

Technology Adoption and the Timing of Environmental Policy: Evidence from Efficient Lighting

Sarah Armitage*

April 25, 2022

JOB MARKET PAPER

[Please click here for the latest version.](#)

Abstract

How does supporting early clean technologies affect the long-run transition away from dirty technologies? Early policy action generates immediate environmental benefits from increased adoption of available efficient products, but may result in intertemporal substitution away from later products with greater potential for reducing externalities. I examine how standards and subsidies supporting early advancements in lighting efficiency (halogens, CFLs) impacted the adoption of later products with higher efficiency (LEDs). I estimate a model of residential lighting demand, using structural methods

*I thank my advisors Jim Stock, Ed Glaeser, Joe Aldy, and Robin Lee for their continual guidance and support. I am also grateful to Ariel Pakes, Rich Sweeney, Kostas Metaxoglou, Jing Li, Frank Pinter, Alex Albright, Tianwang Liu, and Hillary Stein for helpful discussions, as well as to participants at the Harvard Environmental Economics Workshop and Seminar, Labor/ Public Finance Workshop, and Industrial Organization Workshop; the EAERE Summer Conference; the AERE Summer Conference; the EMEE Summer Workshop; and the OSWEET Workshop. I gratefully acknowledge financial support from the NBER Pre-Doctoral Fellowship in Energy Economics and the Joseph Crump Fellowship at Harvard University. I also thank ten experts in industry and government who generously shared their time and information on the lighting market. Researcher own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

adapted from dynamic models to capture how the market size and distribution of consumer heterogeneity depended endogenously on the history of past purchases. Counterfactual simulations suggest that delaying the implementation of standards from 2012 to 2018 would result in 36% greater LED sales over this period, while delaying the phase-out of CFL subsidies from 2012 to 2018 would result in 20% fewer LEDs sold. Across a range of specifications, I find that environmental benefits from some early policy action outweigh the environmental cost of reduced LEDs adoption; the overall environmental externality is minimized when standards are implemented in 2012 and CFL subsidies are phased-out after 2014. Sensitivity analyses around alternative technology lifetimes, externalities, and innovation responses identify conditions under which early policy intervention would be counterproductive.

JEL Codes: H22, L67, Q54, Q55, Q58

1 Introduction

Climate scientists and economists agree that avoiding catastrophic consequences of climate change will require a dramatic transition in energy technologies. Supporting early clean or efficient technologies has been a mainstay of policy interventions towards this goal. While early incentives for clean technologies immediately shift consumption away from dirty products, in the longer term these policies may either “crowd in” or “crowd out” the development of later-generation clean technologies (Sivaram, 2017). For example, early support for hybrid vehicles may have increased consumer awareness of alternative fuel vehicles more broadly or promoted battery research that proved useful for later electric vehicles. However, these policies may have also caused consumers to adopt long-lived vehicles with smaller environmental benefits than electric vehicles or shifted innovation resources to solving hybrid design challenges with little applicability to more advanced clean vehicles.¹ In this paper, I

¹Other anecdotal examples of the potential for long-run “crowd out” of superior environmental technologies include competition between silicon-based solar PV and solar cells made from other materials, between

examine how incentivizing the deployment of earlier-generation clean technologies may affect the long-run transition away from dirty technologies, using the market for efficient lighting as an empirical setting. Given the scale of emissions reductions needed for climate change mitigation, hindering the development of potentially superior clean energy technologies may substantially increase the long-term social cost of climate change mitigation efforts.

Economic theory can illuminate why the timing of environmental policy may have significant implications for the long-run efficiency of climate change mitigation. For one, many clean technologies are durable products, so subsidies or standards that encourage immediate adoption of long-lived products may reduce the market size for later-generation, more efficient technologies or change the distribution of consumers who consider adopting these later technologies.² Even if the market size for later-generation products eventually increases due to stock turnover, a smaller initial market size may slow the process of learning-induced cost reductions (Benkard, 2004; Levitt et al., 2013), and longer time-to-profitability for these later products may diminish firms' willingness to invest in new technologies (Stein, 1989; Budish et al., 2015).³ In efficient policy design, any dynamic inefficiencies in the longer-term transition to clean or efficient products must be weighed against the benefits of achieving immediate reductions in emissions.

To analyze this issue, I study the market for efficient lighting, where the incumbent inefficient technology (incandescents) was replaced by several types of more efficient products. During this transition, the later-generation efficient technology (light emitting diodes, or LEDs) competed with both the earlier-generation efficient technology (compact fluorescent lighting, or CFLs) and a more efficient version of the incumbent technology (halogens). Two key policies influenced the recent development of the efficient lighting market in the United

lithium ion batteries and other battery chemistries, and between ethanol-based biofuels and biofuels derived from other organic matter; in these examples, public support for the early technology may have hindered the development of later technologies with potentially greater long-term technical promise (Sivaram, 2017).

²Solar panels, wind turbines, energy storage, electric vehicles, electric heating systems, industrial process improvements, and efficient appliances are all examples of durable clean technologies.

³Alternatively, knowledge spillovers from the deployment of early efficient products may accelerate the development of later products, given sufficient similarities in their respective technologies, manufacturing processes, or marketing requirements.

States. First, numerous state and utility programs have provided product rebates to counteract the higher upfront cost of efficient bulbs, first for CFLs and later for LEDs. Second, federal lighting efficiency standards removed most traditional incandescent bulbs from the market beginning in 2012, precipitating the large-scale adoption of halogen bulbs and, to a lesser extent, increased adoption of CFL bulbs. Through my analysis of consumer purchase decisions, I simulate the development of the efficient lighting market given different policy scenarios, asking how the timing of these policy instruments influenced their outcomes. The lighting market provides an ideal setting for investigating the technology transition process, given that this process has advanced further than for many other energy technologies. Furthermore, this market is also independently important given the technical potential of LEDs to reduce global electricity consumption — by 7-10% according to one industry estimate (Kooroshy et al., 2016).⁴

To conduct this analysis, I build a structural model of demand in the residential lighting market over 2010 to 2018 to determine the impact of subsidies and standards on the penetration of different technologies. To estimate this model, I rely primarily on retail sales data and estimate a discrete choice demand model in the spirit of Berry et al. (1995). Because light bulb lifetimes vary significantly by technology type, ranging from 1,000 hours for an incandescent bulb to 8,000 hours for a CFL bulb to 15,000 to 25,000 hours or more for an LED bulb, I consider how the distribution and number of potential customers depends on the past history of technology adoption. This component of the analysis is essential to account for the heterogeneous adoption of efficient lighting at the household level, documented in EPA (2010), as the distribution of consumer preferences represented in the market may evolve over the study period. It is also important to account for the nearly 40% decrease in new bulb shipments over the study period, largely a result of longer light bulb lifetimes. To model this endogenous evolution of market size and consumer heterogeneity, I adapt methods from dynamic demand models for my static demand setting; to my knowledge, this paper is one

⁴Efficient lighting has also been a primary driver of reductions in residential electricity consumption over the last decade (Davis, 2017).

of the first to account for endogenous coming-to-market as part of a static demand system, and the approach outlined in this paper may be applicable to other markets undergoing technology transitions.

Counterfactual simulations take a multipronged approach. First, I ask how the impact on the environmental externality would change if federal efficiency standards were phased in or subsidies for CFLs were phased out at an alternate date during the study period. Second, to capture how these welfare impacts might change in the presence of alternate innovation paths, I ask how much faster or slower the marginal cost of LEDs would have needed to fall during the study period to achieve the same total LED market size under each of these counterfactual policy scenarios. Finally, I examine how these results depend on the characteristics of the technology in question, by asking how results would differ if we altered bulb lifetime or externality production by technology.

These counterfactual simulations predict that delaying the beginning of efficiency standards from 2012 to 2018 would result in 29.0% greater LED sales over the study period. This change in total LED sales is equivalent to LED prices falling by half a percentage point faster each quarter during the study period. Likewise, delaying the phase-out of CFL subsidies from 2012 to 2018 would result in 35.6% *fewer* LEDs sold, which is equivalent to a similar in magnitude *decrease* in the rate of LED price changes. However, counterfactual simulations suggest that these reductions in the size of the initial LED market do not outweigh the immediate environmental benefits from at least some early policy intervention. The average externality per hour of lighting sold during the study period is minimized when standards are implemented beginning in 2012 and when CFL subsidies are terminated after 2014. While 2012 was the year of actual standards implementation, 2014 was earlier than many actual lighting rebate programs fully discontinued their support for CFLs.

In examining how these counterfactual results depend on key technology parameters, I find that the relationship between the timing of standards and the average externality produced is remarkably robust to a range of alternative LED price declines. However, the

impact of standards timing is sensitive to the parameters such as the relative lifetime of halogen bulbs, the technology that most directly benefited from standards implementation. The impact of rebates timing is more sensitive to the rate of LED price declines, and ending rebates anywhere from 2011 to 2016 may minimize the average externality, depending on the rate at which LED prices fall. Under alternative assumptions about CFL lifetimes — that frequent on-off switching did not cause realized lifetimes to be unexpectedly lower than rated lifetimes — phasing out subsidies as early as 2012 may have produced the lowest average externality.

This paper connects to several areas of literature. The core idea of competition across multiple generations of clean or efficient technologies is similar in spirit to several papers that have considered dynamic effects of other environmental policies. Most notably, Langer and Lemoine (2018) derive the optimal path of subsidies for a clean technology whose price is declining over time and where consumers exhibit heterogeneous willingness-to-pay for the clean product. Their model, evaluated empirically in the context of the California Solar Initiative, focuses on a single technology for which the regulator seeks to achieve some total quantity of deployment by a terminal date; the model trades off the higher willingness-to-pay of early adopters, which would favor increasing subsidies over time, with higher initial prices, which would favor decreasing subsidies over time. By contrast, my paper focuses on discrete changes in technologies available for decarbonizing a given sector, assuming that the regulator wants to reduce emissions by some amount, and trades off initial efficiency gains with longer-run improvements in technology performance. Other relevant papers that examine dynamic issues in clean technology development include Gerarden (2018) on endogenous investments in solar technology improvement; Covert and Sweeney (2020) on learning rates within and across firms in wind turbine manufacturing; and Bollinger and Gillingham (2019) on learning-by-doing in solar installations.⁵

⁵In many cases, innovation market failures may interact with inefficiencies from second-best environmental policies, given the dual market failure associated with climate change (Jaffe et al., 2005). Also of note, Johnstone et al. (2010) consider the impact of policy instrument choice on energy innovation, using a multi-country panel dataset of patent applications. Their results suggest that technology-specific policies such

This paper also builds on the literature on climate and directed technical change, which examines long-run transition paths from dirty to clean production and posits that the total size of the clean or dirty sector can affect both the direction and magnitude of technological progress (Acemoglu, 2002; Aghion et al., 2009; Acemoglu et al., 2012; Casey, 2019). However, much of this literature focuses on competition between a single “dirty” sector and a single “clean” sector, ignoring whether the dominance of certain clean technologies over others may have implications for the long-run potential growth of the clean sector. One exception is Acemoglu et al. (2019), which studies the impact of using natural gas as a “bridge fuel” on the growth of renewable energy with greater emissions reduction potential.

Third, this paper also contributes to a long evaluation literature on energy efficiency programs, both academic literature and studies commissioned by program administrators (Houde and Aldy, 2017; Allcott and Greenstone, 2012, 2017; Hoffman et al., 2017). In one of the few academic economics papers to examine the efficient lighting market over this period, Allcott and Taubinsky (2015) use online and in-store experiments to quantify consumers’ failure to internalize savings from installing CFLs instead of incandescent bulbs, concluding that information and attention biases are insufficient to justify bans on incandescent bulbs. My paper complements this work by examining the impact of lighting subsidies and standards over a longer time horizon and considering the positive rather than normative implications of consumer preferences over different lighting technologies.

Finally, this paper employs methods from the extensive literature on discrete choice demand estimation, including several papers that estimate demand systems using retail scanner data specifically (Berry et al., 1995; Nevo, 2001, 2003; Hendel and Nevo, 2006b,a, 2013; Villas-Boas, 2007; Dubois and Bonnet, 2010; Asker, 2016). My paper builds on these static demand models by endogenizing the quantity and distribution of consumers in the market in each period, using methods derived from dynamic demand estimation (Lee, 2013;

as feed-in tariffs are useful for encouraging innovation in technologies that are costlier and further from commercial competitiveness, while general policies such as renewable energy credits (RECs) are more likely to induce innovation in more advanced technologies.

Gowrisankaran and Rysman, 2012).

Section 2 develops a simple theory model of multiple clean technology generations to clarify ideas. Section 3 describes the development of the efficient lighting market over 2010 to 2018 and provides reduced form statistics to illustrate the impact of standards and subsidies on lighting technology choice. Section 4 describes the structural model used for demand estimation, including the modeling of endogenous coming-to-market. Section 5 describes the data used in the analysis and the estimation procedure. Section 6 offers findings from counterfactual simulations of alternative policy designs. Finally, Section 7 concludes and offers directions for future research.

2 Theory

To clarify ideas, I construct a simple model of multiple generations of clean technology under varying environmental policy instruments. The model takes as given that policymakers may not be able to implement the optimal carbon tax due to political economy constraints and considers the design of a second-best policy that minimizes distortions in this dynamic setting.⁶ I first discuss how private optimization by consumers and the resulting technology market shares depend on product availability in each period. I then show how the total externality produced over multiple periods can be decomposed into several determinants of consumer demand. Finally, when policymakers are constrained to implement a technology-specific policy on only the first-generation clean technology (e.g., a subsidy), I show that the efficient price on this good is weakly higher when the policymaker must also consider the arrival of the second-generation technology. I discuss how this efficient price on the first-generation technology varies with key product and market characteristics, including the relative externality produced by different technologies and the arrival rate of the later clean

⁶The model is similar in spirit to literature in public finance that begins with the presupposition that the optimal lump sum taxation is not feasible for policymakers and considers the efficient design of distortionary taxation to achieve redistribution objectives or raise government revenue. See Brown (2016) and Hendren and Sprung-Keyser (2020).

technology relative to the lifetime of the early technology.

2.1 Consumer Choice Over Multiple Technologies

I assume that there are three products available to consumers: the dirty technology D , the first-generation clean technology C_1 , and the second-generation clean technology C_2 .⁷ Let per-period indirect utility for individual i consuming product j be given by:

$$U_{ij} = \nu_{ij} - p_j$$

where p_j represents the price for product j , common to all consumers, and ν_{ij} represents i 's idiosyncratic preference for product j , for $j \in \{D, C_1, C_2\}$. Consumers may also choose not to consume any of the three products, with the resulting utility normalized to 0. Each consumer may choose at most one technology in each period.⁸

The second-generation clean technology becomes available to consumers at period $t = A$; the first-generation clean product has a lifetime of length L . I assume that the arrival time of the second-generation technology is weakly less than the lifetime of the first-generation technology ($A \leq L$), so consumers who adopt the long-lived first-generation clean product have left the market when the second-generation becomes available. The dirty technology is assumed to last for only one period, so consumers of this product are still in the market when the second-generation clean technology arrives.⁹ To simplify notation in illustrating the

⁷In this simple model, I assume that each technology corresponds to a single product. In my empirical estimation of lighting demand, I allow for multiple products per technology.

⁸In this model, a “consumer” represents a household choosing a particular technology for a specific use case, such as a household choosing a particular lighting technology for a specific lamp. In this sense, the model is ruling out utility interactions across multiple products (e.g., light bulbs) owned by the same household, which matches the approach taken in my empirical estimation.

⁹A note about timing assumptions is useful here. I assume in this simple model that the first-generation clean product C_1 has a longer lifetime than the dirty product D , which is an appropriate choice for modeling the efficient lighting market, as described further below. While more subtle, this model set-up may also be appropriate for decisions such as installing solar panels, where the “dirty technology” entails continued use of grid electricity generated primarily from fossil fuels, which may be stopped at any time, while the “clean technology” has a lifetime of 10 to 20 years, which entails substantial technological lock-in for a consumer making this adoption decision. Furthermore, similar dynamic effects might be observed when dirty and clean technologies have similar lifetimes but a policy is designed to encourage early replacement of the dirty

consumer's optimization decision, I will set $A = 2$ and $L = 2$ in the derivation that follows, although I discuss how results generalize to arbitrary A and L (provided that $A \leq L$). Given this simplifying assumption, the consumer's potential consumption bundles $(j_{t=1}, j'_{t=2})$ are: (D, D) , (C_1, C_1) , (D, C_2) , $(0, C_2)$, and $(0, 0)$.¹⁰

Technology prices and preferences are assumed to be fixed over all periods. I assume that the efficient technology is more expensive than the dirty one, and the price of the efficient technology is decreasing in subsequent generations, such that $p_d < p_{c_2} < p_{c_1}$. Although I model constraints on consumers' future choice sets based on their previous consumption of the durable first-generation clean product, I ignore the role of lumpy investments in this simple model; p_j represents a per-period rental price rather than an upfront investment cost. Each of the three technologies also creates some negative externality, ζ_j , where the magnitude of this externality is greatest for the dirty product and decreasing with subsequent generations of the clean product ($\zeta_d > \zeta_{c_1} > \zeta_{c_2}$). The externality associated with the outside option is 0. Lastly, I also assume that individuals (and later the regulator) have perfect foresight over the arrival of the second-generation efficient technology.¹¹ While stylized, this assumption allows me to focus on the market characteristics of greatest interest. I discuss in the Appendix how this basic framework could be used to determine the efficient direction for policy updating, given an existing set of second-best policies, as new information about the second-generation technology becomes available.

Consumer i will prefer one bundle to another if the discounted present value of (private) utility from the former is greater than that of the latter. Given consumer discount factor λ ,

product, such as the cash-for-clunkers program.

¹⁰It is trivial to show that if the consumer prefers either D or C_1 over no consumption in the first period, then the same preference relation will also hold in the second period, and vice versa. Therefore, we can ignore the bundles $(0, D)$, $(D, 0)$, $(0, C_1)$, $(C_1, 0)$, (D, C_1) , and (C_1, D) .

¹¹An interesting avenue for future research would be examining how the presence of uncertainty influences the most salient tradeoffs from the regulator's perspective. I discuss this issue further in my counterfactual policy simulations for the lighting market.

we can derive pairwise preference relations across all combinations of consumption bundles:

$$\{j_{t=1}, j'_{t=2}\} \prec \{k_{t=1}, k'_{t=2}\} \leftrightarrow U_{ij_{t=1}} + \lambda U_{ij'_{t=2}} \leq U_{ik_{t=1}} + \lambda U_{ik'_{t=2}} \quad (1)$$

Substituting indirect utility expressions for each consumption bundle allows us to derive threshold values for ν_{ij} above which consumers are willing to adopt one technology over another. I obtain the following pairwise preference relations:

$$\begin{aligned} (D, D) \prec (0, 0) &\leftrightarrow \nu_{id} \leq p_d \\ (C_1, C_1) \prec (D, D) &\leftrightarrow \nu_{ic_1} - \nu_{id} \leq p_{c_1} - p_d \\ (C_1, C_1) \prec (0, 0) &\leftrightarrow \nu_{ic_1} \leq p_{c_1} \\ (D, C_2) \prec (D, D) &\leftrightarrow \nu_{ic_2} - \nu_{id} \leq p_{c_2} - p_d \\ (0, C_2) \prec (0, 0) &\leftrightarrow \nu_{ic_2} \leq p_{c_2} \\ (D, C_2) \prec (C_1, C_1) &\leftrightarrow \nu_{id} - p_d + \lambda(\nu_{ic_2} - p_{c_2}) \leq (1 + \lambda)(\nu_{ic_1} - p_{c_1}) \\ (0, C_2) \prec (C_1, C_1) &\leftrightarrow \lambda(\nu_{ic_2} - p_{c_2}) \leq (1 + \lambda)(\nu_{ic_1} - p_{c_1}) \end{aligned}$$

These preference relations and the resulting technology consumption shares are depicted graphically in Figures 1 and 2. Figure 1 illustrates consumer choice of technology conditional on $\nu_{id} \geq p_d$; that is, consumers in this region prefer the dirty product to the outside option, so the relevant consumption bundles are (D, D) , (C_1, C_1) and (D, C_2) . The set of consumers who prefer (C_1, C_1) to (D, D) is illustrated in Figure 1a, and the set of consumers who prefer (D, C_2) to (D, D) is illustrated in Figure 1b. As is apparent in these figures, a subset of consumers prefer both (C_1, C_1) and (D, C_2) to (D, D) ; some fraction will adopt the first-generation clean technology right away, whereas others will wait for the availability of the second-generation clean technology. The relative size of these two groups is depicted by the triangles in the upper right corner of Figure 1c. The slope of the line separating these two regions is determined by the timing assumptions of the model. In general, the

slope is equal to $\frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L}$; in our simple two period set-up, this slope reduces to $\frac{\lambda}{1+\lambda}$.¹² (For myopic consumers who do not internalize the future arrival of the second-generation technology, this slope becomes a horizontal line, and consumers in the upper right corner all choose (C_1, C_1) .¹³) Conversely, Figure 2 illustrates consumer choices for the region where $\nu_{id} < p_d$, so the relevant consumption bundles are $(0, 0)$, (C_1, C_1) , and $(0, C_2)$. Once again, the triangles in the upper right corner of the diagram depict the relative size of the consumer group that waits for C_2 versus the group that immediately adopts C_1 , with slope $\frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L}$.

