

MIT CEEPR

MIT Center for Energy and Environmental Policy Research

Working Paper Series

Anticipation and Environmental Regulation

KATHERINE RITTENHOUSE AND MATTHEW ZARAGOZA-WATKINS

FEBRUARY 2017

CEEPR WP 2017-004



MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Anticipation and Environmental Regulation

Katherine Rittenhouse and Matthew Zaragoza-Watkins*

February 2017

Abstract

When agents expect a regulation to change the relative price of new equipment, they may shift purchases forward to avoid compliance costs. In the context of new-vehicle emission standards, prior analyses have not considered this adjustment margin. We model the effects of anticipation on freight-truck sales and retirements, and test our theory's predictions empirically. Consistent with our predictions, we find evidence that anticipation caused a sales spike just before new emission standards took effect and a sales slump after implementation. Our findings have important implications for analysis of markets where agents can shift purchases in anticipation of new regulation.

JEL: C41, D92, Q58

I Introduction

In the United States, the transportation sector is the largest source of criteria pollution, and the second largest source of greenhouse gas (GHG) pollution, regulated under the Clean Air Act (CAA).¹ To help achieve NAAQS, EPA regularly updates pollution control standards for mobile and stationary sources.

*(Rittenhouse) Environmental Defense Fund. Email: krittenhouse@edf.org. (Zaragoza-Watkins) Environmental Defense Fund and Research Affiliate, MIT CEEPR. Email: mzaragoza@edf.org. We are grateful for helpful comments from Max Auffhammer, Severin Borenstein, Lucas Davis, Larry Karp, Michael Mehling, Reed Walker, and Catherine Wolfram. We thank seminar participants at EDF, Harvard, MIT, UC Berkeley, and U.S. EPA, and participants at the MIT CEEPR workshop and the AERE Summer Conference for helpful comments. We also thank MIT CEEPR for financial support in acquiring data. All remaining errors are our own.

¹The CAA requires the U.S. Environmental Protection Agency (EPA) to set National Ambient Air Quality Standards (NAAQS) for six categories of common air pollutants. These "criteria" pollutants, which are so named because EPA is required to set NAAQS based

In recent years, concerns about local air pollution, global climate change, and energy security have prompted EPA to adopt more stringent standards for new mobile sources, requiring Original Equipment Manufacturers (OEMs) to deploy control technology, at some cost.

Vintage-differentiated emission standards are important and frequently used tools for controlling pollution (Stavins, 2006). However, they are not theoretically optimal policy instruments;² generally, these regulations do not target pollution from all sources or through all emission pathways.³ When cost increases are unaccompanied by offsetting increases in private benefits, new-vehicle emission standards prompt motorists to operate their older, and higher-emitting, vehicles for longer than they otherwise would have, degrading the short-term environmental benefits of the policy (Gruenspecht, 1982).⁴ Similarly, forward-looking consumers who wish to avoid paying the incremental cost of complying with a new standard may choose to “pre-buy” a new vehicle during the period just before a regulation is implemented. While the incentive for consumers to purchase new vehicles ahead of regulation may be intuitively clear, the impacts on new-vehicle sales and the composition of the

on scientifically-supported health criteria, are particulate matter (PM), ground-level ozone (O₃), carbon monoxide (CO), sulfur oxides (SO_x), nitrogen oxides (NO_x), and lead (Pb). Of criteria pollutants, PM and ground-level O₃ present the broadest and most acute threats to human health and the environment. Transportation is also the largest source of volatile organic compounds, which, along with NO_x, are precursors to photochemical O₃.

²Vintage-differentiated emission standards may be cost-effective in the short run, as retrofitting existing units is often more expensive than requiring that new units be built to a certain standard. However, these regulations create an incentive for operators of older, higher-emitting units to delay retirement. This dynamic inefficiency erodes long-term cost-effectiveness, and may reduce the environmental effectiveness of the standard. In some cases, standards may even backfire if increased lifetime emissions from older units outweigh emissions reductions from new units. That these policies cause a delay in fleet turnover is seen theoretically and empirically in multiple fields of environmental regulation. Two prominent examples include regulation of power generators under the Clean Air Act’s New Source Review program (Maloney and Brady, 1988), and automobile emissions, safety, and fuel-economy standards (Gruenspecht, 1982; Jacobsen and Van Benthem, 2015).

³Optimal pollution control could be achieved by imposing a series of pollution taxes equal to the marginal social damage from emitting a quantity of the relevant pollutant. Pigouvian taxes would induce vehicle operators to drive fewer miles and buy less-polluting vehicles, OEMs to market vehicles with better pollution control equipment, and refiners to produce cleaner fuels (Fullerton and West, 2002). Capturing each of these adjustment margins would yield the broadest and most cost-effective suite of emissions reductions.

⁴Known as the “Gruenspecht” effect, this phenomenon may be particularly relevant for criteria pollutant standards, where the regulation increases production costs for OEMs, and thus private costs to vehicle purchasers, while socializing the benefits (i.e., cleaner air). This is in contrast to, for example, GHG and fuel-economy standards, which provide significant direct consumer benefits in the form of fuel savings, resulting in lower operating costs.

fleet, and the implications for the environmental effectiveness of these policies are essentially unstudied.⁵

To investigate the incentives created by new-vehicle standards, the impact of those incentives on the new-vehicle sales cycle, and the implications for the effectiveness of new-vehicle standards, we analyze the market for new Class-8 heavy-duty vehicles (HDVs or trucks). We choose trucks for several reasons. They are important to the economy; the trucking industry carried 71% of the value and 70% of the weight of U.S. freight in 2007 (U.S. Department of Transportation, 2010). Heavy-duty trucks are responsible for a significant portion of the nation’s air pollution from transportation. EPA has developed and implemented multiple rounds of new-truck emission standards, creating repeated opportunities to analyze the short-run impact of standards on new-truck sales. The trucking industry also has useful properties for analyzing strategic purchasing behavior. New trucks are relatively homogeneous durable goods, which are elastically supplied by competitive OEMs. Decisions about when to buy and retire trucks are generally made by firms seeking to maximize profit. There is relatively free entry into the freight transport market, which is where trucks are primarily used.

We address four specific questions. How does the anticipation of regulation affect the pattern of new-truck sales? How does the pattern of new-truck sales affect the pattern of used-truck retirements? How do purchasing and retirement patterns affect the environmental benefits of standards? Empirically, have recent regulations caused firms to pre-buy trucks?

We begin by developing a dynamic model of a competitive freight truck market, where firms incorporate new-truck prices, operating costs and freight rates (i.e., operating revenue) into their purchasing and retirement decisions. Operating costs of individual trucks are assumed to monotonically increase with age, motivating cycles of retirement and replacement. We derive necessary conditions for trucks to enter and retire from the freight market, and calculate comparative statics for changes in upfront and operating costs. As in many models of capital turnover in competitive markets, trucks enter when

⁵While the pre-buy phenomenon is described in the gray literature – see, for example, Calpin and Plaza-Jennings (2012) – with the exception of Lam and Bausell (2007), this form of “anticipation” is little studied in the mainstream academic literature on the effects of emissions standards. Our work is, however, closely related to a growing literature which examines the effect of anticipated changes in sales and consumption taxes on purchase patterns (see, for example, D’Acunto, Hoang and Weber (2016), Cashin and Unayama (2015) and Crossley, Low and Sleeman (2014)). Most recently, Coglianesse et al. (2016) find that consumers anticipate gasoline tax changes and shift fuel purchases forward when faced with a tax increase.

the net present value (NPV) of future operating profits equals or exceeds the new-truck price, while trucks exit as soon as operating costs exceed revenues. We find that an increase in the upfront cost of new trucks causes an increase in the equilibrium freight rate and vehicle lifetime (consistent with Gruenspecht), while an increase in the operating cost of new vehicles causes an increase in the equilibrium freight rate, but has an ambiguous effect on vehicle lifetime.

We then analyze how incorporating anticipation (i.e., beliefs about future new-truck prices) affects investment and retirement patterns. In particular, we consider how an anticipated increase in the new-truck price affects pre-regulation purchasing decisions (holding constant the new-truck price in the pre-regulation period).⁶ In some contrast to models without anticipation, where higher upfront costs initially deter entry (allowing used durables to earn rents), forward-looking purchasers (chasing the opportunity to earn future rents) buy trucks just ahead of the regulation (fully dissipating those rents). That is, in the period immediately preceding regulation, firms add trucks to the fleet until the anticipated windfall to pre-regulation capital is fully arbitrated.⁷

Strategic pre-buying affects retirements and post-regulation new-truck sales. Injecting new trucks into the freight transport market lowers the current period freight rate, reducing the value of used trucks and causing the oldest trucks to immediately retire. The lower freight rate also extends the post-regulation period during which firms are discouraged from investing in new trucks, because long-run profits are too low to compensate for the higher purchase price. Though pre-buying distorts the transition to the new equilibrium, the predicted effect of anticipation on net sales is a true shift. That is, the increase in pre-regulation purchases induced by rent seeking is predicted to equal the decrease in post-regulation sales, causing no change in the long-run quantity of vehicles sold.⁸

The net environmental effect of anticipation depends on how gains from accelerated turnover compare with losses from more-modest emission-rate improvements. If older trucks are higher emitting, accelerating retirement provides an immediate environmental benefit. Post-regulation, deterred invest-

⁶In our model, as in the real world, trucks are elastically supplied by a competitive industry. Thus, the expected price increase does not create an incentive for OEMs to increase prices in the current period.

⁷The pre-buy is predicted to be strongest during the period just before regulation, because earlier purchases earn fewer revenues over the life of the truck. Strategic timing of purchases is explored in more depth in our theory section.

⁸This effect of anticipation is distinct from the direct effect of the policy on sales, which, as noted above, is predicted to be a decline in sales following the introduction of regulation.

ment reduces the environmental effectiveness of the policy.

