Do Automotive Fuel Economy Standards Increase Rates of Technology Change?

Entry for 2012 Dennis J. OBrien USAEE Best Student Paper Award

Don MacKenzie – Massachusetts Institute of Technology

First version: April 2012
This version: July 2012

Abstract

Corporate Average Fuel Economy (CAFE) standards are popularly presented as technology-forcing regulations, which will accelerate the deployment of efficiency-improving technologies in new automobiles. It is not clear, however, that such regulations actually spur more advanced technology; fuel economy gains may also come through sacrificing other vehicle attributes, such as acceleration performance, size, or features. In this paper I test whether a binding CAFE regulation increases the rate of technology deployment within a firm's fleet of automobiles. I build on recent applications of a product characteristics framework to quantify technological change in automobiles. I use fleet and firm-level regulatory compliance data to identify changes in the rate of technology improvements when a manufacturer's fleet is more tightly constrained by a CAFE standard. In a variety of panel regression specifications, I found little to no significant change in the the rate of technology improvement when fleets were more tightly constrained by a CAFE standard. This was the case for both technology improvement among the menu of cars offered for sale, and the sales-weighted mix of all cars sold. The failure to find a significant effect of the standards on technology change does not preclude the possibility of such an effect, and I discuss several limitations to my results in this paper.

Don MacKenzie, Engineering Systems Division, Massachusetts Institute of Technology, 77 Massachusetts Avenue, 31-141, Cambridge, MA 02139. E-mail: dwm@mit.edu.

I thank Valerie Karplus and Chad Hazlett for helpful comments on drafts on this paper. Any errors are the sole responsibility of the author.

I. Introduction

The development and diffusion of technologies that enhance energy efficiency is a topic of recurring public policy interest, as such technologies offer the promise of reducing resource consumption and externality generation without sacrificing the utility realized by consumers. As Jaffe & Stavins (1994) have argued, understanding the effectiveness of different policies for stimulating the diffusion of energy-saving technologies is essential to sound policymaking in this area. To better understand these issues, those authors identified "two inextricably linked questions: What factors influence the rate of adoption of energy-conserving technologies; and what types of public policy can accelerate their diffusion?" In this paper, I present an empirical analysis of the latter type: testing whether binding automotive fuel economy standards increase the rate of technology change in U.S. cars.

Corporate Average Fuel Economy ("CAFE") standards are the primary policy tool used to curtail petroleum consumption and greenhouse gas emissions from light-duty vehicles (i.e. cars and light trucks) in the U.S. The production and use of fuels for these vehicles accounted for nearly one half of all petroleum consumption and one quarter of greenhouse gas emissions in the U.S. in 2010. CAFE standards require that each fleet of light-duty vehicles sold in the U.S. by each manufacturer in each year meet a minimum average level of fuel economy (sales-weighted, harmonically averaged). Firms have been permitted to bank credits earned through overcompliance, and to borrow credits against promised future overcompliance, for three years.[1]

CAFE standards are commonly characterized in both policy debates and in the literature as "technology-forcing" regulations that accelerate the deployment of efficiency-improving technologies; see, for example, Kleit (2004), and NESCAUM (2008). Numerous studies have estimated the effects of CAFE using economic models that are premised on an assumption that the regulations are technology-forcing. To provide just two examples, Fischer et al. (2007) and Kleit (2004) both permit the level of technology adoption to vary in response to prices or regu-

---

[1] Recently, some additional flexibility mechanisms have been introduced to the CAFE program, but were not in effect during the years covered in this analysis.

lations. Despite the prevalence — and acknowledged plausibility — of this premise, empirical assessments of regulation-induced technological change in automobiles are scarce.

Past work has affirmed the role of CAFE standards in stimulating fuel economy increases. Greene (1990) concluded that the standards were more important than fuel prices in determining fuel economy levels over the first twelve years of the CAFE program. However, this is not the same as saying that the standards spurred more rapid changes in technology. Fuel economy gains may also have come at the expense of other vehicle attributes, such as acceleration performance (Knittel 2011). Or, in the parlance of Newell et al. (1999), it is possible that CAFE standards drove changes in the direction, but not the rate, of technological progress in automobiles.