Because the joint distribution of the idiosyncratic preference parameters $F(\nu_{id}, \nu_{ic_1}, \nu_{ic_2})$ is left unspecified, this model set-up allows for arbitrary correlations in consumer preferences across different technologies. We may expect, for example, some positive correlation in idiosyncratic consumer preferences for the two types of clean technologies, ν_{ic_1} and ν_{ic_2} . This distribution of consumer heterogeneity would result in a larger mass of consumers around the 45 degree line in Figures 1c and 2a, with a smaller mass of consumers in the upper left and lower right corners.

2.2 Impact of Policy on Overall Externality

As illustrated in Figures 1 and 2, the decentralized solution from individual optimization can be expressed as a series of thresholds above which consumers prefer one consumption bundle to another. Any policy intervention in this market will change these adoption thresholds. Due to the presence of externalities in this market, it is optimal to impose a general Pigouvian tax, where the effective tax rate on each technology is exactly equal to the externality associated with that product.

Proposition 1: *When the social planner is able to implement a market-wide tax schedule,*

¹²Given $A = L = 2$, the slope is equal to $\frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L} = \frac{\lambda(1-\lambda)}{(1+\lambda)(1-\lambda)} = \frac{\lambda}{1+\lambda}$

¹³Consumers may also apply a non-zero discount factor that is smaller than the social planner's discount factor (i.e., the consumer's discount rate is higher than would be socially optimal). Under the two-period model presented here, this partial myopia would result in a smaller slope and more consumers choosing (C_1, C_1) over (D, C_2) relative to what a social planner would choose (not accounting for externality-related mispricing). That is, too few consumers would be willing to wait for the improved second-generation technology relative to what would be socially optimal.

the optimal policy sets

$$t_d^* = \zeta_d, \quad t_{c_1}^* = \zeta_{c_1}, \quad \text{and} \quad t_{c_2}^* = \zeta_{c_2}$$

where t_j^* represents the optimal level of the effective tax on technology j .

Proof: See Appendix.¹⁴

Figure 3 shows how this optimal Pigouvian tax changes the threshold at which consumers choose one consumption bundle over another. For example, the threshold above which consumers prefer (C_1, C_1) to (D, D) decreases from $\nu_{ic_1} \geq p_{c_1} - p_d + \nu_{id}$ to $\nu_{ic_1} \geq p_{c_1} - p_d + \nu_{id} + \zeta_{c_1} - \zeta_d$ (Figure 3a).

In practice, however, policymakers may not have access to a general Pigouvian tax due to political economy or other constraints. Instead, policymakers often impose technology-specific policies, such as subsidies on early-generation clean products or standards that are only binding on the dirty product.¹⁵ These technology-specific policies may correct certain distortions in the market resulting from unpriced externalities, but inevitably introduce additional distortions by operating on only certain technologies. The relevant distortions depend not only on the products available to consumers in a particular period, but also on dynamic distortions related to the arrival of new generations of clean or efficient products.

For example, policymakers often use subsidies on early clean technologies as a second-best policy for reducing negative externalities. As illustrated in Figure 4, these subsidies will not only increase the mass of consumers adopting (C_1, C_1) over (D, D) , which is the desired policy outcome, but will also increase (C_1, C_1) adoption relative to $(0, 0)$, $(0, C_2)$, and (D, C_2) . Subsidy-induced adoption of the early clean technology — which still produces some negative externality, albeit less than the dirty technology — is a known phenomenon in a static setting. For example, subsidies for electric vehicles not only increase consumption of electric vehicles relative to gasoline-powered cars, but also increase consumption of electric vehicles relative

¹⁴In this stylized model, I do not separately model consumer decisions to reduce use of an externality-producing product from decisions not to consume the product; both are captured through the outside option.

¹⁵Policies that are *de jure* technology-nonspecific may be *de facto* technology-specific, such as performance standards applied to all products but in practice binding only for the dirty product.

to walking, biking, or using public transportation. However, the impact of subsidizing C_1 on adoption of the later clean technology C_2 only becomes apparent in a dynamic setting. This effect is important for understanding the long-run externality reduction in this market.¹⁶

When these dynamic effects are considered, the impact of a second-best policy on overall externality production may operate through several channels. First, second-best policies change the characteristics of consumers' choice sets in a static sense, such as by changing the price of the early clean technology through subsidies or by changing the characteristics of the dirty technology through standards. This effect is captured by the shifting threshold for adopting (C_1, C_1) over (D, D) or $(0, 0)$ in Figure 4. This channel is well-understood in the literature. Second, when the efficient technology is a durable good, increased adoption of the long-lived early clean product may change the total quantity of consumers in the market when the second-generation, more efficient product becomes available. This effect is captured by the shifting threshold for adopting (C_1, C_1) over (D, C_2) or $(0, C_2)$. Third, increased adoption of the early efficient product may also change the distribution of consumers available in the market for later consumption. This effect depends on the specific distribution of heterogeneous consumers over the preference space depicted in Figure 4. For example, if preferences are positively correlated across the clean products, early policy intervention may disproportionately affect those consumers who would otherwise adopt C_2 when it becomes available. All three of these channels are relevant in my empirical estimation below.¹⁷

¹⁶In another technology-specific policy intervention, standards that bind on the dirty product will change the externality associated with consuming D . As a direct or indirect consequence, these standards may simultaneously alter other characteristics of D or the distribution of idiosyncratic preferences for D , such as by increasing the cost of producing the dirty product and thereby its price. Even if consumers do not internalize the change in the externality resulting from the standard, these changes to other product characteristics will change the thresholds above which consumers are willing to adopt (D, D) over $(0, 0)$, (C_1, C_1) , or (D, C_2) and/or the mass of consumers above or below these thresholds.

¹⁷So far this model holds fixed the marginal cost of C_1 and C_2 ; I discuss below the potential for endogenous innovation response in this stylized set-up.

2.3 Efficient Second-Best Policy

Total social surplus from this market is given by the sum of consumer and producer surplus associated with each of the three technologies, as well as the sum of the total externality produced. For this stylized model, I assume that the price of each technology is equal to its marginal resource cost. The efficient level of a second-best technology-specific policy will balance the marginal benefits associated with correcting existing distortions with the marginal costs of introducing new ones.

To illustrate, I again consider a subsidy on an early-generation clean product. In the absence of any later-generation clean technologies, the efficient second-best policy will weigh the marginal benefits of correcting the relative prices of the dirty product versus the clean product against the marginal cost of further distorting the relative prices of the clean product versus no consumption. Define $s_{j,j'}$ as the share of consumers adopting $\{j_{t=1}, j'_{t=2}\}$.¹⁸

Proposition 2: *When the social planner is constrained to implement a tax schedule on C_1 only and there are two technologies (D and C_1), the efficient second-best pricing policy on C_1 is given by:*

$$\hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) [\underbrace{(\zeta_d - \zeta_{c_1}) \frac{\partial s_{d,d}}{\partial t_{c_1}}}_{\text{Impact on } C_1/D \text{ margin}} + \underbrace{\zeta_{c_1} (\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} + \frac{\partial s_{d,d}}{\partial t_{c_1}})}_{\text{Impact on } C_1/\text{no adoption margin}}]$$

Proof: *See Appendix.*

Given a general Pigouvian tax, the optimal pricing policy on C_1 is always a tax that internalizes the negative externality from consuming C_1 . However, when policymakers are constrained to impose a new price schedule on C_1 only, the optimal pricing policy may instead be a subsidy (a negative tax), provided that the avoided externality from shifting consumers to the clean product instead of the dirty product outweighs the additional externality from shifting consumers to the clean product instead of the outside option.

¹⁸The share of consumers adopting no technology in either period is given by $s_{0,0} = 1 - s_{d,d} - s_{c_1,c_1} - s_{d,c_2} - s_{0,c_2} \geq 0$.

When we also introduce later-generation clean technologies, the efficient second-best policy must also consider the additional cost of shifting consumers to the first-generation clean product instead of the second-generation clean product.

Proposition 3: *When the social planner is constrained to implement a tax schedule on C_1 only and there are three technologies (D , C_1 , and C_2), the efficient second-best pricing policy on C_1 is given by:*

$$\begin{aligned} \hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) [& \underbrace{(\zeta_d - \zeta_{c_1}) \frac{\partial s_{d,d}}{\partial t_{c_1}}}_{\text{Impact on } C_1/D \text{ margin}} + \underbrace{\zeta_{c_1} (\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} + \frac{\partial s_{d,d}}{\partial t_{c_1}} + \frac{\partial s_{d,c_2}}{\partial t_{c_1}} + \frac{\partial s_{0,c_2}}{\partial t_{c_1}})}_{\text{Impact on } C_1/\text{no adoption margin}} \\ & + \underbrace{(\frac{\zeta_d + \lambda \zeta_{c_2}}{1 + \lambda} - \zeta_{c_1}) \frac{\partial s_{d,c_2}}{\partial t_{c_1}} + (\frac{\lambda \zeta_{c_2}}{1 + \lambda} - \zeta_{c_1}) \frac{\partial s_{0,c_2}}{\partial t_{c_1}}}_{\text{Impact on } C_1/C_2 \text{ margin}}] \end{aligned}$$

Proof: *See Appendix.*

From this expression, we see that the efficient subsidy will be smaller when the policymaker must also consider the impact of the early policy on the later-generation clean technology, provided that the total externality from waiting for C_2 is less than the total externality from consuming C_1 .

We can also evaluate how the efficient second-best policy changes with various product and market characteristics. First, consider how \hat{t}_{c_1} changes when consumers are more likely to substitute away from C_1 to C_2 given an increase in the price of C_1 (i.e., $\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}$ and $(\frac{\partial s_{d,c_2}}{\partial t_{c_1}} + \frac{\partial s_{0,c_2}}{\partial t_{c_1}})$ both increase by an equal magnitude, holding all other terms fixed). If the externality from the (D, C_2) bundle is smaller than that of the (C_1, C_1) bundle, then greater substitutability between C_1 and C_2 will decrease the efficient level of a subsidy on C_1 .¹⁹ However, if the externality from the (D, C_2) bundle is larger, then the impact of greater substitutability from C_1 to C_2 will depend on the allocation of consumers between the $(0, C_2)$ and (D, C_2) bundles. Moreover, the efficient price schedule on C_1 is increasing in

¹⁹Since $t_{c_1} < 0$ represents a subsidy while $t_{c_1} > 0$ represents a tax, $\frac{\partial \hat{t}_{c_1}}{\partial z} > 0$ means that a subsidy is decreasing in magnitude or a tax is increasing in magnitude, given an increase in some parameter z .

the externality from consuming C_1 , holding constant the level of the externality associated with D , C_2 , and the outside option. Conversely, the efficient level of the C_1 pricing policy is decreasing in the externalities associated with either D or C_2 , again holding constant the externality associated with all other technologies. Lastly, definitively signing comparative statics on the efficient C_1 pricing policy with respect to the arrival time of C_2 or the lifetime of C_1 is not possible without making further assumptions on the distribution of ν_{id} , ν_{ic_1} and ν_{ic_2} . In the Appendix, I assume that these heterogeneous preference parameters are jointly uniform on the unit cube and discuss the conditions required for \hat{t}_{c_1} to be increasing or decreasing in A or L .

As noted above, I have modeled prices that are fixed for a given technology across all periods. It is also possible that prices in later periods depend on cumulative adoption in previous periods, through own-technology learning or cross-technology spillovers. In general, the presence of own-technology learning would be expected to augment the mechanical impact of first-generation durability on second-generation market size, by increasing the attractiveness of the first-generation technology in later periods as well. By contrast, the presence of cross-technology spillovers would mitigate this direct impact, as early adoption of the first-generation clean technology would make the second-generation technology more attractive for later consumers. While I do not model these impacts explicitly to preserve model tractability, the counterfactual simulations in my empirical analysis enable me to ask how the rate of second-generation price declines would need to change to offset the impact of various policy designs.

3 Efficient Lighting Market

Over the period 2010 to 2018, the market for efficient lighting saw rapid technological innovation and replacement of the incumbent incandescent bulb by three different efficient technologies – halogens, CFLs, and LEDs. Traditional incandescent bulbs had existed for

decades and were well-liked by consumers; these bulbs were also highly inefficient, losing up to 90% of their input energy as waste heat. Halogens used the same underlying technology as traditional incandescent bulbs but were approximately 30% more efficient.²⁰ CFLs represented the first mass market efficient light bulb that used a fundamentally different underlying technology. These bulbs initially emerged in the residential lighting market in the 1990s, with significant improvements in cost and quality over the 2000s. Finally, LEDs entered the residential lighting market in the early 2010s, particularly following the commercial introduction of Philips Lighting’s 60W-equivalent A-shaped LED bulb in early 2012.²¹ The cost and performance of LED bulbs improved considerably over the study period, such that LEDs now dominate CFLs on virtually all characteristics that consumers (or the social planner) might care about.

All four of these technologies could be used for the same general purpose applications, such as household lamps.²² In this paper, I focus specifically on general purpose “A shape” light bulbs, which could be used in a variety of standard applications.²³ A representative example of these bulbs for each of the four lighting technologies is given in Figure 5. Consumers distinguish different types of A-shaped bulbs by their light output, generally denominated in terms of the power consumption of an equivalent incandescent bulb. Common types include 40W-equivalent, 60W-equivalent, 75W-equivalent, and 100W-equivalent bulbs; Table 1 pro-

²⁰A useful analogy here might be a highly fuel-efficient internal combustion engine vehicle – the underlying technology is the same as the traditional product, but various engineering improvements drove efficiency gains (at a cost). Most general purpose halogen bulbs sold during this period exactly met the federal efficiency standards for lighting but no more.

²¹This bulb won the Department of Energy’s L-Prize competition, awarded to the first lighting manufacturer to develop a 60W-equivalent household LED with certain desirable specifications. The scientific breakthrough that made it possible to use LEDs for general purpose lighting was the Nobel Prize-winning discovery of the blue LED in the early 1990s. This discovery opened the door to using LEDs for screens, including in televisions, computers, tablets, and phones, which created spillover benefits in the general purpose lighting market. LEDs use a fundamentally different technology to produce light than traditional incandescent bulbs, converting electricity directly into light particles rather than relying on the intermediate step of heating a filament until it glows, which produces significant waste heat in addition to useful light.

²²As throughout the paper, I use the term “lamp” here in the colloquial sense to refer to the item of furniture containing a light bulb and emitting light. In the lighting industry, however, “lamp” is a technical term that is used to refer specifically to the light bulb, while “luminaire” refers to the lamp plus all the other components required for a functional device, such as the electrical connection, the lamp socket, and so forth.

²³I also include spiral CFLs which, although not officially classified as A-shaped bulbs, had compatible screw bases and were marketed as replacements for traditional A-shaped bulbs.

vides the actual wattage for each of these bulb types by technology. A primary determinant of CFL and LED quality during this period was the light “color” – that is, whether the bulb produced a warm light or a cooler light. Consumers generally preferred warmer colored bulbs, as this light color more closely approximated incandescent light, but warmer light was technologically more difficult to produce using CFL or LED technology. Over the course of the study period, much of the technological innovation in LEDs related to creating warm light bulbs at lower cost.

Bulbs using different technologies also varied in their cost structure. The upfront purchase price was highest for the most efficient light bulbs, CFLs and LEDs, but operating costs (based on electricity consumption from use) were lower and lifetimes were higher compared to the less efficient and less durable incandescent and halogen bulbs. As a result, consumers could generally save on the overall cost of owning and operating a bulb by purchasing more efficient lighting technologies, but it was well documented that consumers often balked at the relatively high upfront cost. The cost of producing LEDs declined significantly over the study period; once available in the market, the median price of LED bulbs declined at a rate of approximately \$3.25 per bulb per year (see Figure 6). The most significant cost reductions in CFLs occurred prior to the study period, and median prices increased slightly as federal efficiency standards were phased in and then declined as LEDs gained greater market share.

During these years, two sets of policies shaped the evolution of the U.S. lighting market in particular: federal efficiency standards and state and utility subsidies for efficient lighting. Provisions in the Energy Independence and Security Act of 2007 banned the manufacture or import of general purpose light bulbs below certain efficiency levels, which no incandescent bulbs were capable of meeting. This federal efficiency standard therefore resulted in the removal of traditional incandescent bulbs from the U.S. market; halogens emerged in their stead, with most halogens sold during this period meeting the exact minimum efficiency permitted under the law. The standard was implemented in three phases, affecting 100W-equivalent bulbs from January 1, 2012; 75W-equivalent bulbs from January 1, 2013; and

40W- and 60W-equivalent bulbs from January 1, 2014.²⁴ California implemented the same set of efficiency standards one year early, between 2011 and 2013, as part of a legislative compromise given that the state had already adopted its own lighting efficiency requirements that were preempted by the EISA of 2007.²⁵ Figure 13 shows the rapid decline in sales of each type of incandescent bulb, as retailers sold any existing inventory after the standards came into effect, paired with the rapid rise in sales of halogen bulbs. A second phase of lighting efficiency standards was set to come into effect on January 1, 2020, which would have effectively removed halogens from the market, leaving only CFLs and LEDs.²⁶ However, these more stringent standards faced legal challenges under the Trump administration, and their future has been uncertain since mid-2019.

In contrast to nationwide efficiency standards, subsidies for efficient lighting were quite local, administered at the state or utility level, with the generosity of subsidy programs varying considerably across the country. These subsidies were often driven by energy efficiency policies at the state level, most notably the passage of dozens of Energy Efficiency Resource Standard (EERS) programs and various changes in cost recovery mechanisms for utilities that incentivized electricity demand reductions. These state and utility programs included a wide range of incentives for efficient appliances, building weatherization, and other measures, targeted at residential, commercial, industrial, and public sector electricity users. Nonetheless, efficient lighting was generally the most cost-effective option available (Hoffman et al., 2017). As a result, subsidies for efficient lighting were often a substantial portion of the purchase price, averaging around \$1 per bulb for general purpose CFLs and \$7 to \$10 per bulb for LEDs.

The particular form of these subsidies differed by utility or state administrator. In some areas, subsidies took the form of manufacturer buydowns or direct discounts, where

²⁴As a largely symbolic gesture, Congress voted to defund the enforcement of these provisions in late 2011, but the major lighting manufacturers all continued to comply with the standard.

²⁵Nevada had also adopted state-wide lighting standards by 2007, but declined to implement the federal standards early.