To empirically test our predictions, we estimate a model of new-truck sales, using monthly Class-8 truck sales in the U.S. over the period 1991-2015. Our preferred specification includes the monthly average real oil price and quarterly U.S. Gross Domestic Product (GDP), as well as year and month-of-year fixed effects. This model explains the vast majority of variation in truck sales. Oil prices and GDP are significant drivers of truck sales, both of which were changing rapidly and dramatically around the time these standards were implemented. Our econometric approach attempts to disentangle the impact of anticipation from those of other sales drivers. We investigate whether anticipation affected sales by examining residual variation in sales (i.e., variation not explained by our model) around the month the standards took effect. If anticipation did impact the sales cycle, our model would predict persistently positive residuals just before the policy start date and persistently negative residuals immediately after the policy start date.

During this time, EPA implemented four rounds of new-truck criteria pollutant standards. We focus on the 2007 new-engine standards, which are widely regarded as the most significant regulatory action (i.e., with respect to trucks) taken by EPA during the 25-year span of our data. Consistent with our prediction, we find robust evidence of a sales spike in the months before implementation of the 2007 criteria pollutant standards, followed by a similarly-sized sales dip in the months after the regulation took effect. Across various specifications, we estimate anticipation of the 2007 criteria pollutant standards caused several thousand more trucks to be sold in each of the months prior to, and approximately the same number fewer trucks to be sold in each of the months after, the introduction of the standards, resulting in a net sales impact (of anticipation) which is statistically indistinguishable from zero.

Our study makes several contributions to the literature on standard-based (environmental) regulation. Our theoretical model introduces the concept of anticipation, distinguishing between the effects of this form of strategic behavior, the direct effect of policy on new-vehicle sales, and the indirect (i.e. Gruenspecht) effect on fleet composition. Our model illuminates two offsetting channels through which anticipation impacts the environmental effectiveness of standards, the net effect of which is *a priori* ambiguous. Empirically, we confirm that strategic responses to regulation affect investment cycles. Our results have important implications for policy design and program evaluation. Confounding the effects of anticipation with the direct effects of policy would, under a variety of identification strategies, result in significantly biased estimates. As noted in the growing literature on tax avoidance, it is critical that analysts account for this behavior when studying markets in which agents can

shift the timing of purchases in anticipation of new regulation.

The paper proceeds as follows. In Section II we provide a brief background on criteria pollutant regulations for HDVs in the U.S. Section III presents our theoretical model of the truck market, and explores the impacts of anticipated and unanticipated regulation-induced changes to upfront and lifetime costs on the cycle of capital turnover and the associated emissions. Section IV describes the econometric framework and Section V the data used to estimate the effect of anticipation on the pattern of new-truck sales. Section VI presents results and Section VII concludes.

II Background

EPA has regulated emissions from HDVs for over forty years (U.S. EPA, “Heavy-Duty Highway Compression-Ignition Engines and Urban Buses – Exhaust Emission Standards”). In 2001, EPA finalized criteria pollutant emission standards for diesel heavy-duty engines and vehicles, model-year 2007 and later, significantly tightening the standards for PM and NO_x emissions (40 C.F.R. § 69, 80, 86). For engines, the standard for PM decreased from 0.1 to 0.01 grams per brake-horsepower hour (g/bhp-hr), NO_x standards decreased from 2.4 to 0.2 g/bhp-hr, and standards for non-methane hydrocarbons (NMHC) were put in place (in isolation from NO_x standards, for the first time) at 0.14 g/bhp-hr. Standards were similarly tightened for overall vehicle performance, varying with gross vehicle weight class. EPA projected that the new standards would reduce annual NO_x, NMHC, and PM emissions by 2.6 million, 115,000 and 109,000 tons respectively, preventing an estimated 8,300 deaths, 9,500 hospitalizations, and 1.5 million lost workdays (40 C.F.R. § 69, 80, 86).

The 2007 criteria pollutant standards were expected to increase the total cost of purchasing and operating model-year 2007 and later HDVs. Though the price for larger HDVs can extend well beyond \$100,000 per vehicle (40 C.F.R. § 69, 80, 86), it was generally accepted that the cost of the 2007 criteria pollutant standards would represent a significant increase in the cost of new HDVs. EPA projected upfront costs for HDVs would increase by \$3,230 in early years, and lifetime operating costs would increase by \$4,600 (EPA, 2000). Industry analysts and OEMs predicted that surcharges of between \$7,000 and \$10,000 would *need* to be applied to model-year 2007 and later HDVs, to cover the cost of installing pollution-control technology. The discrepancy between EPA and industry projections of compliance costs likely contributed to uncertainty among HDV fleet operators. Leading up to the 2007 criteria pollutant

standards, the possibility of a pre-buy was widely discussed in the popular press, while some industry analysts advised purchasers to pre-buy HDVs.⁹

III Theory

A model of capital turnover

To analyze how anticipation of future price changes impacts the sales patterns and fleet composition in the HDV market, we develop a theory model, where profit-maximizing agents choose when to purchase and retire vehicles which they operate in a competitive freight market. We begin by describing the lifetime profit associated with a vehicle in this market:

$$\Pi = \int_0^T (P(Q(t)) - C(t))e^{-rt} dt - M. \quad (1)$$

Where:

T is the lifetime of the vehicle;

$P(Q(t))$ is the freight rate (i.e., operating revenue) at time t , which, fixing demand, is a function of the aggregate supply of vehicles in the market;

$C(t)$ is the operating and maintenance cost for a vehicle at time t , which increases monotonically as t approaches T ;

r is the discount rate; and

M is the purchase price of a new vehicle.

A competitive (i.e. price-taking) firm's profit maximization problem can be thought of in terms of its vehicle purchase and retirement decisions. A firm will choose to purchase and operate a new vehicle if the NPV of operating that vehicle is greater than the purchase price:

$$\int_0^T (P(Q(t)) - C(t))e^{-rt} dt \geq M. \quad (2)$$

⁹See, for examples, "Emission Rule Change the Engine for Truck Sales: But this year's boom to be next year's bust." Chicago Tribune. Web. 23 May 2016., and "To pre-buy or not." Fleet Owner. Web. 23 May 2016., respectively.

A firm will continue to purchase and operate new vehicles until the point where the profit gained from adding the last vehicle equals zero:

$$\int_0^T (P(Q(t)) - C(t))e^{-rt} dt = M. \quad (3)$$

A firm chooses when to retire a vehicle based on the first-order condition for profit maximization with respect to T :

$$\max_T \int_0^T (P(Q(t)) - C(t))e^{-rt} dt - M. \quad (4)$$

Differentiating with respect to T , we get the first-order condition for profit maximization:

$$P(Q(T)) = C(T). \quad (5)$$

Thus, a price-taking firm retires a vehicle at the time, T , when revenue equals costs.

When a truck is retired it is immediately replaced by a new vehicle. It follows directly from our specification of $C(t)$ that the freight-services supply curve is the aggregate cost function of all vehicles in service. In equilibrium (see Figure 1a), both supply of freight services and demand for those services are constant over time, yielding a constant quantity, Q , of vehicles in the market, and price, P , for the services those vehicles provide.

We are interested in the effect of new-vehicle emission standards on the HDV market. Emission standards may affect purchase and retirement decisions through two channels: they may cause OEMs to install additional abatement technology, increasing the purchase price of a new vehicle; and they may change the operating-cost function of a new vehicle. We model each channel, separately and in turn.

The equilibrium effects of regulation

Purchase price

Suppose that the implementation of a new-vehicle standard increases the purchase price, M , but does not change the cost function. How will this change affect P , Q and T ?

Note that a change in M will directly affect a firm's entry decision (Equation (3)) without directly affecting the exit condition for existing vehicles (Equation (5)). We thus begin by evaluating the comparative statics on the entry condition, and find that the equilibrium freight rate is increasing in M .

Equilibrium M can be described in terms of equilibrium price P , and lifetime T (itself a function of P):

$$M(T, P) = \int_0^T (P(Q(t)) - C(t))e^{-rt} dt. \quad (6)$$

Taking the total derivative of $M(T, P)$ above yields:

$$dM = \frac{\partial M}{\partial T} dT + \frac{\partial M}{\partial P} dP. \quad (7)$$

The partial derivative of M with respect to T , $\frac{\partial M}{\partial T}$, is simply the change in NPV of a vehicle evaluated at time T :

$$\frac{\partial M}{\partial T} = (P(T) - C(T))e^{-rT}. \quad (8)$$

From the retirement condition, one can see that, as a consequence of the envelope theorem, $\frac{\partial M}{\partial T}$ reduces to zero. In order to take the partial with respect to P , $\frac{\partial M}{\partial P}$, we first express $M(T, P)$ as:

$$M(T, P) = \int_0^T P(t)e^{-rt} dt - \int_0^T C(t)e^{-rt} dt. \quad (9)$$

In equilibrium, the price P is constant from time $t = 0$ to T . Thus, we view $P(t)$ as a constant P , and can integrate as follows:

$$M(T, P) = \frac{-Pe^{-rt}}{r} \Big|_{t=0}^{t=T} - \int_0^T C(t)e^{-rt} dt \quad (10)$$

$$= \frac{-Pe^{-rT}}{r} + \frac{P}{r} - \int_0^T C(t)e^{-rt} dt \quad (11)$$

$$= \frac{P}{r}(1 - e^{-rT}) - \int_0^T C(t)e^{-rt} dt. \quad (12)$$

Now it is simple to take the partial derivative:

$$\frac{\partial M}{\partial P} = \frac{1 - e^{-rT}}{r}. \quad (13)$$

Which we plug into Equation (7) and rearrange to find:

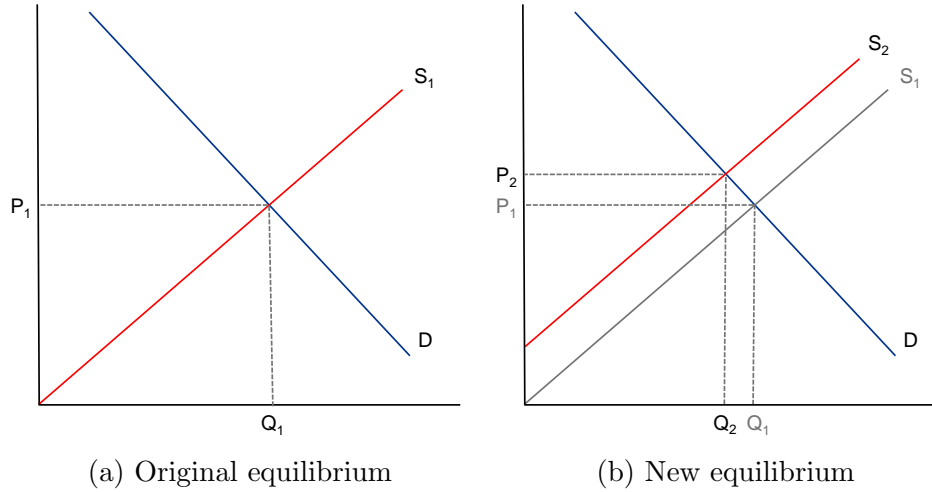
$$\frac{dP}{dM} = \frac{r}{1 - e^{-rT}}. \quad (14)$$

A positive vehicle lifetime T and discount rate r require that $-rT < 0$, and thus $e^{-rT} < 1$ and $\frac{\partial P}{\partial M} > 0$. That is, P is increasing in M . Fixing demand, it follows that a greater equilibrium P implies a lower equilibrium Q , while the exit condition implies a longer vehicle lifetime, T . The equilibria before and after a price change are shown in Figure (1).

