

Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards

Web Appendix

Soren T. Anderson and James M. Sallee

A Appendix

In this appendix, we extend our model to two periods. This allows us to model banking and borrowing and introduce expectations over demand shocks. We use a two-period model for the sake of parsimony; all of our conclusions hold in a model with a finite or infinite number of periods. The interpretation of the shadow price in our static model extends directly to the two-period model. As before, the shadow price we estimate using our loophole methodology will equal the shadow price in the first period. Furthermore, the first period's shadow price will equal the expected shadow price in the second period. Thus, we can interpret the empirical estimates in this paper as both the current and the expected future shadow price of CAFE in a dynamic model with uncertainty.

A.1 Two-period model

In the static model, the demand function was known with certainty. To introduce uncertainty, we now rewrite demand as a function of prices, mileage, and an exogenous shock g_t . The exogenous shock could come from various sources, but we are interested primarily in gasoline prices.

In the static model, we assumed that automakers simultaneously choose mileage, flexible-fuel shares, and prices. In reality, while automakers are able to adjust prices continuously, they must make many engineering choices before production begins. Once these choices are made, however, automakers have considerable flexibility in choosing which options to include

on a vehicle, and they are able to switch between different engine sizes, transmissions, and other characteristics at relatively short notice. For example, Honda must design a 4-cylinder Accord and a 6-cylinder Accord well in advance of sales, but it can adjust the proportion of Accords that have a 4-cylinder engine throughout the model year. The same is true of flexible-fuel capacity. Our understanding of the production process is that firms must incur the fixed cost of designing the flexible-fuel version of a model in advance of sales, but that firms are able to adjust the share of vehicles fitted with flexible-fuel components during the model year.¹

To capture this timing, we assume that firms choose which models will have a flexible-fuel option, as well as the mileage of all models, in the prior period before observing any demand shock. After observing the shock, firms can adjust flexible-fuel shares and prices. To model the choice of which models get the flexible-fuel option, we introduce a new choice variable, I_{tj} , which equals 1 if the firm has already paid the fixed costs of enabling flexible-fuel production for model j in period t , and 0 if it has not. (In the static model, we used a notational shortcut to model the same choice.)

Mathematically, a firm arrives at period 1 in our model having inherited I_{1j} and m_{1j} from an unmodeled prior period. The firm chooses prices p_{1j} and flexible-fuel shares θ_{1j} simultaneously, having observed the demand shock g_1 . The firm also chooses mileage m_{2j} and whether to enable flexible-fuel capacity I_{2j} for period two. In the second period, the firm inherits I_{2j} and m_{2j} , observes gasoline prices, and then chooses prices p_{2j} and flexible-fuel shares θ_{2j} .

To model banking and borrowing rules, we require that the automaker comply with

¹We tested this assumption empirically. Using our micro transaction data, we regressed the flexible-fuel dummy variable on monthly gasoline prices, controlling for model-specific fixed effects and dummy variables indicating the number of months since the beginning of the model year. The results indicate that, for any given model, a \$1 increase in gasoline prices correlates with a 46 percentage-point reduction in flexible-fuel shares. This economically and statistically significant effect has the intuitive sign: an increase in gasoline prices will generally shift demand toward smaller, more efficient vehicles, meaning that firms need fewer flexible-fuel vehicles to achieve compliance. Adding state controls reduces the effect to a statistically insignificant 1.5 percentage points, which suggests that the large response above is likely concentrated in particular states.

CAFE on average over both periods. We use f_t to denote the CAFE credits or debts that a firm inherits at the beginning of each period. In the first period, the CAFE constraint is an equation of motion that determines the firm's balance of credits at the beginning of the second period. In the second period, the CAFE constraint requires that the firm end the period with a nonnegative balance of credits. The backstop constraint remains unchanged from our static model. Since the firm may not carry excess flexible-fuel credits between periods, the backstop constraint applies separately in each period.

We begin by writing the second-period maximization problem, defining V as the value function:

$$\begin{aligned}
V_2(f_2, m_2, I_2) &= \max_{p_2, \theta_2} \sum_{j \in \mathcal{M}} \left(p_{2j} - c_j(m_{2j}) - \alpha_j \theta_{2j} \right) q_j(p_2, m_2, g_2) - \sum_{j \in \mathcal{M}} I_{2j} \cdot F_j \\
\text{s.t. } f_2 &+ \left[\left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1 - \theta_{2j}(1 - \beta)}{m_{2j}} \right)^{-1} - \sigma_2 \right] Q \geq 0 \\
&\left[\left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1 - \theta_{2j}(1 - \beta)}{m_{2j}} \right)^{-1} - \left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1}{m_{2j}} \right)^{-1} \right] Q \leq \phi_2 Q \\
&(1 - I_{2j})\theta_{2j} = 0 \quad \forall j,
\end{aligned}$$

where we have suppressed the arguments of the second period's demand functions (q_j and Q) in all but the first line for convenience. The constraint that $(1 - I_{tj})\theta_{tj} = 0$ forces flexible-fuel shares to zero for any model not enabled with flexible-fuel capacity. Everything else is the same as the static model, with the addition of g and f .