²⁶The second phase of the standards would have also applied efficiency standards to a broader set of light bulbs, such as globe-shaped bulbs, chandelier bulbs, and several types of directional bulbs.

the sponsoring agency would negotiate a lower price with retail partners (and sometimes manufacturers). Consumers would then pay this lower price directly, which would also be reflected in store labels, perhaps with additional signage acknowledging the role of the sponsor. The store would then submit sales data for covered products to the sponsor at the end of some pre-defined period to be compensated for the lower prices charged to consumers. By contrast, under instant rebate programs, consumers would often use in-store coupons or other paperwork to receive discounted prices at the point of purchase. Finally, with mail-in rebates, consumers were required to save receipts showing purchases of covered products and mail them to the sponsoring organization, often to receive a credit on their utility bill.

The combination of widespread policies to promote the adoption of new lighting technologies and rapid technological development in this sector resulted in substantial changes in the composition of lighting sales over 2010 to 2018. While incandescent bulbs constituted 60 to 70% of the general purpose lighting market prior to the federal efficiency standards, with CFLs as the primary technological competitor, halogens constituted 40 to 50% of the general purpose lighting market by 2015. General purpose LEDs did not yet exist as a commercial product at the start of the study period, but commanded approximately 35% of the national market share by 2018. Figure 7 shows the share of general purpose bulb shipments by technology type and year, using data from the National Electrical Manufacturers Association (NEMA). Figures 8, 9, 10, and 11 illustrate how this aggregated time series belies substantial regional variation in the rate of halogen, CFL, and LED adoption, using retailer sales data from Nielsen.

4 Model

4.1 Consumer Utility

To understand the impact of policy design and timing on the evolution of the efficient lighting market, I first model consumer utility from general purpose lighting, focusing on the

residential lighting market. I assume that utility is independent across lightbulb purchases and that a product’s lifetime utility is realized at the moment of purchase. That is, I do not directly model consumer stockpiling and at-home inventories. I define each market as a U.S. county and each period as a quarter of a year; a consumer’s choice set consists of all light bulbs sold in a given county over a given quarter, following Villas-Boas (2007), Asker (2016), and Dubois and Bonnet (2010). Based on survey data, I assume that consumers come to the market to purchase a light bulb when an existing light bulb fails.

The expected indirect lifetime utility for consumer i purchasing product j in market m at period t is therefore given by:

$$U_{ijmt} = \mathbf{x}_{jmt}\boldsymbol{\beta}_i - \alpha p_{jmt} + \xi_{jmt} + \epsilon_{ijmt} \quad (2)$$

Here \mathbf{x}_{jmt} gives the product’s observable characteristics, which I define to be bulb type (40W-equivalent, 60W-equivalent, 75W-equivalent, or 100W-equivalent); technology (incandescent, halogen, CFL, or LED); annual operating costs; and brand, defined as a dummy for whether the product is a private label brand or not.²⁷ Product characteristics that are properties of the bulb technology or wattage-equivalence, such as actual wattage, lifetime, or mercury content (in the case of CFLs), are captured by product category dummies that interact technology and wattage-equivalence.²⁸ p_{jmt} gives the product price (per bulb) charged to consumers, taking the weighted average across multiple stores in a county. For states or utility territories implementing buydown discount programs, the price charged to consumers may be inclusive of these special discounts. ξ_{jmt} gives unobservable product quality at the product-county-quarter level. Finally, ϵ_{ijmt} gives the standard Type 1 extreme value

²⁷Throughout this discussion, I use “product category” to refer to each combination of wattage-equivalence and technology. My sample includes 16 product categories, constructed from four wattage-equivalences and four technologies.

²⁸To first order, these characteristics are all derived directly from the technology and wattage-equivalence of the bulb, and variations within a product category were generally second-order compared to variations across product categories. For example, incandescent lifetimes might vary from 750 hours to 2,000 hours for “double life” bulbs, while LED lifetimes might vary from 15,000 to 25,000 hours – but in all cases, LED and CFL lifetimes were multiple times longer than incandescent and halogen lifetimes.

error term, which I assume to be independent and identically distributed across consumers, products, counties, and periods. The coefficients α and β_i measure the impact on consumer utility of price and other observable characteristics, respectively.

My preferred model specification allows consumer utility from efficient lighting to depend on observable demographic characteristics. I interact various consumer demographic characteristics (education, Democratic voter, concern about climate change, support for Renewable Portfolio Standard policies) with an indicator for whether a given bulb is “efficient,” defined as either a CFL or LED.²⁹ In alternative specifications, I allow for a separate interaction between demographics and CFL and LED indicators, to allow for more flexible preference relationships between CFLs and LEDs, though I do not find that results change meaningfully in this more complex specification.

I also include flexible time variables interacted with technology. My preferred specification includes an indicator for periods before the federal efficiency standards came into effect, interacted with each type of incandescent, halogen, and CFL bulbs. I also include a linear time trend and a quadratic time trend interacted with each type of LED bulb, to account for changes in consumer awareness of LEDs over time and observed changes in LED product quality (e.g., improvements in light color). An alternative specification includes time period dummies instead of time trends interacted with LED product categories.

We can decompose the expression for expected indirect lifetime utility in Equation 2 into components that are common to all consumers within a market and components that are

²⁹Renewable Portfolio Standards were state-level policies for encouraging renewable energy deployment. RPS policies were similar in spirit to state-level Energy Efficiency Resource Standards, which were often the policy impetus for efficient lighting subsidies.

heterogeneous across consumers:³⁰

$$U_{ijmt} = \underbrace{\mathbf{x}_{jmt}\boldsymbol{\beta} - \alpha p_{jmt} + \xi_{jmt}}_{\delta_{jmt}} + \underbrace{\mathbf{D}_i \sum_k \boldsymbol{\pi}_k x_{jkmt}}_{\mu_{ijmt}} + \epsilon_{ijmt}$$

where D_i represents consumer demographic characteristics and k indexes product characteristics. Here I follow the now-standard notation from Berry et al. (1995) in using δ_{jmt} to denote common components of indirect utility and μ_{ijmt} to denote idiosyncratic components. I will use $\boldsymbol{\theta}_1 = [\boldsymbol{\beta}, \alpha]$ to denote the estimated parameters that enter the common component of indirect utility, and $\boldsymbol{\theta}_2 = \boldsymbol{\pi}$ to denote the estimated parameters that enter the idiosyncratic component.

Given this indirect utility specification, consumer i will choose product j over j' whenever $U_{ijmt} > U_{ij'mt}$. Consumers also have the option to choose the outside option, which I define to be lighting purchased at non-Nielsen stores; I assume $U_{i0mt} = \epsilon_{i0mt}$. This set-up yields the following familiar expression for logit choice probabilities, where s_{ijmt} is the probability that consumer i chooses product j :

$$s_{ijmt} = \frac{\exp(\delta_{jmt} + \mu_{ijmt})}{1 + \sum_{j' \in J_m} \exp(\delta_{j'mt} + \mu_{ij'mt})}$$

which yields the following expression for market shares:

$$s_{jmt} = \int s_{ijmt} \cdot f_{mt}(\boldsymbol{\mu}_{imt}) d\boldsymbol{\mu}_{imt}$$

Note that the distribution of idiosyncratic utility $f_{mt}(\boldsymbol{\mu}_{imt})$ depends on both the market and the time period. In standard discrete choice models, it is typical to ignore the time-

³⁰As in the simple theory model presented in Section 2, a consumer's indirect utility from a given product can be decomposed into a component common to all consumers and an idiosyncratic component unique to each consumer. In the empirical model, however, the common component incorporates product characteristics other than price, and the idiosyncratic component depends on interactions between demographics and product characteristics (and potentially random preference draws).

dependence of consumer heterogeneity, or to assume that it arises exogenously to the model. As I describe further below, I explicitly model how the distribution of consumer preferences within a market evolves endogenously as a function of past purchases.

4.2 Consumer Entry into Market

While consumer utility from light bulb purchases is well approximated by a static model, I must also account for the fact that the history of past lighting purchases affected the distribution and rate of consumers coming to the market to purchase new light bulbs. The expected lifetime of general purpose bulbs varied considerably across technologies. While a typical rated lifetime of an incandescent bulb was about 1,000 hours (or 11 months, given three hours of use per day) and a typical rated lifetime of a halogen bulb was about 2,000 hours (1.8 years), the rated lifetime of CFLs and LEDs was considerably longer, at about 8,000 hours (7.3 years) and 15,000 to 25,000 hours (13.7 to 22.8 years), respectively.³¹

This difference in bulb lifetime by technology has two key implications for my model. First, consumers whose homes were filled with short-lived incandescent bulbs could be expected to come to the market to replace expired bulbs far more frequently than consumers whose homes were instead filled with longer-lived CFLs. Consequently, one possible impli-

³¹In my primary model specification, I use 1,200, 2,000, 8,000, and 25,000 hours as representative lifetimes for incandescents, halogens, CFLs, and LEDs, respectively. Two comments are in order about these representative lifetimes. First, actual CFL lifetimes may have been considerably shorter than rated CFL lifetimes, because frequent on-off switching by residential users diminished the effective lifetime, a technological property of CFLs but not other types of light bulbs. Following the U.S. Department of Energy (2016), I use modified failure rates that account for the impact of on-off cycle time on actual CFL lifetimes; I also conduct sensitivity analyses that use the full 8,000-hour lifetime instead. Second, consumer reports generally cited 2,000 hours or 1,000 to 3,000 hours as typical halogen lifetimes, so I have adopted 2,000 hours as a midpoint of these possible lifetimes; see, for example, <https://www.energymaine.com/at-home/lighting-solutions>; <https://web.archive.org/web/20151110161940/https://www.energy.gov/energysaver/how-energy-efficient-light-bulbs-compare-traditional-incandescents>; <https://www.energyrating.gov.au/document/factsheet-light-bulb-buyers-guide>; and https://www.thelightbulb.co.uk/resources/light_bulb_average_rated_life_time_hours. Technologically, halogen lifetimes may differ from incandescent lifetimes for comparable bulbs because they are able to operate with higher filament temperatures and chemical properties of halogen bulbs allow tungsten to be continually redeposited on the filament instead of accumulating inside the bulb shell. An alternative account of halogen lifetimes suggests that by 2019, only 20% of halogens had 2,000-hour lifetimes while 80% had 1,000-hour lifetimes, resulting in an average rated life of 1,200 hours (National Electrical Manufacturers Association, 2019). Shorter halogen lifetimes would lessen the impact of early standards adoption on the total quantity of LEDs sold, by lessening the impact on the total market size later in the study period.

cation of early support for CFLs was removing a substantial share of likely early adopters of LEDs from the market for several years — for example, those consumers with strong preferences for efficient light bulbs. Second, the residential general purpose lighting market saw substantial declines in total light bulb shipments over this period, as bulbs lasted longer and consumers therefore needed to replace their bulbs less frequently, which my model must be able to rationalize. NEMA data included in regulatory filings reveals that total general purpose lighting shipments and imports declined by 38.5% from 2011 to 2017 (NEMA, 2016).

As a result, I model both the absolute number of consumers coming to market in a given quarter and the distribution of consumers’ idiosyncratic preferences as a function of past purchases by technology type. Given my assumption that consumers enter the market primarily to replace failed bulbs, the expected number of bulbs purchased in t is given by:³²

Expected number of bulbs purchased in t = Expected number of bulb failures in t

$$\begin{aligned}
& \sum_i \left(\underbrace{\sum_{j \in \mathcal{I}_t} h(a_{ijt}, l_I, hou_{ij})}_{\text{Incandescent}} + \underbrace{\sum_{j \in \mathcal{H}_t} h(a_{ijt}, l_H, hou_{ij})}_{\text{Halogen}} + \underbrace{\sum_{j \in \mathcal{C}_t} h(a_{ijt}, l_C, hou_{ij})}_{\text{CFL}} + \underbrace{\sum_{j \in \mathcal{L}_t} h(a_{ijt}, l_L, hou_{ij})}_{\text{LED}} \right) \\
&= \frac{\quad}{|\mathcal{I}_t| + |\mathcal{H}_t| + |\mathcal{C}_t| + |\mathcal{L}_t|}
\end{aligned} \tag{3}$$

Here $h(\cdot)$ gives the hazard rate for bulb failure.³³ This probability of failure is in turn a function of a bulb’s age in years (a_{ijt}), which depends on when the bulb was purchased; the bulb’s rated lifetime (l_j); and the number of hours that the bulb is used per day (hou_{ij}).

³²Of course, this set-up does not account for the role of new household creation alongside the replacement of existing household bulbs. Additional robustness tests are underway to account for the rate of population growth in a county, assuming that the “new” consumers exhibit the same distribution of preferences as the underlying population before accounting for selection due to past purchases. In this alternate version of the model, the “existing” consumers continue to behave according to the baseline model described in this section, entering the market to replace failed bulbs.

³³I assume that this hazard rate generally takes the same form across technologies, conditional on bulb lifetime. The exception is CFLs, where effective lifetime was notably shortened by more frequent on-off switching. I incorporate the estimated prevalence of on-off cycles shorter than 30 minutes to adjust downward the expected time-to-failure of CFL bulbs. This characteristic did not affect other bulb technologies. I discuss the construction of bulb lifetimes in greater detail in the data appendix.

I adopt the simplifying assumption that the daily hours-of-use does not depend on bulb technology, while the rated lifetime depends entirely on bulb technology.³⁴ \mathcal{I}_t , \mathcal{H}_t , \mathcal{C}_t , and \mathcal{L}_t denote the existing installed base of incandescent, halogen, CFL, and LED bulbs, respectively; I use an i subscript on these terms to denote household-specific installed bases by technology.³⁵ Insofar as consumer i has a larger base of long-lived bulbs, this consumer will enter the market less frequently in subsequent periods.³⁶

5 Estimation

5.1 Sample Construction

To estimate this demand model for general purpose lighting, I rely primarily on the Nielsen retail scanner dataset over the period 2010 to 2018. This dataset provides prices and quantities information for lightbulbs sold at individual stores on a weekly basis. The dataset covers over half of total national sales at grocery stores and drug stores, and about one-third of sales at mass market retailers, with some heterogeneity by geographic area. Coverage of convenience stores is much more limited, and I drop these stores from my sample. Individual stores are anonymized in the dataset, but geographic information is provided at the county level; the dataset also links stores associated with the same retailer or parent company.

Individual lighting products are identified by their 12- or 13-digit Universal Product

³⁴Here I am ignoring potential correlations between intensity of lighting energy use and willingness-to-adopt energy-efficient lighting technologies; see Dubin and McFadden (1984) for further discussion of potential biases that this simplifying assumption may introduce.

³⁵This assumption about constant hours-of-use across technology types also means that I am effectively assuming away a behavioral response as consumers adopted more efficient lighting. This approach seems reasonable in the short and medium term, as lighting costs constituted a small overall fraction of household budgets and sufficient time had not yet elapsed for residential lighting installations to change substantially. Survey data from the DOE's Lighting Market Characterization Reports suggests that the average daily operating hours per lamp in the residential sector (1.8 hours in 2010 and 1.9 hours in 2015) remained relatively constant over the study period (Navigant Consulting, 2012, 2017). When considered over many decades, however, human consumption of lighting has increased substantially as technology has improved, so we would expect to see a stronger behavioral response over a longer time period; see Nordhaus (1996).

³⁶I assume that each consumer has the same number of bulbs and use the average number of bulbs per household (50.4) to scale counterfactual results; this scaling factor does not influence estimation.

Code (UPC). This code allowed me to connect sales data with a secondary dataset of product characteristics that I assembled from historical product catalogs, third-party retailers, and other sources. For UPCs that could not be externally validated (e.g., for private label brands where UPCs were obscured to maintain store anonymity), I relied on product descriptions contained within the Nielsen data to identify lighting technology, bulb type, and other key characteristics. Following this matching process, I limited the sample for analysis to 100W-equivalent, 75W-equivalent, 60W-equivalent, and 40W-equivalent medium base Type A bulbs or medium base spiral CFLs.³⁷ Additional details about the light bulb sample construction are provided in the data appendix.

Table 2 provides summary statistics for light bulbs included in the final sample, and Figure 12 shows the technology share of sales in the Nielsen data. Note that the technology profile by year does not exactly match national shipments data from NEMA, as presented in Figure 7. Instead, the Nielsen dataset includes a higher share of incandescent and halogen bulbs relative to CFLs and LEDs. There are several possible explanations for this discrepancy. First, the Nielsen data includes a higher share of grocery stores and drug stores relative to mass market retailers, and anecdotally, CFLs and LEDs commanded higher market share at mass market retailers and home improvement stores. Second, the extent of coverage in the Nielsen data varies across geographic area. Finally, some bulbs included in the national shipments data were sold to commercial customers or were sold online rather than in retail stores. While I use population weights in my counterfactual simulations to account for the under- or over-representation of some counties among Nielsen retailers, my results should be interpreted relative to technology shares observed at Nielsen retailers rather than in other sales or shipments data.

³⁷Because more than 95% of Type A lamps are used in the residential sector, it is reasonable to assume that this sample reflects purchases by residential consumers (Navigant Consulting, 2017). In my sample construction, I elected not to include decorative bulbs, such as globe-shaped or candle-shaped bulbs; directional bulbs, such as floodlight or reflector bulbs; or specialty bulbs, such as bug lamps or colored bulbs. These types of light bulbs were not covered by the first phase of the federal efficiency standards and generally had different use cases from standard bulbs. I also drop smart LEDs from my final sample, since these bulbs provided other services such as Wi-Fi connectivity.

I filtered the sample used for estimation by the type of rebate (if any) offered by the utilities serving residential customers in a given county. Using data published annually by the U.S. EPA on Energy Star lighting programs, I manually matched utilities by rebate type to information on the geographic coverage of a given utility territory, reported at the county level on the Energy Information Administration (EIA)’s Form 861.³⁸ In some cases, lighting programs were implemented by state agencies or consortia of several utilities, where the contemporaneous list of member utilities was generally included in the individual EPA reports. I then dropped counties with mail-in or instant rebates for residential customers purchasing general purpose bulbs in order to reduce measurement error in the price variable used for estimation, as these rebates generally would not be reflected in the price paid by the consumer at the register.³⁹

Finally, I filtered the sample by the extent of coverage in the Nielsen data. After aggregating across all stores in a county, I dropped counties below the fifth percentile of bulbs sold per quarter (253 bulbs) or below the fifth percentile for number of unique products sold per quarter (8 products, where a product is defined by technology, wattage-equivalence, and brand). I also dropped counties that did not appear in the Nielsen data in certain years, in order to have a complete time series for modeling the evolution of the installed base of lighting. My final sample includes 852 counties, 36 quarters, and 560,098 product-county-quarter observations.

To the primary data set on light bulb purchases and product characteristics, I bring additional demographic data to capture heterogeneity in the distribution of consumer preferences across geographic areas. I use four different demographic variables in alternative model specifications. Information on the county-level share of college-educated individuals comes from the American Community Survey. Information on the share of individuals expressing belief

³⁸Energy Star was a federal certification program for many types of efficient appliances, including efficient lighting (principally CFLs and LEDs). Data on lighting programs comes from U.S. Environmental Protection Agency (2010-2018).

³⁹In this process, I retained counties with manufacturer buydown programs or no active rebate program, as well as counties with mail-in or instant rebates for commercial or specialty customer classes (e.g., K-12 education) or for specialty bulbs (e.g., holiday lights).

in human-caused climate change or support for state Renewable Portfolio Standards comes from the Yale Program on Climate Change Communication, which estimates county-level variation in climate attitudes (Howe et al., 2015). Lastly, information on the county-level share of Democratic voters is taken from Dave Leip’s Atlas of U.S. Presidential Elections (Leip, 2017).