Remark: Our model confirms two distinct and previously identified effects of an unanticipated price change. First, a price change affects the flow of new vehicles into the market: the quantity of vehicles purchased during a given timeframe is inversely related to the purchase price of the vehicle. We refer to this as the direct effect of the price change. Second, a change in the purchase price affects the stock of vehicles operating in the market: the quantity of vehicles retired during a given timeframe is inversely related to the purchase price of a new vehicle. Equivalently, lifetimes for all vehicles in the fleet are increasing in the new-vehicle purchase price. We refer to this as the Gruenspecht effect.¹⁰ Foreshadowing the discussion below, the existence of the Gruenspecht effect suggests that an increase in the flow of vehicles entering the market which reduces the NPV of used vehicles will drive an increase in retirements.

¹⁰Gruenspecht (1982) presents a structural model of vehicle fleet turnover under vintage-differentiated emission standards. As emission standards become more stringent, lifetime (upfront and possibly maintenance) costs of new vehicles subject to those standards rise. Gruenspecht shows that an increase in the price of new cars induces substitution towards used cars, which increases the value of used vehicles. In Gruenspecht's model, the scrappage rate depends on the value of used vehicles, where the scrappage condition is that the value is below a threshold. Thus, increasing the value of used vehicles drives down scrappage rates for existing vehicles, reducing the environmental effectiveness of the standard. Our model assumes a zero scrappage value, making Gruenspecht's scrappage rate analogous to our retirement rate.

Figure 1: The equilibrium effects of a price change



NOTE: In these figures the y-axis measures the freight rate, P , and the x-axis measures the quantity of vehicles in service, Q , or equivalently the quantity of freight services provided. In the first figure, the supply curve of freight services (red), which corresponds to the aggregate cost function of all trucks, slopes up from the origin to intersect the demand curve (blue) for freight services at the equilibrium P_1 and Q_1 . In the second figure, the supply curve (red) has shifted upward from the origin to intersect the demand curve (blue) at the new equilibrium P_2 and Q_2 . This supply curve corresponds to the uniformly older truck fleet that is in service just before the first vehicle purchased after regulation is put into service.

Cost function

Without loss of generality, suppose regulation decreases the cost function of new vehicles purchased after regulation, without affecting the purchase price. In particular, let $C_2(t) \leq C_1(t)$ for all $t > 0$, where $C_1(t)$ is the cost function of new vehicles purchased prior to the regulation taking effect and $C_2(t)$ the cost function of new vehicles purchased afterward. How will this change affect P , Q , and T ?

A change in $C(t)$ directly affects both the entry and exit conditions for vehicles in this market. Beginning with the entry condition, a lower $C(t)$ implies a new vehicle would earn positive lifetime profit at the existing equilibrium

freight rate P_1 and new-vehicle purchase price M :

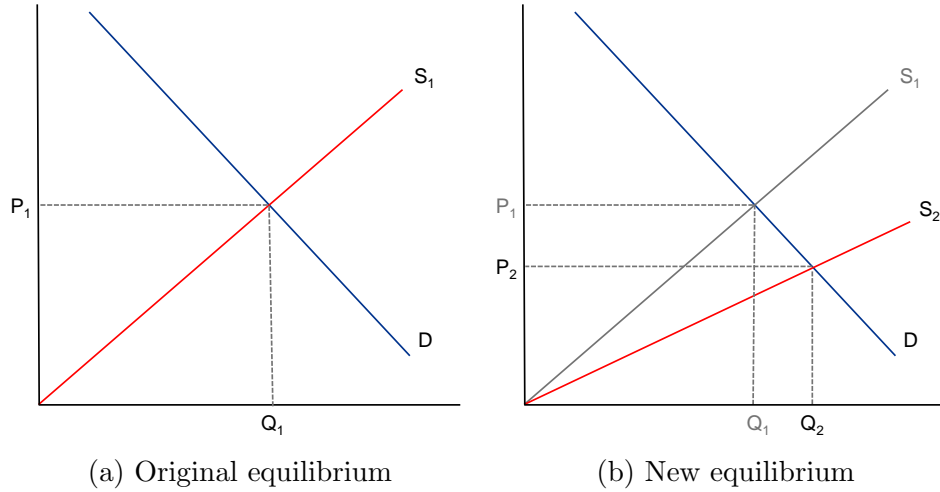
$$\int_0^{T_1} (P_1 - C_1(t))e^{-rt} dt = M < \int_0^{T_1} (P_1 - C_2(t))e^{-rt} dt. \quad (15)$$

The entry condition requires that firms will purchase new vehicles until the point where the NPV of a new vehicle subject to the regulation is exactly equal to the purchase price of a new vehicle. The exit condition requires that the operating cost of a vehicle at the end of its lifetime is equal to the market freight rate. We assume no change in purchase price; thus, the available adjustment margins for meeting these conditions are vehicle lifetime, and market freight rate. Note, there is no change in lifetime T alone that, holding P steady, will satisfy both the entry and exit conditions for an equilibrium. Holding P steady at P_1 , any increase in lifetime would only increase the NPV of a new vehicle. Alternatively, a decrease in T , holding P constant, would not satisfy the exit condition: $C_2(T_2) < C_1(T_2) = P_1$. Thus, it must be the case that such a decrease in the cost function results in a lower equilibrium freight rate and, fixing demand, a higher quantity of vehicles operating in the market.

The effect of the decrease in the cost function on the lifetime T of a vehicle depends on the new shape of the cost function. In the (special) linear case, where the change to $C(t)$ is simply a decrease in slope, lifetime is extended (see Figure (2) for a representation of the change in steady state for a linear $C(t)$). To see this, recall that M must be equal to the discounted value of operating profits over a vehicle's lifetime. If the change in $C(t)$ is simply a change in the slope of a line, M can be represented by the area of the triangle below the equilibrium price and above $C(t)$. In order for M to stay constant as the slope of the cost function declines, T must increase as P falls. In other cases, T may increase or decrease in response to a change in $C(t)$, depending on the relative levels and rates of change of the old and new cost functions.

Remark: We have shown the net change in the equilibrium freight rate, fleet quantity and truck lifetime due to regulation will depend on the changes to the purchase price and the cost function of new vehicles. For each change in $C(t)$ there is an implied change in M which will exactly offset the increase (or decrease) in lifetime costs. In this special case, equilibrium price and quantity are not affected by the regulation, and the effect on lifetime will depend solely on the change in $C(t)$, increasing for a decrease in costs, and decreasing for an increase in costs, as it takes more or less time for $C(t)$ to reach the equilibrium P .

Figure 2: The equilibrium effects of a cost change



NOTE: In these figures the y-axis measures the freight rate, P , and the x-axis measures the quantity of vehicles in service, Q , or equivalently the quantity of freight services provided. In the first figure, the supply curve of freight services (red), which corresponds to the aggregate cost function of all trucks, slopes up from the origin to intersect the demand curve (blue) for freight services at the equilibrium P_1 and Q_1 . In the second figure, the supply curve (red) has pivoted downward to intersect the demand curve (blue) at the new equilibrium P_2 and Q_2 . This supply curve corresponds to the truck fleet which is in service when the first regulation-compliant truck is retired.

Transitions and the effects of anticipation

We have established that regulation which changes the purchase price or cost function of new vehicles results in a new equilibrium freight price, quantity, and vehicle lifetime. The net impact of changes to M and $C(t)$ determines the new equilibrium. If both M and $C(t)$ increase, or if an increase in one outweighs a decrease in the other, we expect to see a higher equilibrium P . If both M and $C(t)$ decrease, or if a decrease in one outweighs an increase in the other, we expect to see a lower equilibrium P .

We now turn to the question of how the market transitions between equilibria. We first assume firms do not anticipate regulation. This is the standard form of analysis in the literature. We then examine how anticipation, through its impact on the transition, affects the pattern of new-vehicle sales and the

fleet composition. Without loss of generality, we analyze the case of an increase in M , from M_1 to $M_2 > M_1$, with no accompanying change in $C(t)$.

Unanticipated regulation

We first consider the case where the price increase is unanticipated, that is, firms respond to a change in price only after it has occurred. An increase in M does not immediately impact either $C(t)$ or P . The exit condition is met, and the first retirement after the price change occurs as it otherwise would have. However, that retired vehicle will not be replaced by a new vehicle purchase, as it would have been in equilibrium. The new entry condition will not be met until the NPV of a new vehicle increases to equal the new purchase price, M_2 . As vehicles retire without replacement, the supply curve shifts inward (or upward), as illustrated in in Figure (1b). As the fleet ages, the quantity of vehicles in service declines and the freight rate rises, increasing the NPV of vehicles in the market. When enough vehicles have retired such that the new entry condition is met, the market reaches its new equilibrium quantity and price, Q_2 and P_2 , respectively. Once this new equilibrium is reached, each retired vehicle is again replaced by a new purchase. The exit condition becomes $C(T_2) = P_2$, giving vehicles a longer lifetime T_2 for a higher P_2 . Empirically, we would expect the period immediately following the price change and during the transition to a new equilibrium to appear as a sales slump.