To determine whether standards have affected the rate of technology change, it is necessary to shift away from a focus on fuel economy (which is but one of many vehicle attributes affected by technological change) and toward a focus on some measure of technological change *per se*. One such approach relies on patent counts as a measure of innovation (Popp 2002). Crabb & Johnson (2010) recently applied this method to energy-saving innovations in U.S. automobiles, concluding that automotive energy efficiency patents are not responsive to CAFE standards. Attempting to square this conclusion with the literature that has found an effect of CAFE, the authors speculate that "Perhaps it is price that drives innovation, as we have shown, and regulation encourages only the final step of adoption by automobile manufacturers." (Crabb & Johnson 2010) This points to another approach to measuring technology change: one that uses a product characteristics space to measure the diffusion of new technologies.

Newell et al. (1999) measured the rate and direction of technological change in air conditioners and water heaters using the product characteristics approach. They included energy prices and policy stringency as regressors, and concluded that these factors influenced the direction, but not the rate, of technology change in air conditioners. They found no significant effects on either the rate or direction of technological change in gas water heaters. Knittel

(2011) recently employed a similar product characteristics approach to estimate the average fleet-wide technology improvements in new automobiles, finding that"Technological progress was most rapid during the early 1980s, a period where CAFE standards were rapidly increasing and gasoline prices were high." Although this suggests that faster rates of technology change are associated with faster rates of standard-tightening (and also higher gas prices), no attempt was made to estimate the effect of CAFE on the rate of technological change, nor to distinguish it from the effect of gasoline prices (Knittel 2011).

In this paper, I build on past investigations by extending the product characteristics framework for automobiles to estimate technological change at the level of individual manufacturers' car fleets. I then use these fleet-specific estimates to test whether being constrained by a CAFE standard is associated with higher rates of technology change in the constrained fleet. In so doing, I exploit variation in the gap between actual fuel economy and applicable standards, both between firms and within the same firm over multiple years, to identify the magnitude of the technology-forcing effect of a binding CAFE standard.

The paper proceeds as follows: in the next section, I discuss my empirical strategy and how it is informed by the structure of the CAFE program; in Section III I discuss the data sources used; in Section IV I present results; and in the final section, I offer some conclusions.

II. Empirical strategy

Several features of the CAFE program are relevant to my empirical strategy. First, CAFE standards are applied separately to each manufacturer's fleets of domestic passenger cars, imported passenger cars, and light trucks, each of which must independently meet an applicable standard.[2] Light trucks are subject to generally looser standards, and have in the past been subdivided into numerous subcategories for compliance purposes. The time-varying classification of light trucks considerably complicates the matter of matching a group of trucks with

---

[2] Many firms produce both a domestic car fleet and an import car fleet, each of which must meet the applicable standard.

a certain rate of technology change to the relevant CAFE fleets, and so the present analysis focuses only on cars. As noted earlier, new provisions have been introduced to allow credit trading between fleets, but these were not in place for the historical period considered in this work.

A second important feature of the CAFE program is its penalty structure. Firms that fail to meet a standard must pay a fine proportional to the number of vehicles in the noncompliant fleet and to the number of miles per gallon (mpg) by which that fleet missed the standard. Although there are penalties for failing to meet the standard, there is no bonus for exceeding it (doing so can generate credits to offset future noncompliance, but until recently these credits could not be sold and had to be used within 3 years). As such, we may expect that firms will respond differently to being above the standard than to being below it.