5.2 Identification

Demand estimation follows the now-standard procedure outlined in Berry et al. (1995), although I vary the nested fixed point routine to account for endogenous coming-to-market, as described further below. The consumer choice of light bulbs is static and utility is assumed to be separable across multiple lighting purchases. I allow for the fact that retailers and/or manufacturers may observe the realization of ξ_{jmt} prior to determining the optimal markup. To address this potential correlation between ξ_{jmt} and p_{jmt} , which could lead to biased estimates of α , I adopt the traditional instrumental variables strategy used in BLP and other papers in this literature. Given a vector of instruments \mathbf{Z}_{jmt} , the identifying assumption is:

$$E[\xi_{jmt}|\mathbf{Z}_{jmt}] = 0$$

In this setting, three different sets of instruments may be appropriate.⁴⁰ First, I use cost shifters at the manufacturer level to capture exogenous changes in input prices for manufacturing light bulbs, including price indices for semiconductors, fluorescent ballasts, and certain rare earth minerals. Semiconductors were a key input in the manufacture of LED chips, and LEDs benefited from the continued decline in semiconductor manufacturing

⁴⁰I also test a fourth set of instruments, following Hausman (1996). I construct these instruments using mean quarterly prices by product category (i.e., technology by wattage-equivalence), using all other census divisions except the one in which a given county is located. The identification assumption here is that there are no national shocks in demand for a particular type of light bulbs, beyond what is captured in my time trends. For my specification with prices measured in levels rather than logs, I find that the estimated coefficient on price is meaningfully reduced in magnitude when I include Hausman instruments, so I elect not to include them in my final specification.

costs over the study period. Fluorescent ballasts were a key input for CFL bulbs, as were rare earth minerals used to make phosphors coating the inside of bulbs.⁴¹ Following Villas-Boas (2007), I interact these instruments with a technology dummy, to reflect the fact that changes in these input prices are likely to affect different bulb technologies in different ways. Second, I use cost shifters at the retailer level, to capture exogenous changes in retailers' operating costs, including average retail wages, commercial real estate price indices, and diesel prices as a proxy for transportation costs. While manufacturing cost shifters varied over time, retail cost shifters varied both over time and across geographic markets. Table 3 provides summary statistics for each of these cost shifters; additional discussion of the data used to construct these cost shifters is provided in the data appendix.

As a third set of instruments, I follow Gandhi and Houde (2019) in building differentiation instruments using product characteristics in a given market. Because the Nielsen data provides a sample of sales in each market, rather than a full census, I use two different approaches for constructing differentiation instruments, both of which are intended to capture a product's proximity to other products in characteristics space. First, I calculate the fraction of all other products in the market, using store-level product offerings, that share a given characteristic (technology, wattage-equivalence, or brand). This measure is intended to capture the frequency with which a consumer might encounter a product with a given characteristic across all stores in the market, which would be expected to influence markups. Second, I count the number of other products in a given store with a particular product characteristic, and then take the average of this number across all stores in the market. This measure is intended to capture a product's similarity to other products available in the same store, which would also be expected to influence markups. Both of these sets of differentiation instruments depend on the assumption that store-level stocking decisions are made before unobserved shocks to demand or product quality are revealed. This assumption may be more appropriate at chain retailers where procurement decisions are made centrally

⁴¹Rare earth minerals were also used in LEDs, but in much lower quantities, so I use CFL-specific weights in constructing these instruments.

but pricing decisions are made locally, or in stores where inventory decisions are made less frequently than pricing decisions.

5.3 Coming-to-Market Algorithm

To account for endogenous coming-to-market, I modify the standard fixed point algorithm commonly used in discrete choice demand estimation (the “inner loop” of BLP-style estimation). In the standard BLP algorithm, distributions of consumer heterogeneity are discretized by taking a finite number (ns) of draws from distributions of observed demographics or unobserved preference heterogeneity. These draws are used to simulate the distribution of individual choice probabilities, which are then summed to calculate predicted market shares. In the standard algorithm, the weights placed on individual choice probabilities are exogenous to the model, with the econometrician either weighting each draw equally (“frequency weights”) or using importance sampling methods to choose weights that improve the performance of the estimator. By contrast, the modified algorithm employed here endogenously updates the weights on each consumer draw, based on the probability that an existing bulb will fail in each period. In the discussion below, I refer to these consumer draws as consumer “types” since each draw represents a mass of consumers with certain characteristics.

I begin with an initial distribution of remaining time-to-failure for each consumer type’s installed base of light bulbs at the start of the study period. Given the assumption that consumers enter the market to replace failed bulbs, I use this initial distribution to simulate the share of each consumer type entering the market in the first period. Let w_{imt} denote the share of consumers of type i entering market m in period t . I then apply the standard BLP fixed point algorithm in the first period, which allows me to predict individual choice probabilities for each consumer type. Next, the individual choice probabilities for each technology are used to update the remaining time-to-failure for the installed base of bulbs for each consumer type, using engineering data on failure rates and assuming that new bulbs are installed in the same period that they are purchased. This updated distribution allows

me to predict the share of each consumer type entering the market in the second period. I repeat this forward simulation for all periods in my study period, and then follow the standard BLP approach of using non-linear search methods to solve for the value of θ_2 that sets the gradient of the objective function as close as possible to 0.⁴² The full details of my estimation algorithm are provided in the Appendix.

As noted above, my estimation algorithm requires as an input the share of each type of consumer entering the market in the first period. To solve this initial condition problem, I simulate the time-to-failure of the installed base of light bulbs at the start of the study period, using quarterly national shipments indices for CFL, halogen, and incandescent bulbs over 2001 to 2010.⁴³ In the absence of data on how these earlier market shares vary by consumer demographics, I adopt the simplifying assumption that these past purchases are uniform across consumer types. Therefore, I assume that there is an equal share of each consumer type entering the market in the first period.⁴⁴

Finally, while my estimation algorithm enables me to predict overall market shares using endogenous weights across consumer types, I am also able to use these weights to calculate how the total quantity of bulbs sold evolves over time in each market. Modeling both the distribution of consumer characteristics represented in the market and the overall number of consumers is essential for my counterfactual policy simulations. Each of my ns draws of consumer types represents some number of real-world households, which each has some number of light bulbs. I use census data on the number of households per county and lighting market characterization data on the number of bulbs per household to determine the total number of bulbs (N_{im}) represented by each of my ns draws.⁴⁵ The total number of type j

⁴²Note that the gradient of the objective function differs from the standard BLP optimization problem because I account for endogenous coming-to-market. In the appendix I derive the analytic gradient of the objective function that incorporates coming-to-market.

⁴³LEDs did not become available until after the start of the study period. To account for the fact that I observe fewer CFL sales in the Nielsen data during the study period, relative to what is reported in the national shipment indices, I assume that Nielsen stores account for a consistently lower portion of CFL sales, and therefore proportionally adjust downwards the market share of CFLs from the national indices during the pre-period. This step smooths the share of bulbs replaced from the pre-period to the study period.

⁴⁴For reference, I use $w_{imt} = 9.48\%$ in the first period, for all consumer types and all markets.

⁴⁵I hold N_{im} constant within each market m over the study period. This approach ignores population

bulbs purchased in a particular market in a particular period is therefore given by:

$$\text{Number of bulbs purchased}_{jmt} = \sum_{i=1}^{ns} N_{im} w_{imt} s_{ijmt}$$

While previous research has considered the impact of endogenous coming-to-market with all the complexity of a dynamic demand model (Lee, 2013; Gowrisankaran and Rysman, 2012), to the best of my knowledge, this paper is the first to integrate evolving consumer heterogeneity into a BLP-style static demand model in an internally consistent manner. Before applying this estimation procedure to lighting data, I use simulated demand data to validate my modified algorithm and examine the impact of evolving consumer heterogeneity on static demand estimation, following the basic simulation procedure outlined in Conlon and Gortmaker (2020). These simulations suggest that the bias in estimated coefficients from ignoring the changing distribution of consumer preferences may be substantial. This procedure may be used to estimate discrete choice demand systems for other product markets during periods of new technology adoption, as well as for goods such as performance tickets where early consumers may differ endogenously from later consumers.

5.4 Estimation and Results

I first present results from logit demand models without consumer heterogeneity, in Table 4. I include several sets of specifications: using price per bulb in levels and in logs; using linear and quadratic time trends interacted with LED product categories and time dummies interacted with LED product categories; and using ordinary least squares and instrumental variables for estimation. I find that coefficients on price and the private label dummy have the expected negative sign. For the instrumental variables specifications, the F-statistic ranges from 270 to 436, indicating a strong first stage. The coefficient on operating costs is

growth or shifting demographics within a county over time. However, these variables are strongly correlated over time during the nine-year study period. For example, the correlation between 2010 county population and 2018 county population is 0.9984.

negative in some specifications, as would be expected, but positive in others; in my preferred logit specification, with prices specified in levels and instrumental variables (Column 2 of Table 4), this coefficient is not statistically distinguishable from zero. Given the inclusion of technology by wattage-equivalent dummy variables in all specifications, the coefficient on operating costs is identified from geographic variation in electricity rates. This variable likely contains measurement error, since utility territories do not exactly coincide with county boundaries and I construct a weighted average of local electricity rates within a county, which is the most granular geographic information provided in the Nielsen scanner data. It is also likely that these differences in operating costs are not fully salient to consumers.⁴⁶

Table 5 presents estimates from seven different model specifications with consumer heterogeneity. Columns 1 through 4 interact demographic variables with an indicator variable for whether a bulb is “efficient,” defined as either a CFL or LED. These demographic variables include expressing a belief in human-caused climate change; being college educated; being a Democratic voter; or expressing support for Renewable Portfolio Standards. Column 5 allows for a separate interaction between expressing belief in human-caused climate change and CFL and LED products, respectively, to allow for a more flexible relationship between demographics and these two types of bulbs. Column 6 uses time dummies rather than time trends interacted with LED product categories. Column 7 uses prices per bulb in logs rather than in levels. In my counterfactual simulations, I adopt the specification in Column 1 as my preferred demand system. In addition to the estimates reported here, I also test the sensitivity of my results to dropping all counties in California, since this state acted before the rest of the U.S. in implementing a second set of stricter lighting efficiency standards

⁴⁶There is a long literature in economics on the salience of operating costs for consumers purchasing durable appliances, vehicles, and other products. Findings from this literature suggest that consumers do not always fully internalize operating cost savings from purchasing more efficient products, though there is heterogeneity across products and consumer characteristics. One explanation is heterogeneity in consumer discount rates (Newell and Siikamäki, 2015). Another explanation is lack of information or inattention, though evidence here is mixed (Davis and Metcalf, 2016; Allcott and Knittel, 2019; Allcott and Taubinsky, 2015). As a robustness test in my setting, I examine whether my counterfactual results change when I drop the operating cost variable from my demand specification. Results are reported in the Appendix; I do not observe any meaningful changes in results relative to my baseline specification.

beginning in 2018. I find that the results from my counterfactual simulations do not change meaningfully from this additional sample restriction, so I retain the full U.S. sample in my preferred specification; these alternative estimates are provided in the Appendix.

In the specifications with consumer heterogeneity, estimated coefficients on price are slightly larger in magnitude than either the OLS or IV logit specifications. Estimated own-price elasticities by technology are presented in Table 6, for both the levels and logs specifications. I find that estimated elasticities for the specification with prices in logs are less dispersed across technologies than for the specification in levels. Estimated own-price elasticities are less than -1 for the specification in logs. Halogen, CFL, and LED own-price elasticities are less than -1 for the specification in levels, but incandescent elasticities are not. There may be several possible explanations for this result. For one, incandescent bulbs were a commodity product for which there were few lower-cost substitutes available; recall that the outside good in this model is purchasing bulbs at non-Nielsen retailers. Second, the definition of product and market used in this model does not capture all the store-to-store or week-to-week variation in incandescent prices over which consumers may have been more responsive. While I use the specification in levels as my primary set of demand estimates, I also present the core counterfactual results using the specification in logs in the Appendix; I find that the central conclusions from my counterfactual analysis do not depend on whether prices are specified in levels or in logs.

6 Counterfactual Simulations

To shed light on how the timing and design of policies affected substitution between the different efficient lighting technologies, I simulate market outcomes from several counterfactual policy regimes. To understand the impact of federal efficiency standards, I simulate the development of the general purpose lighting market if the standards were not implemented beginning in 2012, but instead delayed anywhere from one to six years. To explore the im-

pact of state and local subsidies, I impose a constant \$1 rebate on all CFLs in the market and then systematically vary the year in which these subsidies end, from prior to the study period at one extreme to 2018 at the other. I also examine how the simulated impact of these counterfactual policies would change given alternative lifetimes or externalities associated with each technology; given counterfactual timing of *both* standards and subsidies; and given alternative rates of LED price declines.

I consider two key environmental outcome metrics in evaluating these counterfactual policies: the total quantity of LEDs sold during the study period and the average discounted externality per hour of lighting sold during the study period.⁴⁷ While it is clear why a welfare-maximizing social planner would care about medium-term reductions in negative externalities, the total quantity of LEDs sold in the years immediately following commercialization may also matter for longer-term externality reduction potential. For example, the rate at which LED producers increase quantity sold may influence the rate of learning-by-doing, and expectations of near-term profits may influence incentives to invest in further product improvements.⁴⁸

6.1 Results

6.1.1 Preferred Timing of Standards

In the first set of counterfactual simulations, I ask how the impact of federal efficiency standards would change if their implementation were delayed to later in the study period. On

⁴⁷I focus on the externality associated with the emissions of greenhouse gases during electricity generation. Details of this externality calculation are provided in the Appendix. In choosing my preferred outcome measure around the environmental externality, I weighed two competing considerations. On the one hand, simply calculating the average wattage per bulb sold or the total lighting energy consumption during the study period would likely underestimate the importance of LEDs, which lived longer than their less efficient counterparts and whose benefits would extend beyond the study period. On the other hand, calculating the average wattage per hour of lighting sold or the total lighting energy consumption of bulbs sold during the study period (regardless of when incurred) would likely overestimate the importance of LEDs, whose use might continue years or decades into future. In balancing these considerations, I use the discounted average externality per hour of lighting as my preferred outcome metric. Alternative outcome measures are also possible.

⁴⁸It is well established that market size or time-to-profitability may influence the extent and direction of follow-on innovation. See, for example, Acemoglu et al. (2012) and Budish et al. (2015).

the one hand, later implementation would mean high consumption of inefficient incandescent bulbs in lieu of slightly more efficient halogens, as well as potentially lower consumption of CFLs. On the other hand, later implementation would change the distribution and number of consumers entering the market as LEDs became available. To conduct this set of counterfactual simulations, I assume that the set of incandescent products (and the limited set of halogen products) available in the year before the actual standards were implemented would remain fixed until the counterfactual implementation date. I then shift forward the set of incandescent and halogen offerings from the first year after the actual standards were implemented to whatever is the first year after the counterfactual standards, and likewise for the second year after implementation, third year, and so forth. This approach is grounded in the assumption that general purpose halogens entered the market to comply with federal efficiency standards, and they would not have entered the market until later if standards had been delayed.⁴⁹ I also assume that the arrival of LEDs in the market was exogenous to the timing of these efficiency standards, and therefore I hold fixed the set of LED products and their characteristics in each counterfactual simulation.^{50,51}

⁴⁹Incandescents could not technologically comply with the first-phase standards, so it is straightforward to phase them out of the market in line with whatever is the alternative implementation date.

⁵⁰The assumption that the arrival time of LEDs (and LED product characteristics over the short term) was exogenous to the timing of these policies is grounded in three key institutional details about the early LEDs market. First, the discovery that made possible the use of LEDs in general purpose lighting represented a fundamental scientific breakthrough, winning the Nobel Prize in Physics. Second, rapid reductions in the cost of LED chips during this period stemmed in large part from spillovers from other product markets: improvements in manufacturing LEDs for screens and improvements in semiconductor manufacturing generally. Third, Chinese industrial policy lowered barriers to entry for LED manufacturers, further encouraging the development of this industry. Note that these institutional details are highly specific to this technology setting over this period, and the assumption of exogenous product entry may not be appropriate in other contexts. My counterfactual simulations consider how policy timing affected the initial size of the LED market because of this potential for endogenous innovation responses in other technology settings or in the LED market over the longer term. Additionally, in ongoing work, I am endogenizing the prices of LEDs (and other bulb technologies) in response to alternative policy timing, taking into account the dual pressures of the threat of entry by new LED manufacturers and changes to market size and the distribution of consumer heterogeneity over time.

⁵¹In my main specification, I also hold fixed the set of CFL products in each counterfactual simulation. As a robustness check, I vary CFL product offerings in line with the counterfactual standards, such that the set of CFL products available in the market are held constant from 2011 (or from 2010 in California), until whatever date the counterfactual standards come into effect. Then I shift forward the post-standards product offerings accordingly. I find that the results of the counterfactual simulations do not change in an economically meaningful way from my baseline specification. These results are presented in the Appendix.

To build intuition for my findings, I first review predicted market shares by technology and wattage-equivalent for the counterfactual scenario where standards implemented are implemented beginning in 2016. These simulated market shares are presented in Figures 16a and 16b; simulated baseline results, where standards are implemented beginning in 2012, are presented in Figures 14a and 14b. In validating time trend specifications in my demand model, I required that the simulated counterfactual market shares meet certain criteria. For one, I required that the model predicts the phase-out of incandescent bulbs during whatever year counterfactual standards are assumed to be implemented. I also required that the simulated market shares of bulbs by wattage-equivalent follow a similar pattern as observed in the actual data, with a temporary increase in the market share of 60W-equivalent and 75W-equivalent bulbs prior to the standards implementation. As illustrated in Figure 16, both of these requirements are met in my preferred model specification.

Figures 17a and 17b then present the first of my two environmental outcome metrics: the total quantity of LEDs sold over the study period. The vertical axis specifies the year in which counterfactual standards were first implemented; the bars along the horizontal axis present the simulated quantity of LEDs sold (or the difference in the quantity of LEDs sold relative to the baseline of standards implemented beginning in 2012). I find that the total quantity of LEDs is maximized when standards are implemented at the end of the study period, in 2018. In fact, my model predicts that 35.6% fewer LEDs are sold during the study period when standards are implemented in 2012 relative to 2018.⁵² Implementing standards later increases the total market size in the years where LEDs are widely available in the market, yielding this larger total LED quantity.

Figures 18a and 18b present my second environmental metric: the average discounted environmental externality per hour of lighting sold during the study period. I find that the average environmental externality is minimized when standards are implemented beginning in 2012. However, the average externality is relatively flat across the counterfactual standards

⁵²When CFL product offerings are shifted alongside incandescents and halogens, I find that the total quantity of LEDs sold drops by 29.0%.

scenarios; implementing beginning in 2012 reduces the average discounted externality by 1.8% relative to 2015 and by 2.5% relative to 2018.⁵³

Figure 15 highlights the difference in predicted outcomes for standards implemented beginning in 2012 (i.e., the simulated baseline) using counterfactual simulations that account for the endogenous evolution of market size and consumer heterogeneity (the “full” counterfactual) versus simulations that hold these market characteristics fixed (the “naive” counterfactual). As is evident in the figure, failing to account for the dynamic effects of changing market size and distribution would vastly overstate the benefits of the lighting efficiency standards, overestimating total LEDs quantity by 110% and underestimating the average environmental externality by 19%.

6.1.2 Preferred Timing of Subsidies

In the second set of counterfactual analyses, I impose a constant \$1 per bulb rebate on all CFLs, until some phase-out date. This analysis is complicated by the fact that I do not observe which specific CFL (or LED) products are already discounted in the data, only the price paid by the consumer. Therefore, to predict which prices have already been discounted through manufacturer buydowns, I calculate the maximum price at which each CFL UPC was sold at each store in a given year, and then identify the store-UPC combinations where the product was discounted by \$1 or more per bulb for at least four consecutive weeks.⁵⁴ The goal is to identify bulbs with long-term discounts rather than short-term sales, which might last for only one or two weeks. Next, after identifying this set of bulbs for which I predict buydown rebates are in place, I replace the “rebated” price with the “non-rebated” price, defined to be the maximum price at which that UPC is offered in that store in that year. Finally, I then apply a constant \$1 discount (per bulb) across all CFL products in the dataset, including in counties that did not have active rebate programs in place. For

⁵³When CFLs product offerings are shifted alongside incandescents and halogens, I find that implementing beginning in 2012 reduces the average externality by 2.4% relative to 2015 and by 3.3% relative to 2018.