During this transition period, and once the new equilibrium is reached, the freight rate is greater than P_1 . Accordingly, vehicles purchased prior to the price change are earning greater revenue than expected at the time of their purchase. A vehicle purchased immediately before the implementation of new standards therefore accumulates a windfall for its entire lifetime, earning:

$$\int_0^\tau (P(Q(t)) - C(t))e^{-rt} dt + \int_\tau^{T_2} (P_2 - C(t))e^{-rt} dt. \quad (16)$$

Here, price, $P(Q(t))$, is increasing from P_1 to P_2 between $t = 0$ and $t = \tau$, afterwhich it is steady at P_2 for the remainder of the vehicle's lifetime, until retirement at T_2 . Rents associated with the difference in expected and earned profits due to the unanticipated price change are characterized by:

$$\begin{aligned} Rents = & \int_0^\tau (P(Q(t)) - P_1)e^{-rt} dt + \int_\tau^{T_1} (P_2 - P_1)e^{-rt} dt \\ & + \int_{T_1}^{T_2} (P_2 - C(t))e^{-rt} dt. \quad (17) \end{aligned}$$

Looking at the right-hand side of this equation, the first integral represents the difference in the earned and expected revenues during the period when the market is in transition to a new equilibrium. The second integral represents the difference in earned and expected revenues from the time the market reaches a new equilibrium until the time at which the vehicle would have been retired, T_1 . The third and final integral represents the revenue earned during the period of extended lifetime, from the time T_1 which the vehicle was expected to retire when purchased, until the time T_2 when the vehicle is retired under the new equilibrium price. Each of these integrals is positive, and increasing in P_2 . In a world where firms anticipate, and can respond to future changes in M , the prospect of earning these rents provides an incentive to purchase additional vehicles before the price changes.

Anticipated regulation

In reality, new regulations are implemented through an often-lengthy public process, which provides regulated entities with knowledge about the timing and content of forthcoming standards. When firms anticipate a change in the price of new vehicles, and foresee the rents associated with a higher future equilibrium freight rate, we expect them to arbitrage those rents by adding vehicles to the market. This pre-buy immediately lowers the freight rate, driving down revenues accruing to vehicles in the market. Firms will continue to add vehicles until the point where the negative rents earned when $P(Q(t)) < P_1$ equal, in absolute value, the positive rents earned when $P(Q(t)) > P_1$:

$$\int_0^{\tau_0} (P(Q(t)) - P_1)e^{-rt} dt = \int_{\tau_0}^{\tau} (P(Q(t)) - P_1)e^{-rt} dt + \int_{\tau}^{T_1} (P_2 - P_1)e^{-rt} dt + \int_{T_1}^{T_2} (P_2 - C(t))e^{-rt} dt. \quad (18)$$

The left-hand side of this equation represents the period from the vehicle purchase until the time $t = \tau_0$, when the market price $P(Q(t))$ reaches the price at the old equilibrium, P_1 . During this time, vehicles earn lower revenues than they would have under original equilibrium, giving the integral a negative value. This offsets the higher rents earned as $P(Q(t))$ rises above P_1 for the remainder of the vehicle's lifetime. The first right-hand side integral represents the time between $t = \tau_0$ and $t = \tau$ when the market price is above the original equilibrium price, P_1 , but below the new equilibrium price, P_2 . The second and third right-hand side integrals correspond to the second and third integrals in

Equation (17), above. These three right-hand integrals are positive, and their sum is equal, in absolute value, to the left-hand side integral.

This equality implies a larger pre-buy for a larger price change: The higher is P_2 relative to P_1 , the larger are the rents on the right-hand side of this equation, the more negative must be the rents on the left-hand side and, thus, the larger in volume must be the pre-buy.¹¹

Empirically, we expect to see an increase in vehicle sales directly prior to regulation, followed by a sales slump directly after regulation is implemented. Note, this sales slump is distinct from and additional to the direct effect of the policy on new-vehicle sales.

Anticipation also affects the composition of the fleet through the exit condition. Recall that a firm will retire a vehicle at the time, T , when its cost, $C(T)$, equals the freight rate, $P(Q(T))$. When firms anticipate a price change, the injection of pre-bought vehicles immediately drives down the freight rate below P_1 . Those vehicles whose costs are greater than the new freight rate are no longer profitable to keep in operation, and will be pushed out of the market, into early retirement.¹² Importantly, the quantity of vehicles pushed out must be less than the quantity of pre-bought vehicles. Only with an increase in total quantity will price decline in the short term, dissipating all available rents. We will return to this observation when considering the impact of the pre-buy on emissions.

Remark: Note that the new steady-state quantity and price are the same whether or not the price change is anticipated. This implies that the effect of anticipation on vehicle flows must be a symmetric shift in vehicle sales.

Remark: Our model predicts that anticipation of a price change affects both the flow and stock of vehicles. These effects are additional to the previously identified direct and Gruenspecht effects, respectively. Anticipation of

¹¹We treat the time between when the price change is announced and when it takes effect as a single decision point. In practice, firms often know many months in advance that a price change will occur. One question this raises is whether one would expect to see a pre-buy well in advance of the price change, as firms seek to capture the future stream of rents. While we do not explicitly model this tradeoff, we do consider the following thought experiment. Suppose a firm chooses to purchase a vehicle just ahead of when all other “pre-buy” vehicles are purchased. This vehicle would earn the same stream of revenues as the other “pre-buy” vehicles, except that it would initially earn some $P' < P_1$ (i.e., because it entered, pushing down the price) and not earn P_2 for as long an interval at the end of life. If the average flow of revenues to the other “pre-buy” vehicles is P_1 – which must be the case in order to satisfy the entry conditions with equality – then subtracting a portion $P_2 > P_1$ and replacing it with $P' < P_1$ must yield lower revenues, and therefore be dominated.

¹²Under changing prices, this exit condition becomes more complex, as there may be multiple points where $C(t) = P(Q(t))$, multiple local maxima. A vehicle will exit the market at the time T which produces the global maximum profit.

costly regulation increases the flow of vehicles into the market directly prior to regulation implementation, and symmetrically decreases the flow of compliant vehicles post-implementation, reducing the environmental effectiveness of the policy. Through the same channel as the Gruenspecht effect, anticipation which increases the flow of vehicles ahead of regulation implementation is expected to reduce the remaining lifetime revenues and thus the value of used vehicles. This pushes the oldest, highest-emitting vehicles in the fleet to an early retirement, increasing the environmental effectiveness of the policy. Note, the net effect of anticipation on the emissions impact of the policy is *a priori* ambiguous.

Emissions impact of anticipation

When firms change their purchasing and retirement decisions in anticipation of impending regulation, they impact the environmental outcome of that regulation. Without loss of generality, we continue to examine the case of an increase in purchase price, that is, where anticipation results in a pre-buy. Anticipation affects emissions through two channels – pre-bought vehicles create additional emissions, compared to the policy-compliant vehicles they displace, but emissions are avoided from those vehicles pushed into early retirement by the pre-buy. The net environmental impact depends on the relative size of these countervailing effects. The effect on emissions, evaluated over the lifetime of the pre-bought vehicles, is captured in the following equation:

$$\Delta E = \int_0^{T_2} (e_{pre}\bar{Q}_{pre} - e_{post}Q_{post}(t) - e_{ret}(\bar{Q}_{ret} - Q_{ret}(t)))dt. \quad (19)$$

In discussing the emissions impacts, we will use “Base Case” to refer to the scenario without anticipation, and “Anticipation Case” to refer to the scenario in which firms anticipate future regulation. We define e_{pre} as the average emission rate for a pre-bought vehicle, e_{ret} as the average emission rate for a vehicle which is retired before the market reaches the equilibrium in the Anticipation Case, and e_{post} as the average emission rate for a vehicle which is purchased after the policy is implemented. Finally, let \bar{Q}_{pre} be the magnitude of the pre-buy, \bar{Q}_{ret} be the magnitude of vehicles pushed into early retirement by the pre-buy, $Q_{post}(t)$ be the difference in the quantity of vehicles purchased after the policy is implemented in the Base and the Anticipation cases, and $Q_{ret}(t)$ be the difference in the quantity of vehicles retired after the policy is implemented in the Base and Anticipation cases.¹³

¹³Note that \bar{Q}_{pre} refers to the quantity of vehicles which were purchased prior to the implementation, as a result of expectations about the regulation. This does not include

It is instructive to evaluate the environmental outcome in three distinct periods, during which anticipation affects emissions, and between which that effect differs. Let t_0 be the time of the pre-buy (before policy is implemented), t_1 be time when the market reaches equilibrium (Q_2, P_2) in the Base Case, t_2 be the time when the market reaches equilibrium (Q_2, P_2) in the Anticipation Case (τ in Equations (16) through (18) above), and t_3 be the time when the pre-bought vehicles exit the market in the Anticipation Case. To discuss the dynamics of emissions effects, we will consider, as an illustrative example, the special case where the slopes of the supply and demand curves are constant, as in Figure (1). Constant slopes imply that vehicles are retired at the same rate in the Anticipation Case, between t_0 and t_2 , as they are in the Base Case between t_0 and t_1 .