Finally, it has been documented by Jacobsen (2012) that different firms behave differently in the face of the CAFE standard. One set of firms (the Detroit Three — GM, Ford, and Chrysler) appears to treat it as a binding constraint. Another set of firms (mainly European) routinely violates the standard and pays the penalty, and thus is largely unaffected by the standard (since the penalty rate is extremely low relative to the cost of technical solutions). A final set of firms (mainly Asian) routinely exceeds the standards, and so is also unaffected by them. Thus, we may expect a different response to the standard for different classes of firms. We are especially interested in the response of those firms that treat the standard as a binding constraint, as we may reasonably expect little or no response from the other sets of firms.

A. Fleet-level rates of technological change

Knittel (2011) applied a product-characteristics framework in order to estimate yearly averages for technological progress in automobiles. Using an adaptation of this methodology, I first generate estimates of the rate of technology change for each fleet in each year. Whereas Knittel estimated fixed effects for each year and interpreted these as a measure of cumulative technological change since his base year, I estimate fixed effects for each combination of firm-fleet-year (e.g. GM's import car fleet in 2002). At the same time, I relaxed Knittel's assumption of constant coefficients on weight and power across all firms, allowing for firm-specific coefficients on these variables. Thus, I modeled the fuel economy of model i in fleet j from firm k in year t as:

$$\ln mpg_{ijkt} = T_{jkt} + \beta_{1k} \ln W_{ijkt} + \beta_{2k} \ln P_{ijkt} + Dsl_{ijkt}\beta_3 + Man_{ijkt}\beta_4 + Man_{ijkt} * t\beta_5 + \epsilon_{ijkt} \quad (1)$$

Where T is a set of fixed effects for the fleet-firm-year, W is car weight, P is engine peak power, Dsl is a dummy variable indicating whether a car has a diesel engine, Man is a dummy variable indicating whether it has a manual transmission, and ϵ is an i.i.d. random error term.[3] I took the first differences in the fixed effects T_{jkt} in order to obtain my dependent variable: the year-over-year change in the technology deployed within each fleet by each firm, denoted ΔT_{jkt}.

$$\Delta T_{jkt} = T_{jk,t+1} - T_{jkt} \quad (2)$$

When Equation 1 is estimated using an unweighted regression, the fixed effects can be interpreted as a measure of the average technological sophistication of the cars offered for sale in a given fleet/firm in a given year. Alternatively, Equation 1 can be estimated using weighted least squares, where the weights are equal to the sales of each observed car. In the latter case, the fixed effects may be better interpreted as the average technological sophistication of the cars sold by a given fleet/firm in a given year. I estimated Equation 1 using both unweighted and sales-weighted least squares, and estimated the effects of CAFE on both of these measures of technological change.

Over the years and fleets considered in this work, the mean value of the (unweighted) annual change in technology potential was 0.017, with a standard deviation of 0.032. This means that technological improvements could have increased fuel economy of cars offered for sale

---

[3] This closely resembles Knittel's model specification #3, but with the year fixed effects replaced by firm-fleet-year fixed effects.

by an average of 1.7% annually, conditional on power, weight, transmission types, and diesel share remaining unchanged. On a sales-weighted basis, this value was 0.015, with a standard deviation of 0.040. This indicates that the technological improvements, averaged across the mix of cars actually sold, could have increased average fuel economy by an average of 1.5% per year, conditional on power, weight, transmission types, and diesel share remaining unchanged.

B. Effects of CAFE on technological change

The CAFE penalty system creates an incentive structure in which penalties proportional to the amount by which a fleet falls short of the standard, but fleets that exceed the standard are neither penalized nor rewarded. I therefore define an independent variable which I call shortfall as the gap between the current year's actual fuel economy and the next year's required fuel economy:

$$S_{jkt} = (Std_{jk,t+1} - MPG_{jkt}) \tag{3}$$

The shortfall therefore captures the stringency of a CAFE constraint, in that it represents the amount by which a fleet would have to increase its fuel economy over the coming year to be compliant with the next year's standard. Since there is a penalty for missing the standard but no reward for exceeding it, I also define a binary treatment variable indicating whether or not the fleet is CAFE-constrained (i.e. whether it has to improve its fuel economy to meet the new standard in the coming year):