⁵⁴I conduct this exercise only in counties with active manufacturer buydown programs in place, as documented in the EPA’s annual reports on lighting rebate programs.

this reason, the results of the counterfactual simulations around alternative rebate programs should be interpreted as varying a constant \$1 per bulb discount for CFLs, not as varying the actual patchwork of rebate programs in place across the country.⁵⁵ As an illustration, Figure 19 presents counterfactual market shares when CFL rebates are terminated after 2015. The sharp decline in CFL sales after rebates are phased out is immediately apparent.⁵⁶

Figures 20a and 20b illustrate simulated quantities of LEDs sold under alternative timings of CFL rebates. For this policy, the market size for LEDs is maximized when CFL rebates are never active during the study period – i.e., are terminated in 2009. This result stems not only from direct competition between LEDs and CFLs in later years of the study period, but also from changes to overall market size and to the distribution of consumer preferences for efficient bulbs from early increases in CFL adoption. Phasing out CFL rebates after 2012 results in 4.0% fewer LEDs sold relative to ending rebates in 2009, and 20.4% more LEDs sold relative to ending rebates in 2018.

Figures 21a and 21b then show the impact of rebate timing on the average discounted externality per hour of lighting purchased during the study period. In simulations with my preferred model specification, the average environmental externality is minimized when CFL rebates are phased out after 2014. Ending CFL rebates in 2012 results in 3.0% lower average externality relative to phasing out after 2009, or 2.3% lower externality relative to ending in 2018.

As mentioned above, these counterfactuals vary the timing of a simulated \$1 rebate on all CFL bulbs, not the actual set of rebate policies in place across states and utility territories, and should be interpreted as such. Nevertheless, it is useful to note that many of the actual rebate programs did not discontinue their CFL subsidies until later in the study period. Based on EPA data on lighting rebate programs, 80 of 101 unique programs offering subsidies for general purpose bulbs to residential consumers in 2015 included some form of

⁵⁵It is also possible that I fail to identify certain pre-existing rebates less than \$1 per bulb. In that case, the counterfactual simulations will reflect a rebate greater than \$1 per bulb on some subset of CFL bulbs.

⁵⁶In reviewing these counterfactual simulations, I again validate counterfactual market shares against the criteria described above. My preferred model specification again fulfills both criteria.

CFL subsidies. This number fell to 64 out of 98 programs in 2016, 30 of 104 programs in 2017, and 0 of 102 programs in 2018.⁵⁷

6.1.3 Comparing Policy Instruments

Given these findings on the impact of standards and rebates timing, we can also compare the two different policy instruments.⁵⁸ Under both policies, the total quantity of LEDs sold during the study period is maximized when early policy intervention is minimized – either by delaying standards implementation or by phasing out CFL rebates as early as possible. However, the trajectory of the average environmental externality varies under different policy instruments. For standards, the immediate environmental benefit of early implementation outweighs any reductions in efficiency later in the study period. For rebates, some early intervention is beneficial, but the average environmental externality is minimized when rebates do not remain beyond the first half of the study period.

In this technology setting, the policy timing that minimizes the medium-term average externality differs significantly from that which maximizes LED market size. This is an empirical finding, and these two metrics may be more closely aligned in other technology settings. For this reason, I also compare these two policy instruments in terms of the relative magnitude of the tradeoff between increasing market size of LEDs and reducing the externality associated with bulbs purchased in the study period. Figures 22a and 22b plot the percentage change in the average externality and total quantity of LEDs purchased during the study period, relative to the baseline of standards implemented beginning in 2012 (Figure 22a) or CFL subsidies phased out from 2012 (Figure 22b). The surface created by these counterfactual policy changes represents the frontier facing the policymaker trading off benefits from near-term externality reductions with longer-term potential innovation bene-

⁵⁷Here I define a “unique program” as having a unique program administrator; many of these programs covered a large number of counties. I also do not consider programs offering bulbs to other types of customers, such as commercial entities or K-12 institutions, or subsidies targeted at specialty bulbs, since these fall outside my study purview.

⁵⁸Additional counterfactual simulations also consider the impact of varying the timing of standards implementation or CFL rebate phase-out in tandem. These results are presented in the Appendix.

fits. Since the regulator would generally prefer to increase LED production and decrease the average externality, we expect that welfare increases as we move outwards from the bottom right quadrant. We see immediately, therefore, that there is no “free lunch” for the regulator considering deviations from 2012 policy implementation, for either standards or rebates, as there are no points in this quadrant. The three points in the upper left quadrant of Figure 22b suggest, however, that there are some alternative policies that are strictly dominated by the 2012 baseline. Likewise, we see in Figure 22a that 2015 standards implementation approximately dominates 2013 standards implementation; in both cases, the dominant policy achieves a higher total quantity of LEDs for about the same level of average externality. For other policy comparisons in the upper right quadrant or lower left quadrant, the regulator faces a tradeoff between greater (or lesser) quantity of LEDs and increased (or reduced) medium-term externality. The outcome of this tradeoff ultimately depends on how the regulator weights these outcomes; identifying the efficient weights for the lighting context falls outside the scope of this paper.

6.2 Endogenous Innovation and Regulator Beliefs

My baseline model does not endogenize innovation responses directly and assumes perfect certainty about demand parameters and the evolution of products within each technology category. Relaxing both of these simplifications may be important for understanding the impact of policy timing in other technology contexts. To begin generalizing my model, I pose additional counterfactual questions. First, given two alternative policy timings, what would be the counterfactual rate of LED price declines needed to achieve the same total quantity of LEDs in one policy scenario as predicted in another? Second, what is the range of counterfactual LED price declines for which my earlier findings about the average externality continue to hold? The first of these questions is intended to shed light on how alternative rates of industry learning might influence findings around the total quantity of LEDs sold. The second question is designed to examine how sensitive the conclusions about

average externality are to the rate of LED price declines, around which the regulator may have considerable ex ante uncertainty.

First, let us consider the alternative rates of LED price declines that produce similar changes in the market size of LEDs as what I observe in my policy simulations. I find that delaying standards from 2012 to 2018 is equivalent to LED prices declining by 0.5 percentage point faster per quarter in terms of the impact on the total quantity of LEDs sold (Figure 23a). By the end of the study period, this differential rate of price decline translates to LED prices that are approximately 20% higher relative to what I observe in the data. Likewise, delaying the phase-out of CFL rebates from 2012 to 2018 is equivalent to a comparable change in the rate of LED price declines, but in the opposite direction (i.e., LED prices decline by approximately 0.5 percentage point more slowly per quarter) (Figure 23b).

Next, let us consider how robust are the average externality findings to the rate of LED price declines. For efficiency standards, I find that implementing standards in 2012 minimizes the average externality for a wide range of alternative LED price declines.⁵⁹ By contrast, I find that the impact of CFL rebate phase-out on the average externality is far more sensitive to alternative rates of LED price decline. Figure 24b suggests that ending CFL rebates earlier is beneficial given higher rates of LED price decline. With slower rates of LED price decline, however, retaining CFL rebates for longer is beneficial. Intuitively, the environmental benefits from immediate adoption of the early efficient technology are less likely to outweigh the reduced market size available to the later efficient technology when the latter is able to command a large market share more quickly. As the delay between technology generations increases, earlier intervention is more likely to be justified from the perspective of reducing the overall environmental externality.

⁵⁹As seen in Figure 24a, the predicted average externality plotted against the year of standards implementation is actually flattest around the rates of LED price decline observed in the data (approximately 6.5% per quarter). At both faster and slower rates of LED price decline, the average externality curve is steeper, suggesting greater advantage to implementing standards early.

6.3 Alternative Lifetimes and Externalities

Lastly, to understand how results depend on key parameters, I examine the predicted effects of subsidies and standards under alternative technology lifetimes or externality levels. In particular, I examine how predicted results would differ if CFLs had the same lifetime as halogens, or vice versa, or if CFLs had the same wattage associated with a given light output as halogens, or vice versa. I also examine the extent to which shorter than anticipated CFL lifetimes might have affected outcomes.⁶⁰ I hold all other estimated parameters fixed in the demand system.⁶¹

Results from this exercise are presented in Figure 25 for standards and Figure 26 for rebates. Decreasing the lifetime of CFLs increases the total quantity of LEDs, while increasing the lifetime of halogens decreases the total quantity. The latter effect is smaller in magnitude, as halogens are more distant substitutes for LEDs, and declines to approximately zero as standards are implemented later and halogens become widely available later. By construction, changing the externality values of CFLs or halogens has no impact on the total quantity of LEDs sold.

Impacts on the average environmental externality are more subtle. In the case of standards, the average environmental externality is strongly minimized with early standards when halogens produce much lower externality; in this case, the immediate environmental benefits of replacing incandescent bulbs with halogens are even larger and outweigh any considerations around the later introduction of LEDs. By contrast, the average environmental

⁶⁰As noted throughout this paper, a typical rated lifetime for a general purpose CFL bulb was 8,000 hours. However, frequent on-off switching could diminish the actual lifetime of CFL bulbs. In all of my estimation and counterfactual results presented thus far, I use adjusted CFL lifetimes that account for typical on-off cycles, based on information presented in DOE (2016). In this final set of counterfactual simulations, however, I also ask how results would differ if CFL failure rates followed the same overall distribution as the other technologies, with no adjustments for on-off switching. This change has the effect of increasing average CFL lifetimes from 4,560 hours (or 57% of 8,000 hours) to the full 8,000-hour rated life.

⁶¹Of course, the parameters governing the utility that a consumer derives from purchasing a given bulb might be expected to change as these product characteristics change. However, the product category fixed effects in my demand model do not allow me to disentangle the share of utility from a given product derived from its lifetime, its efficiency, or other characteristics. This exercise should therefore be interpreted as mechanically changing the lifetime and externality parameters that produce certain counterfactual outcomes from a given pattern of consumer demand, while holding fixed the determinants of that consumer demand.

externality is strongly minimized with late standards when halogens are longer lived, as the reduction in the later market size is larger without a concomitant increase in externality benefits. Additional impacts from varying CFL parameters are presented in Figure 25.

For rebates, most of these alternative simulations predict that the average externality is minimized when CFL rebates are phased out sometime in the middle of the study period, between 2012 and 2015. The externality-minimizing year is shifted later when halogens are longer lived; in this scenario, there is no longer a tradeoff between shorter lived but less efficient halogens and longer lived but more efficient CFLs, so there are greater relative benefits to encouraging CFL adoption. By contrast, the externality-minimizing year is shifted earlier when CFLs have longer lifetimes, which increases their impact on later LED adoption. This year is also shifted earlier when halogens produce fewer externalities, which means there are fewer relative benefits from incentivizing CFLs once the efficiency standards are active and halogens are widely available. Lastly, when CFLs have higher externality values, the model suggest that rebates should be phased out as soon as possible, as their impact on the LED market outweighs even the initial externality benefits relative to incandescents.

7 Conclusion

This paper explores how the design and timing of policies – efficiency standards and subsidies for efficient products – affects competition between early and later generations of a clean technology. In modeling consumer demand for general purpose lighting, I find that the overall externality per hour of lighting is minimized when standards are implemented in 2012 and when subsidies for CFLs are phased out after 2014. This finding trades off immediate externality benefits from switching to early efficient technologies (CFLs and halogens) with reductions in LED market size later in the study period, when LEDs would result in even greater externality reduction per hour of lighting. I find that early intervention was warranted for reducing the average medium-term externality in the case of standards, which

primarily targeted halogens. Some early intervention was warranted for subsidies, which primarily targeted CFLs, though terminating the subsidies several periods before LEDs were widely adopted was optimal from an externality reduction perspective. This empirical result depends fundamentally on the relative externalities and lifetimes across technologies in the efficient lighting market, as I show in supplemental counterfactual simulations.

Of course, reductions in the initial LED market size might have consequences not only for the direct impact on the average medium-term externality, but also for endogenous innovation in LED technology. Insofar as further improvements in LEDs are sensitive to rates of learning-by-doing or endogenous investment in technological improvements, policymakers may be concerned with increasing deployment of this later-generation technology because of these *indirect* impacts on the overall externality. Over the long run, the efficient second-best policy design may require both incentives for early technology deployment and support for later-generation research and development; this additional technology support must then be included in the overall cost of the policy. To that end, I show that faster LED price reductions would have been needed to achieve the same initial LED market size relative to a counterfactual world with delayed standards or with CFL subsidies retained through the duration of the study period.

A fruitful direction for future research is investigating this supply-side innovation response more thoroughly. Here the specific technology characteristics will again have first-order importance, with the extent of spillovers between early and later products depending on their technological similarities (or lack thereof). The extent of cross-product spillovers will help to determine both the magnitude and the direction of early policy impacts on later innovation. When policymakers are constrained to use second-best policies, understanding inefficiencies not only in existing technology adoption but also in the development and deployment of technologies over time is essential for achieving long-run climate mitigation goals as efficiently as possible.

References

- D. Acemoglu. Directed technical change. *The Review of Economic Studies*, 69(4):781–809, 2002. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/1556722>.
- D. Acemoglu, P. Aghion, L. Bursztyn, and D. Hemous. The environment and directed technical change. *American Economic Review*, 102(1):131–66, February 2012. doi: 10.1257/aer.102.1.131. URL <https://www.aeaweb.org/articles?id=10.1257/aer.102.1.131>.
- D. Acemoglu, P. Aghion, L. Barrage, and D. Hémous. Climate Change, Directed Innovation, and Energy Transition: The Long-run Consequences of the Shale Gas Revolution. Working paper, February 2019. URL https://scholar.harvard.edu/files/aghion/files/climate_change_directed_innovation.pdf.
- P. Aghion, P. A. David, and D. Foray. Science, technology and innovation for economic growth: Linking policy research and practice in ‘stig systems’. *Research Policy*, 38(4): 681 – 693, 2009. ISSN 0048-7333. doi: <https://doi.org/10.1016/j.respol.2009.01.016>. URL <http://www.sciencedirect.com/science/article/pii/S0048733309000213>. Special Issue: Emerging Challenges for Science, Technology and Innovation Policy Research: A Reflexive Overview.
- H. Allcott and M. Greenstone. Is there an energy efficiency gap? *Journal of Economic Perspectives*, 26(1):3–28, February 2012. doi: 10.1257/jep.26.1.3. URL <https://www.aeaweb.org/articles?id=10.1257/jep.26.1.3>.
- H. Allcott and M. Greenstone. Measuring the welfare effects of residential energy efficiency programs. Working Paper 23386, National Bureau of Economic Research, May 2017. URL <http://www.nber.org/papers/w23386>.
- H. Allcott and C. Knittel. Are consumers poorly informed about fuel economy? evidence from two experiments. *American Economic Journal: Economic Policy*, 11(1):1–37, Febru-

- ary 2019. doi: 10.1257/pol.20170019. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20170019>.
- H. Allcott and D. Taubinsky. Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–38, August 2015. doi: 10.1257/aer.20131564. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20131564>.
- J. Asker. Diagnosing foreclosure due to exclusive dealing. *The Journal of Industrial Economics*, 64(3):375–410, 2016. doi: <https://doi.org/10.1111/joie.12114>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/joie.12114>.
- C. L. Benkard. A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies*, 71(3):581–611, 2004. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/3700737>.
- S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2171802>.
- B. Bollinger and K. Gillingham. Learning-by-Doing in Solar Photovoltaic Installations. Working paper, 2019. URL <http://dx.doi.org/10.2139/ssrn.2342406>.
- S. Borenstein and J. B. Bushnell. Do two electricity pricing wrongs make a right? cost recovery, externalities, and efficiency. Working Paper 24756, National Bureau of Economic Research, June 2018. URL <http://www.nber.org/papers/w24756>.
- J. R. Brown. *Tax Policy and the Economy, Volume 30*. University of Chicago Press, 2016. URL <http://www.nber.org/books/brow-14>.
- E. Budish, B. N. Roin, and H. Williams. Do firms underinvest in long-term research? evidence from cancer clinical trials. *American Economic Review*, 105(7):2044–85, July 2015.

- doi: 10.1257/aer.20131176. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20131176>.
- G. Casey. Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation. Working paper, September 2019. URL <https://unbound.williams.edu/facultypublications/islandora/object/economicsworkingpapers:18>.
- C. Conlon and J. Gortmaker. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161, 2020. doi: <https://doi.org/10.1111/1756-2171.12352>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12352>.
- T. Covert and R. Sweeney. Winds of change: Technical progress and learning in wind power. 2020.
- L. W. Davis. Evidence of a decline in electricity use by U.S. households. *Economics Bulletin*, 37(2):1098–1105, 2017. URL <https://ideas.repec.org/a/ebl/ecbull/eb-17-00365.html>.
- L. W. Davis and G. E. Metcalf. Does better information lead to better choices? evidence from energy-efficiency labels. *Journal of the Association of Environmental and Resource Economists*, 3(3):589–625, 2016. doi: 10.1086/686252. URL <https://doi.org/10.1086/686252>.
- J. A. Dubin and D. L. McFadden. An econometric analysis of residential electric appliance holdings and consumption. *Econometrica*, 52(2):345–362, 1984. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1911493>.
- P. Dubois and C. Bonnet. Inference on vertical contracts between manufacturers and retailers allowing for non linear pricing and resale price maintenance,. *RAND Journal of Economics*, 41:139–164, 03 2010. doi: 10.1111/j.1756-2171.2009.00093.x.

- A. Gandhi and J.-F. Houde. Measuring Substitution Patterns in Differentiated-Products Industries. NBER Working Papers 26375, National Bureau of Economic Research, Inc, Oct. 2019. URL <https://ideas.repec.org/p/nbr/nberwo/26375.html>.
- T. Gerarden. Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs. Working paper, March 2018. URL https://scholar.harvard.edu/files/gerarden/files/gerarden_jmp.pdf.
- G. Gowrisankaran and M. Rysman. Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy*, 120(6):1173–1219, Dec. 2012. ISSN 0022-3808, 1537-534X. doi: 10.1086/669540.
- I. Hendel and A. Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006a. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/4123086>.
- I. Hendel and A. Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006b. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/4123086>.
- I. Hendel and A. Nevo. Intertemporal price discrimination in storable goods markets. *The American Economic Review*, 103(7):2722–2751, 2013. ISSN 00028282. URL <http://www.jstor.org/stable/42920669>.
- N. Hendren and B. Sprung-Keyser. A Unified Welfare Analysis of Government Policies*. *The Quarterly Journal of Economics*, 135(3):1209–1318, 03 2020. ISSN 0033-5533. doi: 10.1093/qje/qjaa006. URL <https://doi.org/10.1093/qje/qjaa006>.
- I. M. Hoffman, C. A. Goldman, G. Rybka, G. Leventis, L. Schwartz, A. H. Sanstad, and S. Schiller. Estimating the cost of saving electricity through u.s. utility customer-funded energy efficiency programs. *Energy Policy*, 104:1 – 12, 2017. ISSN 0301-4215. doi: <https://doi.org/10.1016/j.enpol.2017.04.011>.

[//doi.org/10.1016/j.enpol.2016.12.044](http://doi.org/10.1016/j.enpol.2016.12.044). URL <http://www.sciencedirect.com/science/article/pii/S0301421516307042>.

S. Houde and J. E. Aldy. Consumers' response to state energy efficient appliance rebate programs. *American Economic Journal: Economic Policy*, 9(4):227–55, November 2017. doi: 10.1257/pol.20140383. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20140383>.

P. D. Howe, M. Mildenberger, J. R. Marlon, , and A. Leiserowitz. Geographic variation in opinions on climate change at state and local scales in the usa. *Nature Climate Change*, 5:596–603, 2015. doi: doi:10.1038/nclimate2583.

Interagency Working Group on the Social Cost of Carbon, U.S. Government. Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis, Under Executive Order 12866. Report, July 2015. URL <https://web.archive.org/web/20150917145701/https://www.whitehouse.gov/sites/default/files/omb/inforeg/scc-tsd-final-july-2015.pdf>.