Between t_0 and t_1 , no vehicles are purchased in either the Base Case or the Anticipation Case (i.e. $Q_{post}(t) = 0$). The retirement rate is the same in both cases (i.e. $Q_{ret}(t) = 0$). All emissions from pre-bought vehicles in the Anticipation Case are additional, while emissions from the quantity of vehicles pushed into early retirement in the Anticipation Case are avoided. During this period, the emissions impact of the anticipation is:

$$\Delta E_{0-1} = (t_1 - t_0)(e_{pre}\bar{Q}_{pre} - e_{ret}\bar{Q}_{ret}). \quad (20)$$

Between t_1 and t_2 , the market is in equilibrium in the Base Case, but has not yet reached equilibrium in the Anticipation Case. As vehicles are purchased in the Base Case, $Q_{post}(t)$ increases, eroding the additional emissions from the pre-bought vehicles. At the same time, vehicles are retiring more quickly in the Base Case than in the Anticipation Case, increasing $Q_{ret}(t)$ and reducing the difference between total retirement in the Base and Anticipation Case, and thus reducing total avoided emissions. During this period, the emissions impact of anticipation is:

$$\Delta E_{1-2} = \int_{t_1}^{t_2} (e_{pre}\bar{Q}_{pre} - e_{post}Q_{post}(t) - e_{ret}(\bar{Q}_{ret} - Q_{ret}(t)))dt. \quad (21)$$

By time t_2 , the market has reached equilibrium in both the Base and the Anticipation cases. The supply curve has shifted to the same point, crossing the demand curve at the new equilibrium. The operating cost $C(t_2)$ of the

those vehicles which would have been purchased in the absence of any anticipation of policy changes. Similarly, \bar{Q}_{ret} refers only to the quantity of vehicles which are retired early as an immediate result of the pre-buy. We have used \bar{Q} to refer to a quantity of vehicles which remains constant over time, while $Q(t)$ refers to a quantity of vehicles which changes over time.

next retiring vehicle is equal to P_2 , implying that all vehicles with $C(t_2) > P_2$ have already been retired in both cases. Thus, the total quantity of retired vehicles in the Base Case is now equal to that in the Anticipation Case (i.e. $Q_{ret}(t) = 0$), and the total quantity of purchased vehicles in each case must also be equal. Since no vehicles have been purchased post-policy in the Anticipation Case before t_2 , we have $Q_{post}(t_2) = \bar{Q}_{pre}$. From t_2 to t_3 , vehicles are purchased and retired at the same rate in both cases. The only difference in emissions stems from the higher emission rate of the pre-bought vehicles compared to that of the policy-compliant vehicles purchased in the Base Case. During this period, the emissions impact of anticipation is:

$$\Delta E_{2-3} = (t_3 - t_2)(e_{pre} - e_{post})\bar{Q}_{pre}. \quad (22)$$

The net emissions impact depends on the differences in the emission rates of the three categories of vehicles, and the shape of the supply and demand curves.

IV Econometric Framework

In this section we present our for measuring the effect of anticipation on the new-vehicle sales cycle. Our setting is the 2007 implementation of criteria pollutant standards for HDVs, which are widely regarded as the most significant regulatory action taken by EPA with respect to trucks during the span of our data.¹⁴ Given the expected increases in purchase price and operating costs, our model predicts a sales spike directly prior to implementation, followed by a sales slump in the months after the regulation took effect.

In developing an econometric framework to test the predictions of our theory model, we face several empirical challenges. First, the timing of anticipation's effect on the sales cycle is not immediately obvious. While our model predicts firms will pre-buy in the period immediately before the standards are implemented, concerns around vehicle availability may cause some to shift their purchases further ahead of regulation. In addition, GDP and oil prices – two key drivers of HDV demand – were moving around the time the policy was implemented. Failing to account for these drivers would likely bias our esti-

¹⁴EPA and industry analysts predicted new-vehicle costs would rise by \$3,230 and \$10,000, respectively, due to the installation of new technologies. Uncertainty around cost increases and reliability of new technologies led some industry analysts to recommend pre-buying in response to the new standards.

mates.¹⁵ Finally, we do not find a well-suited comparison population to serve as our counterfactual, so we must rely on time-series variation to identify the effect of anticipation.¹⁶

To address these empirical challenges, we estimate a model of monthly HDV sales and investigate the pattern of residuals (i.e. deviations from the predicted sales pattern) around the time the regulation was implemented for evidence of anticipation.¹⁷ This approach allows us to control for confounding covariates (not driven by regulation), does not require an external counterfactual, and does not require any strong assumptions about the time period in which anticipation might have affected HDV sales.

The intellectual basis for our approach to accounting for anticipation comes, in part, from McCrary (2008). In the context of a regression discontinuity design, McCrary develops a test for manipulation of the running variable, whereby agents may affect their own treatment status.¹⁸ Hausman and Rapson (2017) elaborate on the implications for the McCrary test when time determines treatment status. In these cases, endogenous movements in the outcome variable in response to the timing of treatment present identification challenges. They argue that this makes direct application of the McCrary test impossible, as one cannot separate the presence of manipulation from the effect of treatment. In our context, firms may easily manipulate the time of purchase in anticipation of costly HDV emissions standards, plausibly causing bunching of vehicle sales directly before the price change and invalidating time-series es-

¹⁵This is the flawed approach used in several analyses of the 2007 criteria pollutant standards in the gray literature (See, for examples, Harrison and LeBel (2008) and Calpin and Plaza-Jennings (2012)).

¹⁶We consider and reject three comparison groups: new-HDV sales in Canada, new-HDV sales in Mexico, and new-light duty vehicle (LDV) sales in the U.S. Canada adopted the same HDV standards simultaneously with the U.S., and Mexico adopted HDV standards for the following model-year, within the plausible treatment window of the U.S. policy. Further, new-vehicle sales in Canada and Mexico may have been impacted by U.S. HDV policy, violating the stable unit treatment value assumption. Finally, while LDV sales likely share some drivers with HDV sales (including oil price and the state of the macroeconomy), it is unlikely that those drivers affect HDV and LDV sales in the same way, making the comparison unsuitable for panel-data identification strategies.

¹⁷Under the assumption that HDVs are supplied elastically, these deviations can be interpreted as demand shocks.

¹⁸The regression discontinuity design is a method for identifying the effects of a treatment on individual outcomes when an arbitrary threshold along a continuous “running variable” determines treatment status. McCrary argues that endogenous manipulation of the running variable violates the assumption of continuity (i.e., quasi-random assignment around the treatment threshold) and thus invalidates a regression discontinuity identification strategy in these cases.

timates of the direct effect of the policy on sales. We inspect the econometric residuals for bunching around July 2007 (i.e., when the new-engine standard took effect, causing a discontinuous change in HDV prices).

Similar to Lam and Bausell (2007), we model monthly U.S. HDV sales, Q_t , as a linear function of several economic drivers.¹⁹

$$Q_t = \beta_0 + \beta_1 GDP_t + \beta_2 Real_t + \beta_3 X_t + \varepsilon_t \quad (23)$$

Here, GDP_t and $Real_t$ are GDP in the quarter corresponding to month t and the real price of oil in month t , respectively. X_t is a set of year and month-of-year binary variables equal to one in their designated year or month-of-year, respectively, and zero, otherwise. ε_t is the econometric error term, which accounts for period-specific sales shocks that are orthogonal to other specified sales drivers.

GDP is an important indicator of the strength of the economy, which is used in many industries, including trucking, to forecast the strength of future demand for their products and services. We expect truck operators to purchase more trucks in response to strong economic growth, which implies a positive β_1 in Equation (23).

We use real oil price as a proxy for operating costs.²⁰ Fuel is the largest single cost for trucking fleets, accounting for 38% of the cost of ownership. Lease or purchase payments, by contrast, account for 10% of total average operating costs to owners (Torrey and Murray, 2014). Moreover, other operating costs (e.g., maintenance) are likely correlated with implementation of the regulation. Including them could bias our estimates (Lam and Bausell, 2007). According to our theory model, higher operating costs will reduce lifetime

¹⁹Lam and Bausell (2007) attempt to measure the impact of EPA’s 2002 HDV-engine standards on the timing of truck production, providing the only prior empirical evidence of pre-buying in response to vehicle emissions standards. They specify an econometric model of truck production as a function of GDP, diesel fuel prices, average retail truck prices, quarter-of-year fixed effects and a linear time trend. Lam and Bausell do not consider the effect of anticipation on sales after the implementation of the regulation. To estimate the size of the pre-buy, they use a binary variable that takes a value of one in the six months prior to the implementation of the 2002 regulation, and zero otherwise, and find that production increased by approximately 20% in the six months prior to implementation.

²⁰While diesel prices may be a more direct measure of fuel costs, we find that oil prices are a stronger predictor of vehicle sales, particularly in the presence of month-of-year fixed effects. This appears to be due to a spurious correlation between the annual cyclical of vehicle sales and diesel fuel prices (likely driven by seasonal variation in home-heating-oil demand). Additionally, the global oil price is less likely to be affected by local shocks, which may affect both diesel price and HDV sales and which would potentially bias our estimates of the effect of diesel fuel costs on sales.

profit, reducing demand for new freight trucks, so we expect β_2 in Equation (23) to be negative.²¹

Month-of-year binary variables are included to flexibly model regular variation in truck sales. These monthly differences are likely due to the cyclicity of new model-year releases. Yearly binary variables are included to flexibly model average differences in sales across years, which may be due to differences in product attributes across vehicle model-years.

Vehicle purchase price is excluded from the estimating equation. The pollution-control technology required to comply with the 2007 standards increased the cost of producing a new HDV. If these costs were passed on to HDV purchasers, or if price were endogenous to demand, including price in the regression would absorb some of the sales shock caused by the regulation. Additionally, we don't find a reliable instrument for vehicle supply.²² If the policy drove manufacturers to change prices (i.e. increase prices ahead of regulation in response to the pre-buy), then excluding price from the estimating equation would attenuate our estimates of the effects of anticipation on new-vehicle sales. To the extent that vehicle price varies systematically by year or month-of-year, i.e. with the release of new model-year vehicles, our vector of binary variables likely absorbs the resulting sales variation.