$$D_{jkt} = 1\{S_{jkt} < 0\} \tag{4}$$

I then flexibly model the effect of the shortfall and the CAFE-constrained indicator on the rate of technology change a set of panel regression models utilizing fixed effects for both year and fleet, the general form of which is:

$$\Delta T_{jkt} = \delta_t + \alpha_0 D_{jkt} + \alpha_1 S_{jkt} + \alpha_2 S_{jkt} D_{jkt} + \mu_{0jk} + \mu_{1jk} t + \mu_{2jk} t^2 + X_{jkt}\beta + \epsilon_{jkt} \tag{5}$$

In the above equation, δ_t is a vector of fixed effects for each year; μ_{0jk} is a vector of fixed effects for each fleet; μ_{1jk} and μ_{2jk} allow for fleet-specific time trends; S_{jkt} and D_{jkt} are as defined above; and X_{jkt} is a vector of covariates. I also investigated a specification that included interactions of S_{jkt} and D_{jkt} with a dummy variable for the Detroit Three, allowing for a different response to the standard for this subset of companies.

In some specifications of the model, I also included as covariates average weight, average size (interior volume), average engine peak power, average fuel consumption (gallons per mile), fraction of small cars (mini-compacts, subcompacts, and two-seaters), fraction of cars with 4-wheel drive, fraction of wagons, and fraction of convertibles. The rationale for including these covariates is that the composition of a firm's fleet might conceivably affect the ease with which new technologies can be adopted.

C. Standard errors

The number of fleets in the data set was 29, meaning that conventional methods of estimating clustered standard errors were inappropriate. I therefore used a double bootstrapping procedure to estimate the standard errors. To do this, I first resampled the fleets, with replacement. Next, from the data set created by resampling the fleets, I resampled blocks of years, with replacement, as in block bootstrapping. I used non-overlapping blocks with length 3 years. Finally, I estimated the parameters of interest using the resampled data, repeating the above steps 2000 times to simulate a distribution of parameter values.

D. Identification

It is worth discussing briefly the ability of the panel regression specification to correctly identify the effect of the binding CAFE standard on the rate of technological change. For the panel regression to identify the causal effects of the shortfall S and the state of being constrained by the standard D, we need the error terms to be independent conditional on the fixed effects and the covariates, i.e. $E[\epsilon_{jkt}|X_{jkt}, \delta_t, S_{jkt}, D_{jkt}, t] = 0$ for all t. In other words, we need there to be no unobserved confounders that are varying with time.

One possible reason that the panel regression may not correctly identify the causal effects is if the standards were set at levels that the firms deemed feasible. In this case, there could be an omitted variable, "firm's technology plan," that is correlated with increases in observed technology. If firms had influenced the level of the standard, then this technology plan variable could also be correlated with the next year's CAFE standard. An important point here is that the confounding effect of the technology plan variable would have to be varying over time. To the extent that the the stringency of regulations is shaped by the capabilities and plans of the industry as a whole, this confounding should be soaked up the fixed effects. However, CAFE standards have been set, at least some of the time, according to a "least capable manufacturer" heuristic, in which the standards are set so as not to be overly burdensome on any single firm (NHTSA 2006). This type of process could indeed lead to a situation in which the size of a firm's fuel economy shortfall is correlated with its future technology improvements.

One approach to dealing with this comes from the fact that after 1989, CAFE standards for cars did not change for the remainder of the years in this analysis. The standard for cars was constant at 27.5 mpg, which is essentially an arbitrary value. Since we know that technology continued to improve during this period — as evidenced by the continued changes in the year fixed effects from the fuel economy regression — we can conclude that the level of the CAFE standard in these years was not determined by the technology plans of the manufacturers. So, if we re-estimate our model(s) using only the data from 1989 onwards, we can be more confident that the results are not an artifact of CAFE standards being set based on manufacturer technology plans.