K. Ito. Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–63, February 2014. doi: 10.1257/aer.104.2.537. URL <https://www.aeaweb.org/articles?id=10.1257/aer.104.2.537>.

A. Jaffe, R. Newell, and R. Stavins. A tale of two market failures: Technology and environmental policy. *Ecological Economics*, 54(2-3):164–174, 2005. URL <https://EconPapers.repec.org/RePEc:eee:ecolec:v:54:y:2005:i:2-3:p:164-174>.

N. Johnstone, I. Haščič, and D. Popp. Renewable Energy Policies and Technological Innovation: Evidence Based on Patent Counts. *Environmental and Resource Economics*, 45: 133–155, 2010. URL <https://doi.org/10.1007/s10640-009-9309-1>.

A. Langer and D. Lemoine. Designing dynamic subsidies to spur adoption of new technolo-

- gies. Working Paper 24310, National Bureau of Economic Research, February 2018. URL <http://www.nber.org/papers/w24310>.
- R. S. Lee. Vertical integration and exclusivity in platform and two-sided markets. *The American Economic Review*, 103(7):2960–3000, 2013. ISSN 00028282. URL <http://www.jstor.org/stable/42920677>.
- D. Leip. David Leip’s atlas of U.S. Presidential elections, datasets, 2017. URL <https://doi.org/10.7910/DVN/XX3YJ4>.
- S. D. Levitt, J. A. List, and C. Syverson. Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy*, 121(4):643–681, 2013. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/10.1086/671137>.
- National Electrical Manufacturers Association. NEMA Comments on Energy Conservation Standards for General Service Lamps Notice of Proposed Rulemaking. Technical report, May 2019.
- Navigant Consulting. 2010 u.s. lighting market characterization. Report, U.S. Department of Energy, January 2012. URL <https://www1.eere.energy.gov/buildings/publications/pdfs/ssl/2010-lmc-final-jan-2012.pdf>.
- Navigant Consulting. 2015 u.s. lighting market characterization. Report, U.S. Department of Energy, November 2017. URL https://www.energy.gov/sites/prod/files/2017/12/f46/lmc2015_nov17.pdf.
- A. Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342, 2001. doi: <https://doi.org/10.1111/1468-0262.00194>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00194>.

- A. Nevo. New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand Systems. *The Review of Economics and Statistics*, 85(2):266–275, May 2003. URL <https://ideas.repec.org/a/tpr/restat/v85y2003i2p266-275.html>.
- R. G. Newell and J. Siikamäki. Individual time preferences and energy efficiency. *American Economic Review*, 105(5):196–200, May 2015. doi: 10.1257/aer.p20151010. URL <https://www.aeaweb.org/articles?id=10.1257/aer.p20151010>.
- W. D. Nordhaus. Do real-output and real-wage measures capture reality? the history of lighting suggests not. In T. F. Bresnahan and R. J. Gordon, editors, *The Economics of New Goods*, pages 27–70. University of Chicago Press, Chicago, 1996.
- V. Sivaram. Unlocking clean energy. *Issues in Science and Technology*, 33(2):31–40, Winter 2017.
- J. C. Stein. Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *Quarterly Journal of Economics*, 104(Dec.):655–669, 1989.
- U.S. Department of Energy. Notice of Proposed Rulemaking Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment: General Service Lamps. Working paper, February 2016.
- U.S. Environmental Protection Agency. Energy Star Summary of Lighting Programs. Report, 2010-2018.
- U.S. Environmental Protection Agency. eGRID2014 Summary Tables. Report, January 2017. URL https://19january2017snapshot.epa.gov/sites/production/files/2017-01/documents/egrid_summarytables.pdf.
- S. B. Villas-Boas. Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies*, 74(2):625–652, 2007. ISSN 00346527, 1467937X. URL <http://www.jstor.org/stable/4626153>.

A Theory: Derivation of Optimal Policy and Efficient Second-Best Policies

A.1 Optimal Policy (Proof of Proposition 1)

First, note that the expression for total welfare is given by:

$$\begin{aligned}
W = & \underbrace{s_{d,d}(1+\lambda)(-p_d) + (1+\lambda) \int_{s_{d,d}} \nu_{id} f(\nu_d) d\nu_d}_{\text{(Private) Surplus from } (D, D)} \\
& + \underbrace{s_{c_1,c_1}(1+\lambda)(-p_{c_1}) + (1+\lambda) \int_{s_{c_1,c_1}} \nu_{c_1i} f(\nu_{c_1}) d\nu_{c_1}}_{\text{(Private) Surplus from } (C_1, C_1)} \\
& + \underbrace{s_{d,c_2}(-p_d - \lambda p_{c_2}) + \int_{s_{d,c_2}} \nu_{id} f(\nu_d) d\nu_d + \lambda \int_{s_{d,c_2}} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2}}_{\text{(Private) Surplus from } (D, C_2)} \\
& + \underbrace{s_{0,c_2}(-\lambda p_{c_2}) + \lambda \int_{s_{0,c_2}} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2}}_{\text{(Private) Surplus from } (0, C_2)} \\
& + \underbrace{s_{d,d}(1+\lambda)(-\zeta_d) + s_{c_1,c_1}(1+\lambda)(-\zeta_{c_1}) + s_{d,c_2}(-\zeta_d - \lambda \zeta_{c_2}) + s_{0,c_2}(-\lambda \zeta_{c_2})}_{\text{Externality}}
\end{aligned}$$

In the absence of policy intervention, recall that individual optimization decisions can be expressed as a series of thresholds for ν_{id} , ν_{ic_1} , and ν_{ic_2} that govern whether consumer i prefers certain consumption bundles over others. That is, consumers will prefer (D, D) to $(0, 0)$ whenever $\nu_{id} \geq \bar{\nu}_d = p_d$; (C_1, C_1) to $(0, 0)$ whenever $\nu_{ic_1} \geq \bar{\nu}_{c_1} = p_{c_1}$; and $\{0, C_1\}$ to $(0, 0)$ whenever $\nu_{ic_2} \geq \bar{\nu}_{c_2} = p_{c_2}$; all other preference relations can be expressed in terms of $\bar{\nu}_d$, $\bar{\nu}_{c_1}$, and $\bar{\nu}_{c_2}$. Therefore, we can also model the social planner as choosing a (possibly constrained) price schedule that yields certain values for $\bar{\nu}_d$, $\bar{\nu}_{c_1}$, and $\bar{\nu}_{c_2}$, which defines a certain allocation across consumers.

Therefore, if the social planner is able to impose an unconstrained price schedule, we

have the social planner's problem given by:

$$\begin{aligned}
\max_{t_d, t_{c_1}, t_{c_2}} [& \max_{\bar{\nu}_d, \bar{\nu}_{c_1}, \bar{\nu}_{c_2}} \{ (1 + \lambda) s_{d,d}(\cdot) (-p_d - t_d) + (1 + \lambda) \int_{s_{d,d}(\cdot)} \nu_{id} f(\nu_d) d\nu_d \\
& + (1 + \lambda) s_{c_1,c_1}(\cdot) (-p_{c_1} - t_{c_1}) + (1 + \lambda) \int_{s_{c_1,c_1}(\cdot)} \nu_{ic_1} f(\nu_{c_1}) d\nu_{c_1} \\
& + s_{d,c_2}(\cdot) (-p_d - t_d - \lambda p_{c_2} - \lambda t_{c_2}) + \int_{s_{d,c_2}(\cdot)} \nu_{id} f(\nu_d) d\nu_d + \lambda \int_{s_{d,c_2}(\cdot)} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2} \\
& + s_{0,c_2}(\cdot) (-\lambda p_{c_2} - \lambda t_{c_2}) + \lambda \int_{s_{0,c_2}(\cdot)} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2} \} \\
& - (1 + \lambda) s_{d,d}(\cdot) \zeta_d - (1 + \lambda) s_{c_1,c_1}(\cdot) \zeta_{c_1} - s_{d,c_2}(\cdot) (\zeta_d + \lambda \zeta_{c_2}) - s_{0,c_2}(\cdot) (\lambda \zeta_{c_2}) \\
& + (1 + \lambda) s_{d,d}(\cdot) t_d + (1 + \lambda) s_{c_1,c_1}(\cdot) t_{c_1} + s_{d,c_2}(t_d + \lambda t_{c_2}) + s_{0,c_2}(\lambda t_{c_2})]
\end{aligned}$$

Here each $s_{j,j'}$ is a function of the three thresholds for adoption, which in turn are functions of the effective tax rates chosen by the regulator: $s_{j,j'}(\bar{\nu}_d(t_d, t_{c_1}, t_{c_2}), \bar{\nu}_{c_1}(t_d, t_{c_1}, t_{c_2}), \bar{\nu}_{c_2}(t_d, t_{c_1}, t_{c_2}))$; this notation is suppressed in the expression above for simplicity.

As an illustration, I write out the regulator's full first-order condition with respect to t_{c_1} :

$$\begin{aligned}
\frac{\partial W}{\partial t_{c_1}} = & (1 + \lambda) s_{c_1,c_1}(\cdot) (-1) + (1 + \lambda) s_{c_1,c_1}(\cdot) \\
& + [(1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_1}} (-p_d - t_d) + (1 + \lambda) \frac{\partial}{\partial t_{c_1}} (\int_{s_{d,d}(\cdot)} \nu_{id} f(\nu_d) d\nu_d) \\
& + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} (-p_{c_1} - t_{c_1}) + (1 + \lambda) \frac{\partial}{\partial t_{c_1}} (\int_{s_{c_1,c_1}(\cdot)} \nu_{ic_1} f(\nu_{c_1}) d\nu_{c_1}) \\
& + \frac{\partial s_{d,c_2}}{\partial t_{c_1}} (-p_d - t_d - \lambda p_{c_2} - \lambda t_{c_2}) + \frac{\partial}{\partial t_{c_1}} (\int_{s_{d,c_2}(\cdot)} \nu_{id} f(\nu_d) d\nu_d) + \lambda \frac{\partial}{\partial t_{c_1}} (\int_{s_{d,c_2}(\cdot)} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2}) \\
& + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (-\lambda p_{c_2} - \lambda t_{c_2}) + \lambda \frac{\partial}{\partial t_{c_1}} (\int_{s_{0,c_2}(\cdot)} \nu_{ic_2} f(\nu_{c_2}) d\nu_{c_2})] \\
& + (1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_1}} (t_d - \zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} (t_{c_1} - \zeta_{c_1}) \\
& + \frac{\partial s_{d,c_2}}{\partial t_{c_1}} (t_d - \zeta_d + \lambda(t_{c_2} - \zeta_{c_2})) + \lambda \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (t_{c_2} - \zeta_{c_2}) = 0
\end{aligned}$$

The first and second terms cancel with each other. The subsequent set of terms in square brackets cancel after we substitute the first-order condition from individual optimization.

The third set of terms remains, and the regulator's first-order condition can be rewritten as:

$$\begin{aligned}\frac{\partial W}{\partial t_{c_1}} &= (1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_1}} (t_d - \zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial s_{d,c_2}}{\partial t_{c_1}} (t_d - \zeta_d + \lambda(t_{c_2} - \zeta_{c_2})) + \lambda \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (t_{c_2} - \zeta_{c_2}) = 0\end{aligned}$$

Likewise, the regulator's FOCs with respect to t_d and t_{c_2} are given as follows:

$$\begin{aligned}\frac{\partial W}{\partial t_d} &= (1 + \lambda) \frac{\partial s_{d,d}}{\partial t_d} (t_d - \zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_d} (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial s_{d,c_2}}{\partial t_d} (t_d - \zeta_d + \lambda(t_{c_2} - \zeta_{c_2})) + \lambda \frac{\partial s_{0,c_2}}{\partial t_d} (t_{c_2} - \zeta_{c_2}) = 0\end{aligned}$$

$$\begin{aligned}\frac{\partial W}{\partial t_{c_2}} &= (1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_2}} (t_d - \zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_2}} (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial s_{d,c_2}}{\partial t_{c_2}} (t_d - \zeta_d + \lambda(t_{c_2} - \zeta_{c_2})) + \lambda \frac{\partial s_{0,c_2}}{\partial t_{c_2}} (t_{c_2} - \zeta_{c_2}) = 0\end{aligned}$$

It is immediately apparent that these optimality conditions will be satisfied by setting $t_{c_1}^* = \zeta_{c_1}$, $t_d^* = \zeta_d$, and $t_{c_2}^* = \zeta_{c_2}$.

A.2 Efficient Second-Best Policies

A.2.1 One Clean Technology (Proof of Proposition 2)

Now we turn to deriving the efficient policy when the regulator may implement only a single technology-specific policy on C_1 . First, we consider the case where C_1 is the only clean technology, and the regulatory does not need to account for competition across generations of the clean technology. Under these conditions, the regulator's FOC for the efficient pricing policy on C_1 is given by:

$$\frac{\partial W}{\partial t_{c_1}} = (1 + \lambda)(t_{c_1}) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} + (1 + \lambda)(-\zeta_d) \frac{\partial s_{d,d}}{\partial t_{c_1}} + (1 + \lambda)(-\zeta_{c_1}) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} = 0$$

Rearranging terms to solve for \hat{t}_{c_1} yields:

$$\hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}})[\zeta_d \frac{\partial s_{d,d}}{\partial t_{c_1}} + \zeta_{c_1} \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}]$$

Adding and subtracting $\zeta_{c_1} \frac{\partial s_{d,d}}{\partial t_{c_1}}$ allows us to rewrite this expression in terms of the C_1/D margin and the C_1 /no adoption margin, as presented in the main body of the paper:

$$\hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}})[\underbrace{(\zeta_d - \zeta_{c_1}) \frac{\partial s_{d,d}}{\partial t_{c_1}}}_{\text{Impact on } C_1/D \text{ margin}} + \underbrace{\zeta_{c_1}(\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} + \frac{\partial s_{d,d}}{\partial t_{c_1}})}_{\text{Impact on } C_1/\text{no adoption margin}}]$$

A.2.2 Multiple Clean Technology Generations (Proof of Proposition 3)

Next, we relax the initial assumption that there is only one clean technology and show how the regulator's choice of the efficient pricing policy on C_1 changes when competition across clean technology generations is also a consideration. In this case, the regulator's FOC for the efficient pricing policy on C_1 is given by:

$$\frac{\partial W}{\partial t_{c_1}} = (1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_1}}(-\zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}(t_{c_1} - \zeta_{c_1}) + \frac{\partial s_{d,c_2}}{\partial t_{c_1}}(-\zeta_d - \lambda\zeta_{c_2}) + \lambda \frac{\partial s_{0,c_2}}{\partial t_{c_1}}(-\zeta_{c_2}) = 0$$

Rearranging terms to solve for the modified \hat{t}_{c_1} yields:

$$\hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}})[\frac{\partial s_{d,d}}{\partial t_{c_1}}(\zeta_d) + \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}(\zeta_{c_1}) + \frac{\partial s_{d,c_2}}{\partial t_{c_1}}(\frac{\zeta_d + \lambda\zeta_{c_2}}{1 + \lambda}) + \frac{\partial s_{0,c_2}}{\partial t_{c_1}}(\frac{\lambda\zeta_{c_2}}{1 + \lambda})]$$

Finally, adding and subtracting $\zeta_{c_1} \frac{\partial s_{d,d}}{\partial t_{c_1}}$, $\zeta_{c_1} \frac{\partial s_{d,c_2}}{\partial t_{c_1}}$, and $\zeta_{c_1} \frac{\partial s_{0,c_2}}{\partial t_{c_1}}$ again allows us to rewrite this expression in terms of the margins of decision-making, as presented in the main body of the

paper:

$$\begin{aligned} \hat{t}_{c_1} = (1/\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) [& \underbrace{(\zeta_d - \zeta_{c_1}) \frac{\partial s_{d,d}}{\partial t_{c_1}}}_{\text{Impact on } C_1/D \text{ margin}} + \underbrace{\zeta_{c_1} (\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} + \frac{\partial s_{d,d}}{\partial t_{c_1}} + \frac{\partial s_{d,c_2}}{\partial t_{c_1}} + \frac{\partial s_{0,c_2}}{\partial t_{c_1}})}_{\text{Impact on } C_1/\text{no adoption margin}} \\ & + \underbrace{(\frac{\zeta_d + \lambda \zeta_{c_2}}{1 + \lambda} - \zeta_{c_1}) \frac{\partial s_{d,c_2}}{\partial t_{c_1}} + (\frac{\lambda \zeta_{c_2}}{1 + \lambda} - \zeta_{c_1}) \frac{\partial s_{0,c_2}}{\partial t_{c_1}}}_{\text{Impact on } C_1/C_2 \text{ margin}}] \end{aligned}$$

A.3 Policy Updating

Return to the regulator's first-order optimization conditions for constrained and unconstrained instruments. By reordering terms, we can see that these conditions require the regulator to equate the marginal abatement cost of a given policy with its marginal benefits from externality reduction. In the case of the optimal Pigouvian tax, the first-order condition for t_j can be rewritten as:

$$\begin{aligned} \frac{\partial W}{\partial t_j} = & \underbrace{(1 + \lambda)(t_d) \frac{\partial s_{d,d}}{\partial t_j} + (1 + \lambda)(t_{c_1}) \frac{\partial s_{c_1,c_1}}{\partial t_j} + (t_d + \lambda t_{c_2}) \frac{\partial s_{d,c_2}}{\partial t_j} + (\lambda t_{c_2}) \frac{\partial s_{0,c_2}}{\partial t_j}}_{\text{Marginal Abatement Cost}} \\ & + \underbrace{(1 + \lambda)(-\zeta_d) \frac{\partial s_{d,d}}{\partial t_j} + (1 + \lambda)(-\zeta_{c_1}) \frac{\partial s_{c_1,c_1}}{\partial t_j} + (-\zeta_d - \lambda \zeta_{c_2}) \frac{\partial s_{d,c_2}}{\partial t_j} + (-\lambda \zeta_{c_2}) \frac{\partial s_{0,c_2}}{\partial t_j}}_{\text{Marginal Benefit}} = 0 \end{aligned}$$

Likewise, for the efficient second-best policy when there are two generations of clean technology, the regulator's first-order condition for t_{c_1} can be expressed as:

$$\begin{aligned} \frac{\partial W}{\partial t_{c_1}} = & \underbrace{(1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}(t_{c_1})}_{\text{Marginal Abatement Cost}} \\ & + \underbrace{(1 + \lambda) \frac{\partial s_{d,d}}{\partial t_{c_1}}(-\zeta_d) + (1 + \lambda) \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}(-\zeta_{c_1}) + \frac{\partial s_{d,c_2}}{\partial t_{c_1}}(-\zeta_d - \lambda \zeta_{c_2}) + \lambda \frac{\partial s_{0,c_2}}{\partial t_{c_1}}(-\zeta_{c_2})}_{\text{Marginal Benefit}} = 0 \end{aligned}$$

This formulation allows us to consider the case where the regulator has already imposed some t_{c_1} and is now considering marginal deviations from the original tax (or subsidy), for

example in response to new information about a second-generation clean technology that was not available when the policy was initially developed. We can use the ratio of marginal benefits to marginal abatement costs to identify the direction in which the current policy should be updated. For the efficient second-best policy, this ratio should be approximately equal to 1:

$$\frac{\text{Marginal benefit}}{\text{Marginal abatement cost}} = \frac{\zeta_{c_1} + (\frac{\partial s_{d,d}}{\partial t_{c_1}} / \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) \zeta_d + (\frac{\partial s_{d,c_2}}{\partial t_{c_1}} / \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) (\frac{\zeta_d + \lambda \zeta_{c_2}}{1+\lambda}) + (\frac{\partial s_{0,c_2}}{\partial t_{c_1}} / \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}) (\frac{\lambda \zeta_{c_2}}{1+\lambda})}{t_{c_1}}$$

A.4 Comparative Statics

Externality Values

Given the assumptions in the text that $A = L = 2$, we have the following comparative statics around the relative externality values:

$$\frac{\partial \hat{t}_{c_1}}{\partial \zeta_d} \propto -\frac{\partial s_{d,d}}{\partial t_{c_1}}(1 + \lambda) - \frac{\partial s_{d,c_2}}{\partial t_{c_1}} \leq 0$$

$$\frac{\partial \hat{t}_{c_1}}{\partial \zeta_{c_1}} \propto -\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}}(1 + \lambda) \geq 0$$

$$\frac{\partial \hat{t}_{c_1}}{\partial \zeta_{c_2}} \propto -\frac{\partial s_{d,c_2}}{\partial t_{c_1}} \lambda - \frac{\partial s_{0,c_2}}{\partial t_{c_1}} \lambda \leq 0$$

C_1 Lifetime and C_2 Arrival Rate

With more general timing assumptions, the regulator's first-order condition for the efficient price on C_1 is given by:

$$\begin{aligned} \frac{\partial W}{\partial t_{c_1}} &= \frac{\partial s_{d,d}}{\partial t_{c_1}} \left(\frac{1 - \lambda^L}{1 - \lambda} \right) (-\zeta_d) + \frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} \left(\frac{1 - \lambda^L}{1 - \lambda} \right) (t_{c_1} - \zeta_{c_1}) \\ &+ \frac{\partial s_{d,c_2}}{\partial t_{c_1}} \left[\left(\frac{1 - \lambda^{A-1}}{1 - \lambda} \right) (-\zeta_d) + \left(\frac{1 - \lambda^{L-A+1}}{1 - \lambda} \right) (-\lambda^{A-1} \zeta_{c_2}) \right] + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} \left(\frac{1 - \lambda^{L-A+1}}{1 - \lambda} \right) (-\lambda^{A-1} \zeta_{c_2}) = 0 \end{aligned}$$

where A is the arrival time of C_2 and L is the lifetime of C_1 ; we assume $L \geq A$.