The econometric error – variation in the data not explained by the econometric model – can give insight into underlying patterns of sales shocks. If the model fits the data, then residuals will vary quasi-randomly around zero. We seek to identify the effect of anticipation by visually inspecting the residuals for significant shocks around the time the regulation took effect. In addition to a graphical analysis, we test for the presence of anticipation using regression analysis. First we regress the residuals on two binary variables:

$$\varepsilon_t = Pre_t + Post_t + \epsilon_t. \quad (24)$$

Here, ε_t is the econometric error term from Equation (23). Pre_t takes the value of one in the several months before the regulation is implemented, and zero otherwise; $Post_t$ takes the value of one in the several months after the regulation was implemented, and zero otherwise. ϵ_t is the remaining econometric

²¹Our investigation of the relationship between oil prices and the monthly new-truck producer price index (PPI) suggests new-truck prices are relatively insensitive to oil-price fluctuations. To the extent that oil price enters the new-truck supply function, we expect higher oil prices to drive higher new-truck prices, resulting in a reduced quantity demanded.

²²Lam and Bausell (2007) instrument for the price of freight trucks with the price of steel, a significant input to truck production. However, closer examination of this relationship shows that the truck PPI is relatively insensitive to changes in the steel PPI. Further, the steel PPI is driven by changes in oil price, violating the exclusion restriction.

error. We then add Pre_t and $Post_t$ to our main specification:

$$Q_t = \beta_0 + \beta_1 GDP_t + \beta_2 Real_t + \beta_3 X_t + Pre_t + Post_t + \varepsilon_t. \quad (25)$$

In addition to our main specifications, we perform several robustness checks to test the sensitivity of our results to the specification of treatment-period length, and to the specification of the relationship between oil prices and sales.

V Data

Monthly Class-8 HDV sales for the U.S., from 1991 through 2015, were purchased from Ward’s Automotive, Inc. Monthly real oil prices (in 1991 dollars) for the same time period were collected from the Energy Information Administration (EIA). Quarterly GDP data for the same period were obtained from Bureau of Economic Analysis (BEA), and matched to monthly observations in the same quarter and year. Table (1) presents descriptive statistics for our data.

Table 1: Descriptive Statistics of Variables used in our Analysis

Variable	Obs	Mean	Std. Dev.	Min	Max
HDV Sales	285	14367	4529	6232	26380
Oil Price (1991 \$)	285	52.69	30.96	13.45	137.39
GDP (\$ Billion)	285	11480	3429	6054	17600

VI Results and Discussion

In this section we report results for our econometric model of sales, investigate residual variation not explained by our econometric model for evidence of anticipation using graphical and regression analysis, and test the sensitivity of our results to alternative specifications.

Sales model

Table (2) reports coefficients and standard errors for several specifications of our econometric model of truck sales. Column (1) reports results for our base specification, excluding year and month-of-year fixed effects. While the coefficients on real oil price and GDP each take the predicted sign, and are highly

statistically significant, they account for a relatively modest share of sales variation. To explain more of the variation, and to control for seasonality and year-specific disturbances that could confound our results, we add annual and month-of-year fixed effects. Column (2) reports results for this specification. Fixed effects significantly improve the fit of our econometric model. Coefficients on real oil price and GDP have the same signs as in Column (1) and are statistically significant at the one percent level. Adding fixed effects attenuates the impact of the real oil price on sales and amplifies the effect of GDP. We infer that the fixed effects are flexibly controlling for some variation in sales that covaries with, but is not driven by, the real oil price. At the mean, a one percent increase in GDP is associated with a four percent increase in HDV sales, while a one-hundred percent increase in the real oil price is associated with a twenty percent decline in sales.

Table 2: Demand Regression Results

VARIABLES	(1) Sales	(2) Sales	(3) Residuals	(4) Sales
Real Oil Price	-83.94*** (15.83)	-54.38*** (13.72)		-28.98** (11.22)
GDP	0.813*** (0.139)	4.650*** (1.559)		5.387*** (1.260)
Pre-treatment			4,010*** (525.0)	4,452*** (928.8)
Post-treatment			-4,681*** (525.0)	-5,226*** (853.8)
Observations	290	290	290	290
R-squared	0.107	0.864	0.330	0.914
Year FE	N	Y	N	Y
Month FE	N	Y	N	Y


This table reports coefficients and standard errors for four specifications of our econometric model of truck sales. In Columns (1), (2) and (4), the dependant variable is monthly Class-8 HDV sales. Columns (1) and (2) present results from a specification which uses Real Oil Price and GDP as independent variables; the specification reported in Column (2) includes year and month-of-year fixed effects. In Column (3), the dependant variable is the difference in actual Class-8 HDV sales and Class-8 HDV sales predicted by the fixed-effects model presented in Column (2), and the independent variables are two binary variables which respectively take the value of one during the seven months prior to and the seven months after regulation takes effect. Column (4) presents results from a specification which includes these two binary variables in the fixed-effects specification presented in Column (2). Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Graphical analysis

Figure (3) plots monthly residuals from our fixed-effects specification, the coefficients and standard errors for which are reported in Column (2) of Table (2). The y-axis reports the quantity of sales not explained by our econometric model and the x-axis reports the date (i.e., month of sample). Each (blue) dot is a monthly observation and the (red) reference line is aligned with the month the regulation is implemented (July 2007).

It is clear from the residuals plot that there was some anticipation of the policy. During the months immediately prior to the implementation of the

Figure 3: Plot of monthly sales residuals



NOTE: In this figure the y-axis reports the difference between actual monthly Class-8 HDV sales and monthly Class-8 HDV sales predicted by our fixed-effects regression, the coefficients and standard errors for which are presented in Column (2) of Table (2). The x-axis reports the date. Each (blue) dot is a monthly observation and the (red) reference line is aligned with the month the regulation took effect (July 2007).

regulation the residuals are strongly and consistently positive. Immediately following the implementation, residuals are strongly and consistently negative. As our model predicts, these effects appear to be approximately symmetric. Away from the policy start date, the distribution of residuals appears quasi-random.

Econometric analysis

We test for the significance of anticipation's impact on sales by incorporating two binary variables, Pre_t and $Post_t$, into our regression framework. Pre_t takes the value of one during the seven months prior to the regulation and zero otherwise, while $Post_t$ takes the value of one in the seven months after the

regulation takes effect and zero otherwise. Both Pre_t and $Post_t$ take the value of zero for the month in which the regulation was implemented (July 2007). The pattern of variation in Figure (3) suggests that sales were affected between four and eight months on either side of the policy. We choose seven months as the treatment period for our main specification, but test the sensitivity of our results to alternative treatment-period specifications in robustness checks, below.

Column (3) reports the results for a regression of the residuals from our fixed-effects model, shown in Figure (3) above, on the two binary variables, Pre_t and $Post_t$. We confirm that the coefficients are statistically different from zero, with a strongly positive coefficient on Pre_t and strongly negative coefficient on $Post_t$. As predicted, we are unable to statistically distinguish between the absolute values of these coefficients. However, the more-negative point estimate on $Post_t$ may indicate the variable is absorbing some of the direct effect of the policy on sales.

Column (4) reports results for a model that includes Pre_t and $Post_t$ in the main regression. These results are qualitatively similar – a strongly positive coefficient on Pre_t and strongly negative coefficient on $Post_t$, with statistically indistinguishable absolute values – though this specification attenuates the coefficient on oil, and yields marginally stronger, though not statistically different, coefficients on both of the policy variables. That the coefficient on oil price changes suggests that oil prices covary with the treatment period. In Column (3), more of the sales variation around the time of the policy is attributed to changing oil prices, while in Column (4) Pre_t and $Post_t$ soak up more of the sales variation. We explore the sensitivity of our results both to alternative specifications of the treatment period, and to alternative assumptions regarding the relationship between oil prices and sales, below.

Sensitivity to treatment-period specification

Table (3) presents results for our fixed-effects model, including Pre_t and $Post_t$, across a range of treatment-period specifications. Column (4) corresponds to our central specification (also Column (4) in Table (2)).

Our results are robust to alternative specifications of the pre and post period. The coefficients on Pre_t and $Post_t$ are statistically significant and generally consistent across these specifications. The five- and six-month specifications produce the most markedly different coefficients on the policy variables. In addition, these specifications break the pattern of an otherwise-stable oil-price coefficient; oil prices and policy variables drive sales to different extents in those two specifications. One explanation for the pattern of results

Table 3: Robustness to Treatment-period Specification

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Pre: 4m Post: 4m	Pre: 5m Post: 5m	Pre: 6m Post: 6m	Pre: 7m Post: 7m	Pre: 8m Post: 8m
Real Oil Price	-33.08*** (12.72)	-18.24 (11.69)	-19.76* (11.61)	-27.57** (12.06)	-27.31** (12.27)
GDP	4.242*** (1.389)	3.925*** (1.259)	4.519*** (1.255)	5.286*** (1.320)	5.609*** (1.354)
Pre-treatment	4,209*** (1,111)	2,248* (1,188)	7,376*** (1,229)	4,480*** (938.9)	4,082*** (820.3)
Post-treatment	-6,763*** (1,115)	-9,507*** (1,204)	-4,005*** (1,071)	-5,215*** (861.0)	-4,925*** (779.2)
Observations	285	285	285	285	285
R-squared	0.901	0.918	0.920	0.913	0.910
Year FE	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y