III. DATA

I used two principal data sources in this work. Data on the annual fuel economy performance and applicable CAFE standards came from CAFE compliance reports held by the National Highway Traffic Safety Administration (NHTSA), which administers the CAFE program. These were used to calculate whether a fleet was above or below the next year's standard, and by how much. Another database maintained by NHTSA provided model-level data on car attributes including fuel economy, weight, engine type and power, transmission type, drive type, and body style. I used these data to estimate the rates of technological progress in each year (the outcome of interest) as well as various covariates. I used data from the years 1978-2008, and omitted only manufacturers of limited-volume, specialty vehicles.

IV. RESULTS

The results of the series of panel regressions for predictors of technological change in cars offered for sale are shown in Table 1. The bootstrapped standard errors are reported in parentheses. If the CAFE standard had a forcing effect on technology, we would expect the coefficients on the shortfall variable, S, and on the shortfall–CAFE-constrained interaction term, $S * D$, to be positive and significant. However, from columns 1–4 in Table 1 it is clear that essentially none of the coefficients on the binary (constrained/unconstrained) or continuous treatment intensity (shortfall) variables are significant. This remained true regardless of the inclusion of firm-specific time trends and covariates, and of the restriction of the data to only those years when standards were arbitrarily fixed at 27.5 MPG.

As noted earlier, we would expect the effects of a CAFE standard to be larger for those firms (the Detroit Three) that treat the standards as a binding constraint. It is possible that the estimated effects in columns 1–4 were dampened by the inclusion of firms that tend to ignore their CAFE obligations. Therefore, I re-estimated the models in columns 1–4 while allowing for a different CAFE response for the Detroit firms. These results are shown in columns 5–8, and column 9 shows the results when only the Detroit firms were included in the data set. The coefficient on shortfall for Detroit firms, when they are constrained by the CAFE standard, was positive and significant in model specification 6. This is in line with predictions of technology-forcing, but the result was not robust to different model specifications and is not compelling on its own.

Automobile manufacturers generally make product planning decisions several years ahead. Therefore, it is plausible that being constrained by a CAFE standard in one year would not be sufficient to alter the average technological sophistication of the mix of cars offered in the next year. One approach to dealing with this is to recognize that being constrained by the standard might alter the mix of products actually sold, an effect which would show up in the (sales-weighted) measure of technological sophistication. Table 2 shows the same set of models as before, but with the sales-weighted technological progress as the dependent variable. The results are slightly different than in Table 1. As can be seen in columns 1 and 2, the estimated coefficients on the shortfall variable are significant, but the estimate is not stable across these different specifications and becomes insignificant when covariates are included.

An alternative approach to accommodating the dynamics of the product cycle is to estimate rates of technological change and stringency of the CAFE constraint over periods longer than one year. Table 3 summarizes the results of several regressions like those already presented. However, these results were based on technological change and CAFE shortfall measured over three-year intervals instead of one-year intervals. As with the earlier results, no significant effect of CAFE on the rate of technological change is evident.

Table 1: Results for year over year technological change in cars offered for sale (unweighted).

	1	2	3	4	5	6	7	8	9
CAFE Constrained	0.0035	0.0046	0.0068	0.0058	-0.0012	0.0081	0.0059	-0.0091	-
	(0.0086)	(0.0113)	(0.0116)	(0.0167)	(0.0142)	(0.0233)	(0.0286)	(0.0235)	-
CAFE Shortfall	0.0019	0.0022	0.0030	-0.0001	0.0012	0.0038	0.0034	0.0009	-
	(0.0015)	(0.0023)	(0.0049)	(0.0085)	(0.0013)	(0.0027)	(0.0053)	(0.0092)	-
Constrained * Shortfall	0.0008	0.0058	0.0042	-0.0024	0.0012	0.0042	0.0028	-0.0054	-
	(0.0037)	(0.0052)	(0.0072)	(0.0283)	(0.0089)	(0.0233)	(0.0123)	(0.0250)	-
Detroit * CAFE Constrained	-	-	-	-	0.0008	-0.0017	0.002	0.0196	0.0103
	-	-	-	-	(0.0202)	(0.0027)	(0.0056)	(0.0570)	(0.0555)
Detroit * CAFE Shortfall	-	-	-	-	0.0022	-0.0046	-0.0025	-0.0046	0.0118
	-	-	-	-	(0.0038)	(0.0233)	(0.0286)	(0.0119)	(0.0291)
Detroit * Constrained * Shortfall	-	-	-	-	0.0019	0.0065*	0.0046	0.0109	0.0115
	-	-	-	-	(0.0168)	(0.0027)	(0.0053)	(0.0684)	(0.0679)
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