Comparative statics with respect to A are proportional to:

$$\begin{aligned}\frac{\partial \hat{t}_{c_1}}{\partial A} &\propto \frac{\lambda^{A-1} \ln(\lambda)}{(1-\lambda)} \left[\frac{\partial s_{d,c_2}}{\partial t_{c_1}} (\zeta_d - \zeta_{c_2}) + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (-\zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{d,d}}{\partial t_{c_1} \partial A} \left(\frac{1-\lambda^L}{1-\lambda} \right) (-\zeta_d) + \frac{\partial^2 s_{c_1,c_1}}{\partial t_{c_1} \partial A} \left(\frac{1-\lambda^L}{1-\lambda} \right) (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial^2 s_{d,c_2}}{\partial t_{c_1} \partial A} \left[\left(\frac{1-\lambda^{A-1}}{1-\lambda} \right) (-\zeta_d) + \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{0,c_2}}{\partial t_{c_1} \partial A} \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2})\end{aligned}$$

Likewise, comparative statics with respect to L are proportional to:

$$\begin{aligned}\frac{\partial \hat{t}_{c_1}}{\partial L} &\propto \frac{-\lambda^L \ln(L)}{1-\lambda} \left[\left(\frac{\partial s_{d,d}}{\partial t_{c_1}} \right) (-\zeta_d) + \left(\frac{\partial s_{c_1,c_1}}{\partial t_{c_1}} \right) (-\zeta_{c_1}) + \left(\frac{\partial s_{d,c_2}}{\partial t_{c_1}} \right) (-\zeta_{c_2}) + \left(\frac{\partial s_{0,c_2}}{\partial t_{c_1}} \right) (-\zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{d,d}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^L}{1-\lambda} \right) (-\zeta_d) + \frac{\partial^2 s_{c_1,c_1}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^L}{1-\lambda} \right) (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial^2 s_{d,c_2}}{\partial t_{c_1} \partial L} \left[\left(\frac{1-\lambda^{A-1}}{1-\lambda} \right) (-\zeta_d) + \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{0,c_2}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2})\end{aligned}$$

Note that we can substitute in the original first-order condition to rewrite this expression as:

$$\begin{aligned}\frac{\partial \hat{t}_{c_1}}{\partial L} &\propto \frac{-\lambda^L \ln(\lambda)(1-\lambda^{A-1})}{(1-\lambda)(1-\lambda^L)} \left[\frac{\partial s_{d,c_2}}{\partial t_{c_1}} (\zeta_d - \zeta_{c_2}) + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (-\zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{d,d}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^L}{1-\lambda} \right) (-\zeta_d) + \frac{\partial^2 s_{c_1,c_1}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^L}{1-\lambda} \right) (t_{c_1} - \zeta_{c_1}) \\ &\quad + \frac{\partial^2 s_{d,c_2}}{\partial t_{c_1} \partial L} \left[\left(\frac{1-\lambda^{A-1}}{1-\lambda} \right) (-\zeta_d) + \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2}) \right] \\ &\quad + \frac{\partial^2 s_{0,c_2}}{\partial t_{c_1} \partial L} \left(\frac{1-\lambda^{L-A+1}}{1-\lambda} \right) (-\lambda^{A-1} \zeta_{c_2})\end{aligned}$$

In order to obtain closed-form expressions for the derivatives of $\frac{\partial s_{i,j'}}{\partial t_{c_1}}$ with respect to A and L , we need distributional assumptions for heterogeneous preference parameters. For analytic tractability, I assume that ν_{c_1} , ν_{c_2} , and ν_d are jointly uniform on the unit cube;

other distributional assumptions are also possible.

Under this distributional assumption, we must consider three ranges of parameter values, which in turn determine the expressions for share derivatives. To simplify notation, I define $p_{c_1}^t = p_{c_1} + t_{c_1}$. These cases are presented in the table below:

Condition	$\frac{1-p_{c_1}^t}{1-p_{c_2}} \leq \frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L}$	$\frac{1-\hat{\nu}_d+p_d-p_{c_1}^t}{1-\hat{\nu}_d+p_d-p_{c_2}} = \frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L}$ for $\hat{\nu}_d \in [p_d, 1]$	$\frac{p_d-p_{c_1}^t}{p_d-p_{c_2}} \geq \frac{\lambda^{A-1}-\lambda^L}{1-\lambda^L}$
$\frac{\partial^2 s_{d,d}}{\partial t_{c_1} \partial A}$	0	0	0
$\frac{\partial^2 s_{d,d}}{\partial t_{c_1} \partial L}$	0	0	0
$\frac{\partial^2 s_{c_1,c_1}}{\partial t_{c_1} \partial A}$	$(\frac{1}{2} + p_d - p_{c_1}^t - \frac{p_d^2}{2}) \cdot a$	$\frac{(1-\hat{\nu}_d)(1-\hat{\nu}_d+2p_d-2p_{c_1}^t)}{2} \cdot a$ $-f(\frac{\partial \hat{\nu}_d}{\partial p_{c_1}^t}, \frac{\partial \hat{\nu}_d}{\partial A}, \frac{\partial^2 \hat{\nu}_d}{\partial p_{c_1}^t \partial A})$	0
$\frac{\partial^2 s_{c_1,c_1}}{\partial t_{c_1} \partial L}$	$(\frac{1}{2} + p_d - p_{c_1}^t - \frac{p_d^2}{2}) \cdot l$	$\frac{(1-\hat{\nu}_d)(1-\hat{\nu}_d+2p_d-2p_{c_1}^t)}{2} \cdot l$ $-f(\frac{\partial \hat{\nu}_d}{\partial p_{c_1}^t}, \frac{\partial \hat{\nu}_d}{\partial L}, \frac{\partial^2 \hat{\nu}_d}{\partial p_{c_1}^t \partial L})$	0
$\frac{\partial^2 s_{d,c_2}}{\partial t_{c_1} \partial A}$	$\frac{-(1-p_d)(1+p_d-2p_{c_1}^t)}{2} \cdot a$	$\frac{-(1-\hat{\nu}_d)(1-\hat{\nu}_d+2p_d-2p_{c_1}^t)}{2} \cdot a$ $+f(\frac{\partial \hat{\nu}_d}{\partial p_{c_1}^t}, \frac{\partial \hat{\nu}_d}{\partial A}, \frac{\partial^2 \hat{\nu}_d}{\partial p_{c_1}^t \partial A})$	0
$\frac{\partial^2 s_{d,c_2}}{\partial t_{c_1} \partial L}$	$\frac{-(1-p_d)(1+p_d-2p_{c_1}^t)}{2} \cdot l$	$\frac{-(1-\hat{\nu}_d)(1-\hat{\nu}_d+2p_d-2p_{c_1}^t)}{2} \cdot l$ $+f(\frac{\partial \hat{\nu}_d}{\partial p_{c_1}^t}, \frac{\partial \hat{\nu}_d}{\partial L}, \frac{\partial^2 \hat{\nu}_d}{\partial p_{c_1}^t \partial L})$	0
$\frac{\partial^2 s_{0,c_2}}{\partial t_{c_1} \partial A}$	$-p_d(1-p_{c_1}^t) \cdot a$	0	0
$\frac{\partial^2 s_{0,c_2}}{\partial t_{c_1} \partial L}$	$-p_d(1-p_{c_1}^t) \cdot l$	0	0

where $a = \frac{(1-\lambda^L)(\lambda^{A-1})\ln(\lambda)}{(\lambda^{A-1}-\lambda^L)^2}$ and $l = \frac{(-\lambda^L)(1-\lambda^{A-1})\ln(\lambda)}{(\lambda^{A-1}-\lambda^L)^2}$. In the middle column, $f(\cdot)$ captures the impact of t_{c_1} and A or L on the value of $\hat{\nu}_d$.

We can substitute these terms into the general expressions above for $\frac{\partial \hat{t}_{c_1}}{\partial A}$ and $\frac{\partial \hat{t}_{c_1}}{\partial L}$. The signs of these comparative statics depend on the parameter values. However, it is easiest to

gain intuition in the case where $\frac{p_d - p_{c_1}^t}{p_d - p_{c_2}} \geq \frac{\lambda^{A-1} - \lambda^L}{1 - \lambda^L}$, which yields:

$$\frac{\partial \hat{t}_{c_1}}{\partial A} \propto \underbrace{\ln \lambda}_{\leq 0} \underbrace{\left[\frac{\partial s_{d,c_2}}{\partial t_{c_1}} (\zeta_d - \zeta_{c_2}) + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (-\zeta_{c_2}) \right]}_{\leq 0 \text{ or } \geq 0}$$

$$\frac{\partial \hat{t}_{c_1}}{\partial L} \propto \underbrace{-\lambda^L \ln \lambda}_{\geq 0} \underbrace{\left[\frac{\partial s_{d,c_2}}{\partial t_{c_1}} (\zeta_d - \zeta_{c_2}) + \frac{\partial s_{0,c_2}}{\partial t_{c_1}} (-\zeta_{c_2}) \right]}_{\leq 0 \text{ or } \geq 0}$$

We see here that the efficient second-best subsidy for C_1 moves in opposite directions when the arrival rate of C_2 is increasing versus the lifetime of C_1 is increasing. Under this parameterization, I find that as long as the advent of C_2 induces a sufficiently large shift away from consuming the dirty product, then the efficient second-best subsidy is increasing (i.e., t_{c_1} becomes more negative) in the arrival time of the second-generation technology. That is, as the arrival rate of the improved efficient technology recedes into the longer term, the near-term benefits of switching consumers from the dirty technology to the early efficient technology increasingly outweigh the benefits of waiting. Conversely, as long as C_2 induces more sufficient switching from the dirty product instead of the outside option, the efficient second-best subsidy decreases as the lifetime of the first-generation efficient product lengthens. That is, the efficient subsidy decreases as early deployment of the first-generation product represents greater lock-in away from the later technology.

B Data Appendix

B.1 Lighting Sales Data and Product Characteristics

I collected data on UPC product characteristics from several third-party aggregators, including Semantics3, Barcode Lookup, and UPCitemdb. In other cases, I looked up UPC product characteristics using search engines and archived product catalogs. I collected data on bulb shape (e.g., A, PAR, G), base type (e.g., E12, E17, E26), wattage, wattage-

equivalent, technology, and light color. As reported in the main text, I filtered the sample of bulbs to include only the general purpose bulbs subject to the first phase of the federal efficiency standard. This step entailed dropping directional bulbs (e.g., floodlights or reflector bulbs), decorative bulbs (e.g., globe or candle-shaped bulbs), bulbs with candelabra or intermediate bases, appliance bulbs, and specialty bulbs (e.g., bug lamps, colored bulbs, rough service bulbs, or three-way bulbs).

I initially sought to collect information on whether a product was Energy Star rated or dimmable, since these characteristics may also have been important to consumers. However, I found that reporting of these characteristics was too inconsistent to be included in my demand model without introducing substantial measurement error.⁶² These unobserved characteristics are therefore captured in the ξ_{jmt} econometric error. In my final specification, I also do not include light color directly, since I found that many products were marketed as “soft white” with a large range in the associated correlated color temperature (CCT). Because many products reported *only* the “soft white” label without the accompanying value of CCT, I determined that there was extensive measurement error in this variable as well.⁶³

B.2 Demographic Characteristics and Lighting Market Characteristics

Annual demographic information at the county level is taken from the American Community Survey, including number of households, median household income, share of college educated individuals. As noted in the main text, in my specification with endogenous coming-to-market, I hold fixed cross-sectional variation in the demographics of the overall county population, varying only the demographics of the sub-population entering the market in each

⁶²Energy Star bulbs also had to be approved as such, so a given UPC might be Energy Star-rated in one period but not in another.

⁶³Light color is expressed in degrees Kelvin; “warm” light has fewer degrees Kelvin while “cool” light has more; an open flame produces light around 1900K.

period. I use data on the share of college educated individuals from 2013.

I use estimated county-level data on climate change attitudes from the Yale Project on Climate Change Communication (Howe et al., 2015). In my primary specification, I use the estimated percentage who think that global warming is caused mostly by human activities. In alternative model specifications, I use the estimated percentage who somewhat or strongly support requiring utilities to produce 20% of electricity from renewable sources. These variables are intended to capture the share of the population that linked climate change to human activities such as electricity use and the share that supported contemporary policy responses. Since these county-level estimates are only available for certain years, I use 2014 data on climate attitudes for my coming-to-market specifications.

Lastly, I use information on the share of Democratic voters from Dave Leip’s Atlas of Presidential Elections (Leip, 2017). Because party voter registration rules vary by state, I use the share of Democratic votes cast in a particular election as a proxy for the share of Democratic voters. I use the 2012 U.S. House of Representatives election since it occurred in the middle of my study period, had higher turnout from coinciding with a presidential election, and was arguably more likely to be driven by party affiliations than personalities or other idiosyncratic factors that may drive presidential or senatorial election results.

I use survey data on the average number of lightbulbs per household, from the DOE’s regular Lighting Market Characterization Reports, plus information on the number of households per county to estimate the total potential residential lighting market in a particular county.

Information on lighting operating costs is taken from the EIA Form 861, which all utilities are required to complete. Using this data, I calculate the average residential electricity rate by dividing total revenue in the sector by total sales volume; I follow the procedure outlined in Borenstein and Bushnell (2018) to estimate the combined energy and delivery charge for areas with deregulation at the retail level.⁶⁴

⁶⁴As noted, this procedure yields estimates of *average* residential rates for a given utility, where rational consumers with perfect attention and perfect information would instead optimize based on *marginal* rates.

B.3 Light Bulb Failure Rates

I use information on light bulb failure rates from the spreadsheet accompanying the U.S. DOE’s Technical Support Document for analyzing the impact of lighting efficiency standards (U.S. Department of Energy, 2016). These failure rates are calculated using data on the distribution of daily hours-of-use for each light bulb in the residential sector and engineering data on the rate of bulb failure as a function of elapsed hours of use. Light bulb survival is assumed to follow a Weibull distribution, expressed as a function of the percentage of rated lifetime for which the bulb has been used. Failure rates for CFLs are further adjusted by observed cycles of on-off switching in the residential sector, which resulted in actual lifetimes that were shorter than rated lifetimes. No adjustments are made for LEDs, in the absence of evidence to suggest that on-off cycle time affected realized lifetimes.

The DOE’s Technical Support Document provides failure rates for CFLs and LEDs but not for incandescents and halogens. As such, I assume that the shape of the survival distribution is the same for incandescents and halogens as for LEDs, again in the absence of evidence to suggest that on-off switching affected the actual lifetimes for those bulbs. That is, I assume that the parameters of the Weibull distribution for incandescents and halogens are the same as for LEDs and simply apply a different technology-specific lifetime (1200 hours for incandescents and 2000 hours for halogens). Following the notation in DOE’s Technical Support Document, the probability of survival is given by the following expression:

$$P_{surv}(A) = \exp \left[- \left(\frac{L_{cons}(A) \cdot 100}{\lambda} \right)^k \right]$$

where

$$L_{cons}(A) = \sum_{i=1}^n F_{sect}(hb_i) \cdot hb_i \cdot \left(\frac{365 \cdot A}{l_{rated}} \right)$$

A gives the age of the bulb in years; k gives the shape parameter of the Weibull distribution

However, marginal rates are often not fully transparent for residential customers, and other research suggests that average rates may be more salient (Ito, 2014).

and λ gives the scale parameter; hb_i gives daily hours-of-use (binned); $F_{sect}(hb_i)$ gives the sector-specific frequency rates for each hours-of-use bin; i indexes each hours-of-use bin; and l_{rated} gives the rated lifetime in hours. In my main specification, I adopt the following Weibull parameters by technology:

Technology	Shape Parameter (k)	Scale Parameter (λ)	Rated Lifetime (l_{cons})
LED	1.718	118.1	25,000 hours
Incandescent	1.718	118.1	1,200 hours
Halogen	1.718	118.1	2,000 hours
CFL	1.253	61.0	8,000 hours

The Weibull parameters used in the DOE’s Technical Support Document are originally estimated from California Public Utilities Commission data on failure rates of CFLs as a function of average on-cycle time. For residential LEDs, the DOE analysis adopts the parameters associated with average on-cycle time of 180 minutes, noting that “GSL survival data is commonly presented using three-hour on-time cycle lengths.” (The final Weibull parameters adopted for residential CFLs are weighted averages of parameters associated with various on-cycle lengths, given the average on-cycle length by room type and distribution of bulbs by room type.) However, using these parameters originally estimated from CFL failure rates for LEDs (and in my analysis, for incandescents and halogens as well) results in mean and median effective lifetimes that are considerably longer than the rated lifetime.⁶⁵ Consequently, I also estimate my demand model using the Weibull parameters associated with 90-minute average cycle length (shape parameter = 1.805 and scale parameter = 118.1), since this specification produces mean and median effective lifetimes that are much closer to the expected service length. These estimates are presented in Column 3 of Table 7.

⁶⁵Based on the parameterization in the DOE’s Technical Support Document, the effective mean lifetime for LEDs is 138% of the expected service length and the effective median lifetime is 125% of expected service length.

B.4 Externality Calculation

To calculate the average externality per hour of lighting sold during the study period, I adopt several simplifying assumptions. First, I assume that the externality per kWh of electricity is constant across the U.S. and over the study period. This assumption ignores the improvements in the emissions intensity of U.S. electricity generation that occurred during the study period. This assumption also ignores the heterogeneity of emissions intensity across the U.S., which means we might overestimate the relative externality reduction from purchasing efficient lighting technologies, since they were heavily purchased in areas of the country with cleaner electricity grids (e.g., California and the Northeast). I use data on the emissions intensity of the U.S. electric grid in 2014 from U.S. Environmental Protection Agency (2017), reflecting the midpoint of my study period. I convert total CO_2 emissions to dollar values using estimates of the Social Cost of Carbon (SCC) from Interagency Working Group on the Social Cost of Carbon, U.S. Government (2015), assuming a 2.5% discount rate. I do not consider the impact of local air pollutants associated with electricity generation.