Next, we write the first-period maximization problem, using $V_2(f_2, m_2, I_2)$ as an abbre-

viation for the second-period value function:

$$\begin{aligned}
V_1(f_1, m_1, I_1) = & \max_{p_1, \theta_1, m_2, I_2} \sum_{j \in \mathcal{M}} \left(p_{1j} - c_j(m_{1j}) - \alpha_j \theta_{1j} \right) q_j(p_1, m_1, g_1) - \sum_{j \in \mathcal{M}} I_{1j} \cdot F_j \\
& + \delta E_1 \left[V_2(f_2, m_2, I_2) \right] \\
\text{s.t. } f_1 + & \left[\left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1 - \theta_{1j}(1 - \beta)}{m_{1j}} \right)^{-1} - \sigma_1 \right] Q = f_2 \\
& \left[\left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1 - \theta_{1j}(1 - \beta)}{m_{1j}} \right)^{-1} - \left(\sum_{j \in \mathcal{M}} \frac{q_j}{Q} \cdot \frac{1}{m_{1j}} \right)^{-1} \right] Q \leq \phi_1 Q \\
& (1 - I_{1j}) \theta_{1j} = 0 \quad \forall j,
\end{aligned}$$

where E_1 is the expectation in period 1 over the stochastic shock g_2 , δ is a discount factor, and we have again suppressed the arguments of the first period's demand functions for convenience. A standard Lagrangian can be written for each period's maximization problem. We label shadow prices on the CAFE constraints in periods 1 and 2 as λ_1 and λ_2 , while the corresponding shadow prices on the backstop constraints are μ_1 and μ_2 .

In period one, the first-order condition of the Lagrangian with respect to the flexible-fuel share of vehicle k (for which we assume $I_{1k} = 1$) is directly analogous to our static condition:

$$-\alpha_k + (\lambda_1 - \mu_1) \frac{1 - \beta}{m_{1k}} M_1^2 = 0. \tag{14}$$

In turn, this leads to an expression for the marginal compliance cost per vehicle that is exactly the same as our static result, with the only notable difference being the time subscript on the shadow price.

The automaker will also attempt to equate marginal compliance costs across time. This is made clear by differentiating the Lagrangian with respect to f_2 , which yields the Euler

equation characterizing dynamic efficiency:

$$\begin{aligned}\lambda_1 &= \delta E_1 \left[\frac{\partial V_1}{\partial f_2} \right] \\ &= \delta E_1 \lambda_2.\end{aligned}$$

That is, the shadow price in period 1 is equal to the expected shadow price in period 2. Thus, our estimates, which we derived from the static model, can be interpreted in a dynamic context as both the current and the expected future shadow price. In any given period, the cost of exploiting the flexible-fuel loophole will reveal the shadow price on the CAFE constraint, assuming the backstop constraint is not binding. This shadow price will equal the marginal cost of compliance based on changing current vehicle prices. Furthermore, the loophole will reveal the expected shadow price in the next period, which will equal expected marginal compliance costs based on modifying next period's models to be more efficient.

A.2 Discussion

Recasting our static model in this dynamic stochastic framework illustrates several key points. First and most importantly, this formulation makes clear that banking and borrowing do not change the interpretation of our estimates. The estimates we report can be interpreted as both the current and the expected future marginal cost per vehicle of complying with CAFE standards.²

Second, this framework allows us to understand and interpret the stability of our parameter estimates over time. Note that the first-order condition with respect to flexible-fuel shares in the second period produces an expression that is analogous to equation (14) above. Even if demand shifts unexpectedly between periods 1 and 2, if the firm is at an interior solution in both periods, the shadow price will remain roughly stable over time. The shadow

²Note that this model still represents a simplification of banking and borrowing, since credits actually expire after three years. A model that includes these details would be cumbersome and would provide no substantial difference of interpretation.

price will change only to the extent that $\alpha \cdot m / (1 - \beta) / M^2$ changes. Since α and β are fixed parameters, the only change will come from changes in the mileage statistics, which change very little from year to year. Thus, as long as firms are at an interior solution, the realized shadow price will remain stable in the face of small shifts in demand. This does not mean that the shadow price would be stable if there were no loophole. Rather, stability follows from the (nearly) constant marginal cost of exploiting the loophole in different periods.

Third, because firms in this framework are constrained by the flexible-fuel models chosen in the previous period, a firm might fail to exhaust the loophole, even though it would like to. One scenario is the following. An automaker plans to use only half the loophole next period, and thus enables flexible-fuel capacity on a model with low to moderate sales. Then, gasoline prices drop sharply, increasing demand for large, inefficient vehicles. The automaker would like to compensate by installing flexible-fuel capacity on many more vehicles than planned, but is unable, because sales of its flexible-fuel model are too low.

An alternative scenario involves the reverse. Suppose the automaker enables the flexible-fuel option on a large, inefficient model and expects to exhaust the loophole next period. Then, gasoline prices rise sharply, and demand for large vehicles plummets. The automaker fails to exhaust the loophole, but only because it is unable to sell as many large flexible-fuel vehicles as it had planned.

In either scenario, we would expect to see flexible-fuel *shares* on these models reach very high levels at the end of the model year. This is not evident in our data. In fact, the data indicate that flexible-fuel shares are quite stable on average (about two-thirds) over the last twelve months of production. Furthermore, when we examine flexible-fuel shares in the last several months of the model year separately by fleet and year, we see no obvious pattern to indicate that firms would have trouble further boosting flexible-fuel shares late in the year. This weighs heavily against any concern that automakers were often constrained in this way.