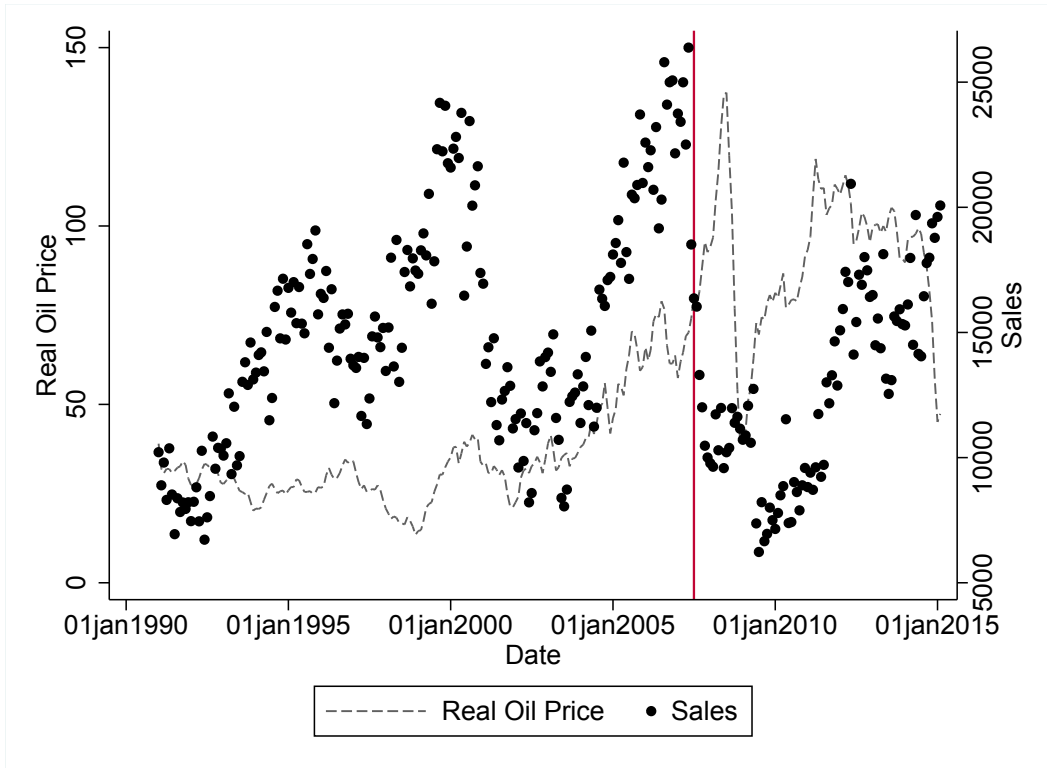
NOTE: This table reports coefficients and standard errors for five specifications of our econometric model of truck sales. In each column, the dependant variable is monthly Class-8 HDV sales. Each specification includes Real Oil Price, GDP, year and month-of-year fixed effects, and two binary variables (Pre-treatment and Post-treatment). The column titles report the number of months prior to regulation for which Pre-treatment takes the value of one, and the number of months after regulation takes effect for which Post-treatment takes the value of one. Pre- and Post-treatment each take the value of zero in all other months, including the month the regulation takes effect (July 2007). Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

reported in Columns (2) and (3) is that the five- and six-month specifications covary particularly closely with periods of extreme oil-price variation, entangling, in those specifications, the effect of the regulation with the effect of oil-price variation.

Sensitivity to oil-price specification

A challenge to our identification strategy is to correctly identify the effect of the real oil price on sales. The oil price changes dramatically and unprecedentedly around the time the 2007 criteria pollutant standards were implemented (see Figure (4)). We have observed that the coefficient on oil prices is sensitive

Figure 4: Monthly HDV sales and oil prices



NOTE: Monthly sales are for Class-8 HDVs, and are purchased from Ward’s Automotive, Inc. Oil prices are real monthly average imported crude oil prices, in 1982-1984 dollars, as reported by the Energy Information Administration (EIA). The red reference line corresponds to the month the regulation took effect (July 2007)

both to the addition of Pre_t and $Post_t$, and to the specification of those variables, consistent with the fact that treatment covaries with the oil prices. As a result of this covariance, it appears that in certain specifications the policy variables soak up sales variation that may, in reality, be driven by the oil price. We explore the consequences of this covariance by testing the sensitivity of our results to alternative specifications of the impact of oil price on sales, identifying the effect of the policy by regressing the sales residuals on the policy variables.

Table 4: Sensitivity to Oil-price Specification

VARIABLES	(1) Sales (all)	(2) Residuals (all)	(3) Sales (Excl. 2007)	(4) Residuals (Excl. 2007)	(5) Sales (shift)	(6) Residuals (shift)	(7) Sales (shift Excl. 2007)	(8) Residuals (shift Excl. 2007)
Real Oil Price	-54.38*** (13.72)		-11.23 (10.70)					
GDP	4.650*** (1.559)		3.919*** (1.165)		4.397*** (1.548)		3.750*** (1.158)	
Pre-treatment		4,450*** (516.9)		4,450*** (516.9)		3,929*** (520.1)		4,450*** (516.9)
Post-treatment		-5,026*** (516.9)		-5,026*** (516.9)		-4,620*** (520.1)		-5,026*** (516.9)
Oil (pre)					64.22 (50.68)		69.70* (37.85)	
Oil (post)					-60.97*** (13.86)		-16.22 (10.85)	
Observations	290	290	278	290	290	290	278	290
R-squared	0.864	0.376	0.921	0.376	0.867	0.327	0.922	0.376
Year FE	Y	Y	Y	N	Y	N	Y	N
Month FE	Y	Y	Y	N	Y	N	Y	N

NOTE: This table reports coefficients and standard errors for four two-stage specifications of truck sales. Columns (1), (3), (5) and (7) each report results for a first-stage fixed-effects regression of monthly Class-8 HDV sales (excluding Pre-treatment and Post-treatment variables). Columns (2), (4), (6) and (8) report results of second-stage regressions of the residuals from the previous column's specification on Pre-treatment and Post-treatment. For each of these second-stage regressions, the independent variables are two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Columns (1) and (2) reproduce Columns (2) and (3) from Table (2). Columns (3) and (4) report results for a specification that excludes observations from 2007 in the first-stage regression. Columns (5) and (6) report results for a specification which includes two separate independent variables for monthly real oil price before 2002, and after 2001, respectively. Columns (7) and (8) report results for a specification which both excludes observations from 2007, and includes the two separate variables for monthly real oil price in the first-stage regression. Triple, double, and single asterisks denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table (4) reports coefficients and standard errors for four two-stage specifications of truck sales. Columns (1), (3), (5) and (7) each report results for a first-stage fixed-effects regression of monthly Class-8 HDV sales (excluding Pre_t and $Post_t$). Columns (2), (4), (6) and (8) report results of second-stage regressions of the residuals from the previous column’s specification on the policy variables. For each of these second-stage regressions, the independent variables are two binary variables, Pre-treatment and Post-treatment, which respectively take the value of one during the seven months prior to regulation, and the seven months after regulation took effect. Columns (1) and (2) reproduce Columns (2) and (3) from Table (2).

We’ve seen that oil prices co-vary with the treatment period. It’s possible that our estimated coefficient on oil prices in previous specifications is soaking up variation which is in fact driven by the policy variables. To test this hypothesis, we suppose that the relationship between oil prices and vehicle sales is equal in 2007 to the mean value of that relationship in all other years and calculate the implied residuals for 2007 by differencing actual and predicted monthly sales. Columns (3) and (4) report results for this specification. Somewhat surprisingly, the coefficient on oil in this specification becomes insignificant. The coefficients on Pre_t and $Post_t$ break significantly from the relatively tight range we have observed over other specifications. These results suggest that the specified effect of oil price on sales may, here, be overly restrictive. Real oil prices rose 66% in 2007, to \$95/barrel, the highest observed price in our data up to that point. If firms responded to the significant rise in oil price differently than to more moderate price changes during periods of relative price stability, the relationship between oil price and truck sales during this volatile period might be importantly different.

If the effect of oil price on sales did change over time, a natural moment for that change to begin might have been during the post-2001 run up in oil prices. Between 1991 and 2001, the mean real oil price was \$28/barrel, with a standard deviation of \$6/barrel. Between 2002 and 2015, the mean real oil price was \$74/barrel, with a standard deviation of \$27/barrel. It is conceivable that this period of increasing, and increasingly volatile, oil prices marked a change in the relationship between the real oil price and HDV sales. To test this hypothesis, we interact oil price with indicator variables for pre-2002 and post-2001; Oil_{pre} takes the value of the monthly real oil price between 1991 and 2001 and zero otherwise, while Oil_{post} takes the value of the monthly real oil price between 2002 and 2015, and zero otherwise. Columns (5) and (6) report first- and second-stage regression results for this specification, respectively. Our results suggest that there was a structural shift in the relationship after 2001. The coefficient on Oil_{pre} is positive and insignificant, and the coefficient on

Oil_{post} is negative and significant. The coefficients on the policy variables are statistically indistinguishable from those in Column (2). While these results suggest the effect of oil prices on sales did change significantly post-2001, as before, it may be that the coefficient on Oil_{post} is affected by sales variation that is, in fact, caused by the policy.

To address this potential source of bias, we combine the two previous specifications, estimating coefficients for Oil_{pre} and Oil_{post} , excluding data from 2007. Here, the exclusion of 2007 is likely to be somewhat less restrictive than in the previous specification. If there was a structural break, the mean response to oil price changes in 2007 is likely closer to the mean response around that time than it is to the mean response over the entire time period. Columns (7) and (8) report first- and second-stage regression results for this specification, respectively. In this specification, oil price is a positive, significant (at the 10% level) driver of sales before 2002, and a negative, insignificant driver after 2001 (though this point estimate is not statistically different from the point estimate in Column (4) of Table (2)). The estimated value of Pre_t and $Post_t$ fall within the range of estimates observed in Table (2).

While we are unable to definitively disentangle the effects of anticipation from the (changing) effect of oil prices, we find that the estimated impact of anticipation is relatively stable over a range of plausible assumptions.

VII Conclusion

Vintage-differentiated emission standards are widely used to regulate pollution from mobile and stationary sources. When a new emission standard is expected to discontinuously change the purchase price or lifetime cost of a new piece of equipment, forward-looking agents may shift the timing of purchases in order to avoid compliance costs. To investigate the incentives created by vintage-differentiated standards, the impact of those incentives on the new-equipment sales cycle, and the implications for the effectiveness of new-equipment standards, we analyze the market for new Class-8 HDVs. In the context of new-vehicle emissions standards, prior analyses have not considered anticipation as an adjustment margin. In this paper, we address four specific questions: How does the anticipation of regulation affect the pattern of new-truck sales? How does the pattern of new-truck sales affect the pattern of used-truck retirements? How do purchasing and retirement patterns affect the environmental benefits of standards? Empirically, have recent regulations caused firms to pre-buy trucks? To answer these questions, we first develop a theoretical model, which incorporates the effects of anticipation on

new-vehicle sales and the used-vehicle fleet, and differentiates those impacts from the previously identified direct and Gruenspecht effects. We test our predictions using a data set of monthly U.S. sales of new freight trucks around the time of EPA's 2007 implementation of HDV criteria pollutant standards, widely regarded as the most significant action taken by EPA (i.e., with respect to trucks) during the 25-year span of our data. Consistent with our predictions, we find evidence that anticipation caused a sales spike in the months before the policy took effect and a sales slump after implementation. For analysts using time-series variation to study the effects of standards, failing to account for anticipation likely results in significantly biased estimates of the direct effect of the policy on sales. More broadly, our findings have important implications for the analysis of markets in which agents can shift the timing of purchases in anticipation of new regulation.

