq \frac{1}{2} (V''(\frac{dQ}{dG})^2 - \frac{\mu^{*2}}{V''})\sigma^2 - \frac{C''}{2} (\frac{de}{dG})^2 \sigma^2 + \mu^* (\frac{\mu^*}{V''} - \frac{dQ}{dG})\sigma^2 - Q^* \frac{de}{dG}\sigma^2 - \frac{dQ}{dG} \frac{de}{dG} \int_{G_L}^{G_H} (G - \bar{G})^2 (G + \phi_g) f(G) dG.$$
(77)

Next, note that:

$$\frac{V''}{2}\left(\frac{dQ}{dG}\right)^2 - \frac{\mu^{*2}}{2V''} + \mu^*\left(\frac{\mu^*}{V''} - \frac{dQ}{dG}\right) = \frac{\mu^{*2}}{2V''} - \mu^*\frac{dQ}{dG} + \frac{V''}{2}\left(\frac{dQ}{dG}\right)^2 \\ = \frac{1}{2V''}(\mu^* - V''\frac{dQ}{dG})^2 \ge 0.$$
(78)

⁴⁸Equation (76) is derived by eliminating all terms involving multiplication of $(G - \bar{G})$ by a constant.

And further note that:

$$-\frac{C''}{2}\left(\frac{de}{dG}\right)^2 - Q^*\frac{de}{dG} + \frac{Q^*}{2}\frac{de}{dG} = \frac{de}{dG}\left(-\frac{C''}{2}\frac{de}{dG} - \frac{Q^*}{2}\right) \ge 0,$$
(79)

where the inequality comes from the fact that $de/dG \ge -Q^*/C''$ per equation (70). Using equations (78) and (79), we may then write:

$$\Delta \geq -\frac{Q^*}{2} \frac{de}{dG} \sigma^2 - \frac{dQ}{dG} \frac{de}{dG} \int_{G_L}^{G_H} (G - \bar{G})^2 (G + \phi_g) f(G) dG$$

$$\geq -\frac{de}{dG} \int_{G_L}^{G_H} (G - \bar{G})^2 \left(\frac{Q^*}{2} + \frac{dQ}{dG} (G + \phi_g)\right) f(G) dG$$

$$\geq -\frac{1}{2} \frac{de}{dG} \int_{G_L}^{G_H} (G - \bar{G})^2 Q^* \left(1 + 2\frac{Q}{Q^*} \varepsilon_{G + \phi_g}\right) f(G) dG, \tag{80}$$

where $\varepsilon_{G+\phi_g}$ is the gasoline price elasticity of miles driven, holding e constant, at a gasoline price of $G + \phi_g$. Equation (80) suggests that the optimal fuel economy tax will yield greater expected welfare than the optimal fuel economy standard so long as the magnitude of $\varepsilon_{G+\phi_g}$ is less than roughly 0.5. This value exceeds recent elasticity estimates with magnitudes between 0.27 and 0.35 from Levin, Lewis, and Wolak (forthcoming), which are the highest among recent studies. Note also that if $\varepsilon_{G+\phi_g} = 0$ (i.e., $V'' \to -\infty$), then this expression reduces to $\Delta = Q^2 \sigma^2 / 2C''$, consistent with the main text.⁴⁹

Why does the advantage of the tax over the standard diminish with the rebound effect? The intuition flows from the fact that, under the tax, fuel economy responds more strongly to gasoline price shocks than what is socially optimal.⁵⁰ The magnitude of this over-response increases with the responsiveness of miles driven to the cost of driving. Eventually, if the magnitude of $\varepsilon_{G+\phi_g}$ is large enough, it is better to have a fixed fuel economy standard (which yields insufficient responsiveness of e to G) than a fuel economy tax (too much responsiveness of e to G).