* Significant at the 0.05 level

Table 2: Results for year over year technological change in mix of cars sold (sales-weighted).

	1	2	3	4	5	6	7	8	9
CAFE Constrained	0.0001	-0.0003	0.0036	0.0078	-0.0058	0.0021	-0.0004	-0.0036	-
	0.0105	0.0132	0.0126	0.1567	0.0165	0.0195	0.0372	0.0176	-
CAFE Shortfall	0.0036*	0.0066*	0.0067	0.0042	0.0032	0.0032	0.0032	0.0032	-
	0.0018	0.0033	0.0068	0.0161	0.0018	0.0018	0.0018	0.0018	-
Constrained * Shortfall	0.0008	0.0041	0	-0.0027	0.0011	0.0023	-0.0027	-0.0057	-
	0.0049	0.0067	0.0078	0.4362	0.0093	0.0146	0.0125	0.0202	-
Detroit * CAFE Constrained	-	-	-	-	0.0042	0.0024	0.0061	0.0154	0.0063
	-	-	-	-	0.1097	0.0057	0.0372	0.0875	0.0353
Detroit * CAFE Shortfall	-	-	-	-	0.001	-0.0065	-0.004	-0.0054	0.0118
	-	-	-	-	0.0032	0.0195	0.0068	0.01	0.0201
Detroit * Constrained * Shortfall	-	-	-	-	0.0031	0.0075	0.0091	0.0116	0.0117
	-	-	-	-	0.0738	0.0037	0.0125	0.0777	0.027
Firms Included	All	All	All	All	All	All	All	All	Detroit Only
Years Included	All	All	All	1989 on	All	All	All	1989 on	All
Firm-Specific Time Trends	No	Yes	Yes	Yes	No	Yes	Yes	Yes	Yes
Covariates Included	No	No	Yes	Yes	No	No	Yes	Yes	Yes

* Significant at the 0.05 level

Table 3: Results for technological change in cars offered for sale (unweighted), over three-year intervals.

	1	2	5	6
CAFE Constrained	0.0017	-0.0123	-0.0085	-0.0421
	0.0299		0.0524	0.0503
CAFE Shortfall	0.0021	0.0053	0.0017	0.0045
	0.0034	0.0167	0.004	0.0503
Constrained * Shortfall	0.005	0.011	0.0126	0.0187
	0.0112		0.0207	0.0503
Detroit * CAFE Constrained	-	-	0.0228	0.0493
	-	-	0.0816	0.0503
Detroit * CAFE Shortfall	-	-	0.0072	0.0035
	-	-	0.0121	0.0503
Detroit * Constrained * Shortfall	-	-	-0.0311	-0.0303
	-	-	0.0429	0.0503
Firms Included	All	All	All	All
Years Included	All	All	All	All
Firm-Specific Time Trends	No	Yes	No	Yes
Covariates Included	No	No	No	No

* Significant at the 0.05 level

V. Conclusions

In this work, I found little to no evidence that either the state of being constrained by a CAFE standard in a given year, or the stringency of that constraint, increases the rate of technological change in the coming year. In the few cases where significance was found, the results were not robust to alternative specifications. This held true when the rate of technology change was defined in terms of (1) the year to year change in the technology of cars offered for sale, (2) the year to year change in the sales-weighted average technology level, and (3) the change in technology of cars offered for sale over three-year intervals.