We begin by specifying a dynamic model of a competitive freight truck market, where firms incorporate new-truck prices, operating costs and freight rates (i.e., operating revenue) into their purchasing and retirement decisions, and calculate comparative statics for changes in upfront and operating costs. We find that an increase in the upfront cost of new trucks causes an increase in the equilibrium freight rate and vehicle lifetime (consistent with Gruenspecht), while an increase in the operating cost of new vehicles causes an increase in the equilibrium freight rate, but has an ambiguous effect on vehicle lifetime. We then analyze how incorporating anticipation (i.e., beliefs about future new-truck prices) affects investment and retirement patterns. We find that, if firms are given the opportunity to buy trucks ahead of costly regulation, they will shift purchases forward, increasing demand for new trucks before regulation is implemented, symmetrically decreasing demand after implementation, and pushing out the oldest (highest-emitting) vehicles in the fleet. The net environmental effect of anticipation depends on how gains from accelerated turnover compare with losses from more-modest emission-rate improvements.

We test our predictions by estimating an econometric model of new-truck sales, using monthly HDV sales in the U.S. over the period 1991-2015. We investigate whether anticipation affected sales by examining residual variation in sales around the month the standards took effect. Consistent with our theory model, we find evidence of anticipation of the 2007 criteria pollutant standards. We estimate anticipation of the standards caused several thousand more trucks to be sold in each of the months prior to, and approximately the same number fewer trucks to be sold in each of the months after, the introduction of the standards. Our results are relatively stable across various specifications.

Our results have important implications for policy design and program

evaluation. Confounding the effects of anticipation with the direct effects of policy would, under a variety of identification strategies, result in significantly biased estimates. Ex ante, policy-makers should account for the effects of anticipation, and minimize the costs associated with it. For example, they may choose to phase in new standards (or award credits for early compliance), eliminating the discontinuous price change which induces a pre-buy. Ex post, analysis that does not account for anticipation risks mischaracterizing the effects of policy. Anticipation is not unique to emissions standards in the HDV industry; similar behavior has recently received increasing attention in the tax avoidance literature. Going forward, it will be important to consider and identify the effects of anticipation across a wide range of policy areas. Whenever regulation is expected to result in a discontinuous change, and agents affected by the regulation are able to adjust the timing of their behavior, we should expect to see some form of anticipation.

Although EPA has updated HDV emission standards at multiple dates covered in our data, we did not observe similarly strong evidence of sales shocks around the implementation of these standards (in 1998, 2002 and 2010). There are several reasons firms may not have pre-bought in anticipation of these regulations. Not only were the 2007 standards the most costly for OEMs to comply with, but the technology required to comply with these standards was, in many cases, unfamiliar to trucking firms. The uncertainty around the new technology may have strengthened the incentive to purchase pre-policy vehicles, which may have been perceived to be more reliable. In contrast, other emission standards generally required that OEMs use technology which was already available, reducing the uncertainty around the reliability and operating costs of policy-compliant vehicles.

In addition to criteria pollutants, EPA regulates GHG emissions. EPA and the National Highway Transportation Safety Administration have recently finalized Phase 2 GHG and fuel efficiency standards for model-years 2019-2027, which standards are expected to significantly increase the upfront costs of new vehicles. Given the pre-buy observed in response to criteria pollutant standards, one might predict a similar response to these Phase 2 standards. However, our theoretical model suggests these standards are not likely to create conditions that would induce a pre-buy. While the societal benefits of criteria pollutant standards diffuse throughout the population, the bulk of the benefits of GHG standards accrue to vehicle operators. EPA projects, within several years, the proposed Phase 2 GHG standards will fully compensate HDV operators for the increased upfront cost, and that the lifetime savings accruing to those operators will total \$170 billion (U.S. EPA, “Cutting Carbon Pollution, Improving Fuel Efficiency, Saving Money and Supporting Innovation

for Trucks”). In the context of our theory model, decreased lifetime operating costs, which outweigh the increase in upfront costs, would eliminate the incentive to pre-buy vehicles.

Future work may consider the implications of alternative market structures. We model competitive suppliers in a competitive freight market. Results of a model with less competition in either sector would surely yield qualitatively different results. In particular, less competition upstream would add endogeneity to the upfront price, and less competition downstream would alter the entry and exit conditions in the freight market. In addition, it would be worthwhile to verify our model empirically in a setting where macroeconomic conditions are more stable around the time of the policy change, and in a setting where data and a clear counterfactual would allow for an identification of the direct effect of the policy. Finally, while we did not empirically estimate the emissions impact of anticipation, making those estimations, and clarifying the conditions under which anticipation increases and decreases a policy’s net emissions impact, would be a valuable contribution to the literature.

References

- 40 C.F.R. § 9, 22, 85, et al. Greenhouse Gas Emissions and Fuel Efficiency Standards for Medium and Heavy-Duty Engines and Vehicles—Phase 2; Proposed Rule (July 2015).
- 40 C.F.R. § 69, 80, 86. Control of Air Pollution from New Motor Vehicles: Heavy-Duty Engine and Vehicle Standards and Highway Diesel Fuel Sulfur Control Requirements.
- Calpin, Patrick and Esteban Plaza-Jennings. “A look back at EPA’s cost and other impact projections for MY 2004-2010 heavy-duty truck emissions standards.” *NADA Comments to NHTSA/EPA Re: MY 2017-2025 Proposed Standards* (2012).
- Cashin, D. and T. Unayama. “Measuring intertemporal substitution in consumption: Evidence from a VAT increase in Japan.” *Review of Economics and Statistics* (2016).
- Coglianesi, John, Lucas W. Davis, Lutz Kilian, and James H. Stock. “Anticipation, tax avoidance, and the price elasticity of gasoline demand.” *Journal of Applied Econometrics* (2016).
- Crossley, T., H. Low, and C. Sleeman. “Using a temporary indirect tax cut as a fiscal stimulus: evidence from the UK.” (2014). Unpublished Manuscript, University of Essex.
- D’Acunto, Francesco, Daniel Hoang, and Michael Weber. “The effect of unconventional fiscal policy on consumption expenditure.” NBER Working Paper 22563 (2016).
- “Emission Rule Change the Engine for Truck Sales: But this year’s boom to be next year’s bust.” Chicago Tribune. Web. 23 May 2016.
- Fullerton, Don, and Sarah E. West. “Can taxes on cars and on gasoline mimic an unavailable tax on emissions?.” *Journal of Environmental Economics and Management* 43, no. 1 (2002).
- Gruenspecht, Howard K. “Differentiated regulation: The case of auto emissions standards.” *The American Economic Review* (1982).
- Harrison, David and Mark LeBel. “Customer behavior in response to the 2007 heavy-duty engine emission standards: Implications for the 2010 NO_x standard.” *NERA Economic Consulting* (2008).
- Hausman, Catherine and David. S. Rapson. “Regression discontinuity in time: Considerations for Environmental Economists” Working Paper (2017).
- Jacobsen, Mark R., and Arthur A. Van Benthem. “Vehicle scrappage and gasoline policy.” *The American Economic Review* (2015).
- Lam, Terence, and Charles Bausell. “Strategic behaviors toward environmental regulation: The case of the trucking industry.” *Contemporary Economic*

- Policy* (2007).
- Maloney, Michael T., and Gordon L. Brady. "Capital turnover and marketable pollution rights." *Journal of Law and Economics* (1988).
- McCrary, Justin. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of Econometrics* (2008).
- Stavins, Robert N. "Vintage-differentiated environmental regulation." *Stanford Environmental Law Journal* (2006).
- "To pre-buy or not." Fleet Owner. Web. 23 May 2016.
- Torrey, W. Ford and Dan Murray. "An Analysis of the Operational Costs of Trucking: A 2014 Update." American Transportation Research Institute (2014).
- U.S. Department of Transportation Research and Innovative Technology Administration Bureau of Transportation Statistics. *State Summaries: 2007 Commodity Flow Survey* (2010).
- U.S. Environmental Protection Agency. "Heavy trucks, buses and engines." Available at <http://www3.epa.gov/otaq/hd-hwy.htm>.
- U.S. Environmental Protection Agency. *Regulatory Impact Analysis: Heavy-Duty Engine and Vehicle Standards and Highway Diesel Fuel Sulfur Control Requirements* (2000).
- U.S. Environmental Protection Agency. "Cutting Carbon Pollution, Improving Fuel Efficiency, Saving Money, and Supporting Innovation for Trucks." Regulatory Announcement EPA-420-F-15-900 (June 2015). Available at <http://www3.epa.gov/otaq/climate/documents/420f15900.pdf>.
- U.S. Environmental Protection Agency. "Heavy-Duty Highway Compression-Ignition Engines and Urban Buses – Exhaust Emission Standards." Available at <http://www3.epa.gov/otaq/standards/heavy-duty/hdci-exhaust.htm>.



MIT Center for Energy and Environmental Policy Research

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

**MIT Center for Energy and
Environmental Policy Research**
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139
USA

Website: ceepr.mit.edu

MIT CEEPR Working Paper Series is published by
the MIT Center for Energy and Environmental
Policy Research from submissions by affiliated
researchers.

Copyright © 2017
Massachusetts Institute of Technology

For inquiries and/or for permission to reproduce
material in this working paper, please contact:

Email ceepr@mit.edu
Phone (617) 253-3551
Fax (617) 253-9845