D Banking and borrowing

This appendix presents a model involving two compliance periods, in which firms can bank or borrow emissions credits across the two periods. The main result is that if fuel price shocks are persistent across periods, there is no incentive to bank or borrow credits. Mean-reverting fuel price expectations will, however, lead to some banking or borrowing.

Let μ denote the imposed standard, which must hold on average across the two periods. Let G_1 and G_2 denote the first and second period gasoline prices, and let e_1 and e_2 denote the first and second period choices of fuel consumption per mile. The policy requires that

⁴⁹Recall that when miles driven are perfectly inelastic at Q, B_{eG} and B_{ee} from the model in the main text are equivalent to -Q and -C'', respectively. Also, the formula for Δ is an equality rather than an inequality because all approximations used in deriving equation (80) become equalities when $V'' \to -\infty$.

⁵⁰This effect is driven by the fact that consumers drive more miles than what is socially optimal, so that their fuel economy choices are too sensitive to the gasoline price.

 $\mu = (e_1 + e_2)/2$. Permits may be freely banked or borrowed across periods, and for simplicity I assume that the discount rate between periods is zero. The unconditional expected gasoline price in both periods is \bar{G} .

Let the standard μ be set optimally, and suppose for simplicity that the standard is guaranteed to bind. It must then be that the Taylor approximation for the derivative of the single-period private benefits function with respect to e is given by:

$$B_e(e,G) = \phi + B_{ee}(e-\mu) + B_{eG}(G-\bar{G}),$$
(81)

so that $B_e(\mu, \bar{G}) = \phi$.

In the first period, agents choose e_1 knowing the realization of G_1 (where $E[G_1] = \bar{G}$). This choice determines e_2 via $e_2 = 2\mu - e_1$. Let \tilde{G} denote the expectation of G_2 given G_1 ; i.e., $\tilde{G} = E[G_2|G_1]$. I parameterize this expectation by $\tilde{G} = \alpha G_1 + (1 - \alpha)\bar{G}$. Thus, $\alpha = 1$ represents an expectation that fuel price shocks are persistent, and $\alpha = 0$ represents an expectation that fuel price shocks are white noise around \bar{G} . Intermediate values of α denote a positively serially correlated AR(1) process for G, with a long-run mean of \bar{G} .

Now consider the agents' optimal choice of e_1 given G_1 . The total private benefits function, noting that $e_2 = 2\mu - e_1$, is given by:

$$B(e_1, G_1) + E[B(e_2, G_2)] = \phi(e_1 - \mu) + \frac{1}{2} B_{ee}(e_1 - \mu)^2 + B_{eG}(e_1 - \mu)(G_1 - \bar{G}) + \phi(\mu - e_1) + \frac{1}{2} B_{ee}(e_1 - \mu)^2 + B_{eG}(\mu - e_1)(\tilde{G} - \bar{G}) = B_{ee}(e_1 - \mu)^2 + B_{eG}\left((e_1 - \mu)(G_1 - \bar{G}) + (\mu - e_1)(\tilde{G} - \bar{G})\right)$$
(82)

The FOC is then:

$$2B_{ee}(e_1 - \mu) + B_{eG}(G_1 - \tilde{G}) = 0$$

$$\leftrightarrow 2B_{ee}(e_1 - \mu) = -B_{eG}(G_1 - \alpha G_1 - (1 - \alpha)\bar{G})$$

$$\leftrightarrow e_1 - \mu = \frac{-B_{eG}}{2B_{ee}} \left(G_1 - \alpha G_1 - (1 - \alpha)\bar{G}\right).$$
(83)

Thus, if fuel prices are persistent (consistent with section 2.2) so that $\alpha = 1$, then agents' optimal choice of e_1 is μ regardless of the realization of G_1 . Thus, fuel economy matches the standard in both periods, and there is no banking or borrowing. In contrast, if fuel prices are mean reverting so that $\alpha < 1$, then e_1 will be greater (less) than μ if G_1 is less (greater) than \overline{G} , signifying borrowing (banking).