Although little prior empirical work has been reported on the effects of CAFE standards on rates of technology change in automobiles, the negative results reported here do run counter to a good deal of well-grounded microeconomic modeling. This modeling work (for example, Kleit (2004), Fischer et al. (2007), Shiau et al. (2009)) shows that there are good reasons to expect that a CAFE standard would lead to faster technology adoption. Therefore, several caveats to the present results are in order, and the absence of proof should not be misconstrued as proof of absence of an effect of CAFE on technology change.

The failure to find a significant effect of the standards on technology change does not preclude the possibility of such an effect, and there are at several compelling reasons why an effect might not show up in my model. First, to the extent that standards might be affecting technology improvements across the industry as a whole, such effects would be soaked up in the time fixed effects included in my specifications. In general, these effects would be difficult to identify separately from effects of gasoline prices and political support for CAFE standards, since all of these variables are likely to move together over time. Second, even during the period covered in this analysis, manufacturers had the ability to bank and borrow credits if they over- or under-complied with CAFE in a given year. This would be expected to dampen the firm's response to a CAFE shortfall in any single year. Finally, there is a considerable amount of noise in the estimates of technology change, and the standard errors of the coefficient estimates are substantial. It is possible that a larger sample could reveal a significant effect of CAFE, but we are already using all of the available years and major manufacturers for which data are available.

Notwithstanding the above caveats, it is still interesting that there was no evidence for binding CAFE standards increasing rates of technological change. Taken with the results of Greene (1990), which found that CAFE significantly influenced fuel economy levels, the results reported here suggest a pattern similar for cars to that reported by Newell et al. (1999) for air conditioners: namely, that standards affect the direction, but not the rate, of technological change. Opportunities remain to apply the product characteristics framework to automobiles in order to test this hypothesis more thoroughly.

REFERENCES

Crabb, J. & Johnson, D. (2010), 'Fueling innovation: The impact of oil prices and cafe standards on energy-efficient automotive technology', *The Energy Journal* **31**(1), 199–216.

Fischer, C., Harrington, W. & Parry, I. (2007), 'Should automobile fuel economy standards be tightened?', *The Energy Journal* **28**, 1–29.

Greene, D. (1990), 'CAFE or price?: An analysis of the effects of federal fuel economy regulations and gasoline price on new car mpg, 1978-89', *The Energy Journal* **11**(3), 37–57.

Jacobsen, M. (2012), Evaluating us fuel economy standards in a model with producer and household heterogeneity. `http://econ.ucsd.edu/~m3jacobs/Jacobsen_CAFE.pdf`.

Jaffe, A. & Stavins, R. (1994), 'Energy efficiency investments and public policy', *The Energy Journal* **15**(2), 43–43.

Kleit, A. (2004), 'Impacts of long-range increases in the fuel economy (cafe) standard', *Economic Inquiry* **42**(2), 279–294.

Knittel, C. R. (2011), 'Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector', *American Economic Review* **101**(7), 3368–99.

NESCAUM (2008), 'Northeast States for Coordinated Air Use Management: Submission to federal docket id no. NHTSA-2008-0089', `www.nescaum.org/documents/nescaum-cafe-comment-ltr-final-070108.pdf`.

Newell, R., Jaffe, A. & Stavins, R. (1999), 'The induced innovation hypothesis and energy-saving technological change', *The Quarterly Journal of Economics* **114**(3), 941–975.

NHTSA (2006), 'Average fuel economy standards for light trucks model years 20082011: Final rule', Federal Register, 71, 66 / Thursday, April 6, 2006 / Rules and Regulations.

Popp, D. (2002), 'Induced innovation and energy prices', *The American Economic Review* **92**(1), 160–180.

Shiau, C., Michalek, J. & Hendrickson, C. (2009), 'A structural analysis of vehicle design responses to corporate average fuel economy policy', *Transportation Research Part A: Policy and Practice* **43**(9), 814–828.