I then calculate the externality associated with each bulb category (technology by wattage-equivalent) purchased in a given period. To do so, I must determine in which periods and for how long a given bulb will be used. Here I make two simplifying assumptions. First, I assume that all bulbs last for exactly their rated lifetime, with a downward adjustment for CFLs due to on-off switching in my main specification. Second, I assume that all bulbs purchased during the study period are used for an average of 5.56 hours per day. Given these assumptions, the wattages listed in Table 1, and the externality values described above, I can then calculate the discounted externality associated with each bulb as a function of the period in which it is sold. I then divide by the total number of lifetime hours to obtain the average discounted externality per hour of lighting sold during the study period, which is the outcome metric reported throughout the main text.

A few comments are warranted about my assumptions around bulb lifetime and hours of use. My assumed average hours of daily use is significantly higher than the average hours of

daily use across the installed base of light bulbs in the residential sector (see, for example, Navigant Consulting (2012, 2017)). This higher value is intended to account for the fact that the hours-of-use conditional on being replaced is higher than the unconditional hours-of-use across the entire population of light bulbs, since bulbs that are used more intensively will fail more frequently. Since I want to capture the average hours-of-use across the set of new bulbs purchased, rather than across the population of installed bulbs, it is appropriate to use this conditional mean. I use the distribution of residential hours-of-use provided in the spreadsheet accompanying the DOE’s Technical Support Document to simulate average hours of use conditional on bulb failure. Nonetheless, my approach does not account for any correlations between more intensive lighting use and willingness to invest in energy efficiency lighting technologies, in the spirit of Dubin and McFadden (1984); I further discuss this issue in the main text.

B.5 Instruments for Demand Estimation

Data for demand-side instruments is taken from a variety of sources. Price indices for semiconductors, reflecting the price of manufacturing LED chips, and fluorescent ballasts, a key component in CFL manufacturing, were both taken from the Bureau of Labor Statistics (BLS) Import and Export Price Indices. Rare earth prices are taken from the U.S. Geological Survey’s National Minerals Information Center, which provides annual price indices for a variety of rare earth minerals; I then computed a weighted average of the relevant mineral prices, using weights that reflect the composition of rare earths used in CFLs, taken from the DOE (2016) for the second phase of the lighting efficiency standards. On the retail cost side, county-level average wage data in the retail sector – including wages specific to food and beverage stores (NAICS code 445), health and personal care stores (446), and general merchandise retailers (452) – were taken from the BLS Quarterly Census of Employment and Wages. Quarterly data on diesel prices, a proxy for retail transportation costs, was collected from the U.S. EIA by regional Petroleum Administration for Defense Districts (PADDs).

Finally, data on commercial real estate indices by U.S. region was taken from the Society of Industrial and Office Realtors (SIOR).

C Coming-to-Market

C.1 Estimation Details

Below I describe my modified estimation algorithm. I describe in detail the steps of the algorithm applied in the first period of my study period, which are then repeated for all subsequent periods:

1. I begin with an initial distribution of remaining time-to-failure for each consumer type's installed base of light bulbs at the start of the study period. Given the assumption that consumers enter the market to replace failed bulbs, I use this initial distribution to simulate the share of each consumer type entering the market in the first period. Let w_{imt} denote the share of consumers of type i entering market m in period t .
2. For each consumer type, I calculate $\mu_{ijmt}(\boldsymbol{\theta}_2)$ for all products available in the first period, given some $\boldsymbol{\theta}_2$.
3. Given starting values for the common component of utility δ_{jmt}^0 , I use μ_{ijmt} and δ_{jmt}^0 to predict individual choice probabilities s_{ijmt}^h , again for the first period ($t = 1$). Here h indexes each iteration of the algorithm, so $h = 1$ for this first iteration.
4. Next I predict overall market shares in the first period by calculating the weighted average of predicted choice probabilities by type, where weights are equal to the share of a given type that enters the market:

$$s_{jmt}^h = \frac{\sum_{i=1}^{ns} w_{imt}^h \cdot s_{ijmt}^h}{\sum_{i=1}^{ns} w_{imt}^h}$$

5. I use the vector of predicted market shares $s_{mt}^h = \{s_{jmt}^h\}$ to update mean utility levels using the BLP fixed point algorithm, δ_{mt}^{h+1} , for all products in the first period.
6. I repeat steps 3-5 until fixed point convergence, with h increasing by 1 for each iteration. (After convergence, I drop h superscripts for the final set of predicted individual choice probabilities and coming-to-market weights.)
7. I use individual choice probabilities for each technology to update the distribution of remaining time-to-failure for the installed base of bulbs. I use engineering data on failure rates by technology and assume that new bulbs are installed in the same period that they are purchased. The distribution of remaining time-to-failure therefore evolves heterogeneously by consumer type. I use this updated distribution to determine the share of each type that enters the market in the next period $w_{im,t+1}$.
8. I repeat steps 2-7 for period $t + 1$, and then for each remaining period.

I follow the standard BLP approach of using non-linear search methods to solve for the value of θ_2 that sets the gradient of the objective function as close as possible to 0. I estimate the model using one-step GMM; I use $W = [Z'Z]^{-1}$ to weight my over-identified moment conditions, where Z represents a matrix of instrumental variables. To the extent possible, I follow the best practices for discrete choice demand estimation outlined in Conlon and Gortmaker (2020). I use a numerical gradient-based optimization routine, with tolerance for the inner loop set to 10^{-14} and tolerance for the outer loop set to 10^{-5} . I use multiple starting values for θ_2 in my optimization routine. Standard errors are bootstrapped at the county level (forthcoming).

C.2 Analytic Gradient

C.2.1 Analytic Gradient for Standard BLP

In this section, I follow the notation in the main body of the paper, where θ_2 represents non-linear parameters in the demand system; δ represents a vector of the common components of indirect utility from choosing each product; and ξ represents a vector of the residuals in the common component of utility. For standard demand-side-only BLP, the gradient of the GMM objective function is given by:

$$\nabla q(\theta_2) = 2G(\theta_2)'Wg(\theta_2)$$

where $g(\theta_2)$ is the moment function, W is a weighting matrix, and $G(\theta_2)$ is the gradient of the moment function.

$$G(\theta_2) = \frac{1}{N}Z' \frac{\partial \xi}{\partial \theta_2}$$

where N gives the total number of observations. Because residuals are linear, we can write:

$$\frac{\partial \xi}{\partial \theta_2} = \frac{\partial \delta}{\partial \theta_2}$$

Noting that shares s , δ , and θ_2 are related through the inversion function $s(\delta_{mt}; \theta_2) = s_{mt}$ for market m in period t , we can use the implicit function theorem to rewrite this derivative market-by-market and period-by-period as follows:

$$\frac{\partial \delta_{mt}}{\partial \theta_2}(\theta_2) = - \left[\frac{\partial s_{mt}}{\partial \delta_{mt}}(\theta_2) \right]^{-1} \left[\frac{\partial s_{mt}}{\partial \theta_2}(\theta_2) \right]$$

Alternatively, to see the elements of these matrices, we index each product by j and

expand as follows:

$$\begin{bmatrix} \frac{\partial \delta_{1,mt}}{\partial (\theta_2)_1} & \cdots & \frac{\partial \delta_{1,mt}}{\partial (\theta_2)_L} \\ \vdots & & \vdots \\ \frac{\partial \delta_{J,mt}}{\partial (\theta_2)_1} & \cdots & \frac{\partial \delta_{J,mt}}{\partial (\theta_2)_L} \end{bmatrix} = - \begin{bmatrix} \frac{\partial s_{1,mt}}{\partial \delta_{1,mt}} & \cdots & \frac{\partial s_{1,mt}}{\partial \delta_{J,mt}} \\ \vdots & & \vdots \\ \frac{\partial s_{J,mt}}{\partial \delta_{1,mt}} & \cdots & \frac{\partial s_{J,mt}}{\partial \delta_{J,mt}} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial s_{1,mt}}{\partial (\theta_2)_1} & \cdots & \frac{\partial s_{1,mt}}{\partial (\theta_2)_L} \\ \vdots & & \vdots \\ \frac{\partial s_{J,mt}}{\partial (\theta_2)_1} & \cdots & \frac{\partial s_{J,mt}}{\partial (\theta_2)_L} \end{bmatrix}$$

Recall from the main text that the non-linear parameters θ_2 consist of parameters on demographics interacted with product characteristics (π_{kd}) and on unobservable preference heterogeneity interacted with product characteristics (σ_k). Following earlier notation, let D_{id} represent demographic characteristics and let ν_i^k represent random preference draws, where k indexes product characteristics and d indexes demographic characteristics. The elements of the above matrices are then defined as follows:

$$\frac{\partial s_{jmt}}{\partial \delta_{jmt}} = \frac{1}{ns} \sum_{i=1}^{ns} \frac{\partial s_{ijmt}}{\partial \delta_{jmt}} = \frac{1}{ns} \sum_{i=1}^{ns} s_{ijmt} (1 - s_{ijmt})$$

$$\frac{\partial s_{jmt}}{\partial \delta_{j'mt}} = \frac{1}{ns} \sum_{i=1}^{ns} \frac{\partial s_{ijmt}}{\partial \delta_{j'mt}} = -\frac{1}{ns} \sum_{i=1}^{ns} s_{ijmt} s_{ij'mt}$$

$$\begin{aligned} \frac{\partial s_{jmt}}{\partial \sigma_k} &= \frac{1}{ns} \sum_{i=1}^{ns} \frac{\partial s_{ijmt}}{\partial \sigma_k} = \frac{1}{ns} \sum_{i=1}^{ns} s_{ijmt} (x_{jmt}^k \nu_i^k - \sum_{j'=1}^J x_{j'mt}^k \nu_i^k s_{ij'mt}) \\ &= \frac{1}{ns} \sum_{i=1}^{ns} \nu_i^k s_{ijmt} (x_{jmt}^k - \sum_{j'=1}^J x_{j'mt}^k s_{ij'mt}) \end{aligned}$$

$$\begin{aligned} \frac{\partial s_{jmt}}{\partial \pi_{kd}} &= \frac{1}{ns} \sum_{i=1}^{ns} \frac{\partial s_{ijmt}}{\partial \pi_{kd}} = \frac{1}{ns} \sum_{i=1}^{ns} s_{ijmt} (x_{jmt}^k D_{id} - \sum_{j'=1}^J x_{j'mt}^k D_{id} s_{ij'mt}) \\ &= \frac{1}{ns} \sum_{i=1}^{ns} D_{id} s_{ijmt} (x_{jmt}^k - \sum_{j'=1}^J x_{j'mt}^k s_{ij'mt}) \end{aligned}$$

C.2.2 Analytic Gradient for BLP with Coming-to-Market

When we modify the BLP procedure to incorporate coming-to-market, we still have the same basic expression for the gradient of the GMM objective function, and we still use

the implicit function theorem to rewrite the derivative of mean utilities with respect to the nonlinear parameters (as shown above). Now, however, we can only separate these derivatives market by market – but not period by period. We now have:

$$\begin{bmatrix} \frac{\partial s_{1m1}}{\partial \theta_{21}} & \dots & \frac{\partial s_{1m1}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{Jm1}}{\partial \theta_{21}} & \dots & \frac{\partial s_{Jm1}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{1mT}}{\partial \theta_{21}} & \dots & \frac{\partial s_{1mT}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{JmT}}{\partial \theta_{21}} & \dots & \frac{\partial s_{JmT}}{\partial (\theta_2)_L} \end{bmatrix} = - \begin{bmatrix} \frac{\partial s_{1m1}}{\partial \delta_{1m1}} & \dots & \frac{\partial s_{1m1}}{\partial \delta_{Jm1}} & \dots & \frac{\partial s_{1m1}}{\partial \delta_{1mT}} & \dots & \frac{\partial s_{1m1}}{\partial \delta_{JmT}} \\ \dots & & \dots & & \dots & & \dots \\ \frac{\partial s_{Jm1}}{\partial \delta_{1m1}} & \dots & \frac{\partial s_{Jm1}}{\partial \delta_{Jm1}} & \dots & \frac{\partial s_{Jm1}}{\partial \delta_{1mT}} & \dots & \frac{\partial s_{Jm1}}{\partial \delta_{JmT}} \\ \dots & & \dots & & \dots & & \dots \\ \frac{\partial s_{1mT}}{\partial \delta_{1m1}} & \dots & \frac{\partial s_{1mT}}{\partial \delta_{Jm1}} & \dots & \frac{\partial s_{1mT}}{\partial \delta_{1mT}} & \dots & \frac{\partial s_{1mT}}{\partial \delta_{JmT}} \\ \dots & & \dots & & \dots & & \dots \\ \frac{\partial s_{JmT}}{\partial \delta_{1m1}} & \dots & \frac{\partial s_{JmT}}{\partial \delta_{Jm1}} & \dots & \frac{\partial s_{JmT}}{\partial \delta_{1mT}} & \dots & \frac{\partial s_{JmT}}{\partial \delta_{JmT}} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial s_{1m1}}{\partial (\theta_2)_1} & \dots & \frac{\partial s_{1m1}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{Jm1}}{\partial (\theta_2)_1} & \dots & \frac{\partial s_{Jm1}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{1mT}}{\partial (\theta_2)_1} & \dots & \frac{\partial s_{1mT}}{\partial (\theta_2)_L} \\ \dots & & \dots \\ \frac{\partial s_{JmT}}{\partial (\theta_2)_1} & \dots & \frac{\partial s_{JmT}}{\partial (\theta_2)_L} \end{bmatrix}$$

That is, we must now consider how s_{jmt} depends on changes in mean utility for products offered in the same period and products offered in previous periods but in the same market. Letting w_{imt} denote the probability that each of the ns consumer types enters market m and period t , the individual elements of these Jacobians are now defined as follows:

$$\begin{aligned} \frac{\partial s_{jmt}}{\partial \delta_{jmt}} &= \frac{\sum_{i=1}^{ns} w_{imt} \frac{\partial s_{ijmt}}{\partial \delta_{jmt}}}{\sum_{i=1}^{ns} w_{imt}} = \frac{\sum_{i=1}^{ns} w_{imt} s_{ijmt} (1 - s_{ijmt})}{\sum_{i=1}^{ns} w_{imt}} \\ \frac{\partial s_{jmt}}{\partial \delta_{j'mt}} &= \frac{\sum_{i=1}^{ns} w_{imt} \frac{\partial s_{ijmt}}{\partial \delta_{j'mt}}}{\sum_{i=1}^{ns} w_{imt}} = - \frac{\sum_{i=1}^{ns} w_{imt} s_{ijmt} s_{ij'mt}}{\sum_{i=1}^{ns} w_{imt}} \\ \frac{\partial s_{jmt}}{\partial \delta_{j'mt'}} &= \frac{\sum_{i=1}^{ns} s_{ijmt} \frac{\partial w_{imt}}{\partial \delta_{j'mt'}}}{\sum_{i=1}^{ns} w_{imt}} - \frac{\sum_{i=1}^{ns} w_{imt} s_{ijmt}}{\sum_{i=1}^{ns} w_{imt}} \cdot \frac{\sum_{i=1}^{ns} \frac{\partial w_{imt}}{\partial \delta_{j'mt'}}}{\sum_{i=1}^{ns} w_{imt}}, \text{ for } t' < t; 0 \text{ otherwise} \\ \frac{\partial s_{jmt}}{\partial \sigma_k} &= \frac{\sum_{i=1}^{ns} w_{imt} \frac{\partial s_{ijmt}}{\partial \sigma_k}}{\sum_{i=1}^{ns} w_{imt}} + \frac{\sum_{i=1}^{ns} s_{ijmt} \frac{\partial w_{imt}}{\partial \sigma_k}}{\sum_{i=1}^{ns} w_{imt}} - \frac{\sum_{i=1}^{ns} w_{imt} s_{ijmt}}{\sum_{i=1}^{ns} w_{imt}} \cdot \frac{\sum_{i=1}^{ns} \frac{\partial w_{imt}}{\partial \sigma_k}}{\sum_{i=1}^{ns} w_{imt}} \\ \frac{\partial s_{jmt}}{\partial \pi_{kd}} &= \frac{\sum_{i=1}^{ns} w_{imt} \frac{\partial s_{ijmt}}{\partial \pi_{kd}}}{\sum_{i=1}^{ns} w_{imt}} + \frac{\sum_{i=1}^{ns} s_{ijmt} \frac{\partial w_{imt}}{\partial \pi_{kd}}}{\sum_{i=1}^{ns} w_{imt}} - \frac{\sum_{i=1}^{ns} w_{imt} s_{ijmt}}{\sum_{i=1}^{ns} w_{imt}} \cdot \frac{\sum_{i=1}^{ns} \frac{\partial w_{imt}}{\partial \pi_{kd}}}{\sum_{i=1}^{ns} w_{imt}} \end{aligned}$$

A few notes on these expressions. The first two expressions, which capture the effect of

changing one product's mean utility on the shares of that product and others offered in the same market and period, is exactly analogous to the expressions from the standard BLP model, except now we do not weight each type draw equally but according to the share of that type in the market in that period. These weights are already calculated as part of the modified fixed point routine for coming-to-market, and so are straightforward to add to the calculation of the analytic gradient.

The third expression is new, and captures the effect of changing a past product's mean utility on the shares of products offered in that market in future periods. This effect occurs entirely through the impact on the distribution of types coming to the market in the future period. We calculate $\frac{\partial w_{imt}}{\partial \delta_{j'mt'}}$ as follows:

$$\frac{\partial w_{imt}}{\partial \delta_{j'mt'}} = w_{imt'} \sum_{j=1}^{J_{t'}} \frac{\partial s_{ijmt'}}{\partial \delta_{j'mt'}} \cdot \text{Prob}(\text{in market in period } t \mid \text{consume } j \text{ in } t')$$

The probability of being in the market in period t conditional on having purchased product j in past period t' is calculated as follows. Consider a simple example with 3 periods and 2 technologies of different lifetimes. Then perform the following calculations:

- Use engineering lifetimes to calculate: Probability(Product 1 fails in 1 period), Probability(Product 2 fails in 1 period), Probability(Product 1 fails in 2 periods), Probability(Product 2 fails in 2 periods)
- We can then recover directly:
 - Probability(In market period 3 | consume product 1 in period 2) = Probability(Product 1 fails in 1 period)
 - Probability(In market period 3 | consume product 2 in period 2) = Probability(Product 2 fails in 1 period)
 - Probability(In market period 2 | consume product 1 in period 1) = Probability(Product 1 fails in 1 period)

- Probability(In market period 2 | consume product 2 in period 1) = Probability(Product 2 fails in 1 period)
- We then use the demand model to calculate: Probability(Consume product 1 | in market period 2) and Probability(Consume product 2 | in market period 2)
- We can then recover:
 - Probability(In market period 3 | consume product 1 in period 1) = Probability(Product 1 fails in 2 periods) + Probability(Product 1 fails in 1 periods) · Probability(Consume product 1 | in market period 2) · Probability(Product 1 fails in 1 periods) + Probability(Product 1 fails in 1 periods) · Probability(Consume product 2 | in market period 2) · Probability(Product 2 fails in 1 periods)
 - And analogously for Probability(In market period 3 | consume product 1 in period 1).

In this simple example, we have now recovered Probability(in market in period t | consume j in t') for all j and all $t > t'$.

For the fourth and fifth terms above, we see that the overall derivative is composed of two effects, a direct effect of changing σ_k or π_{kd} on s_{imjk} for each i , and an indirect effect on the distribution of types in the market for period t . The direct effect is calculated in essentially the same manner as in the original BLP, except that type draws are now weighted according to their probability of being in the market.

However, the indirect impact of σ_k or π_{kd} on the distribution of types in period t is more difficult to calculate. We would essentially use our calculations of $\frac{\partial s_{ijmt'}}{\partial \sigma_k}$ for all j and all $t' < t$, combined with probabilities of being in the market conditional on choosing some product in some period, to determine the overall impact on the distribution of type draws. The dimensionality of this calculation may be quite high, however.

D Additional Model Specifications

D.1 Additional Demand Estimates


This section presents demand estimates from several additional model specifications in Table 7.

D.2 Additional Counterfactual Results

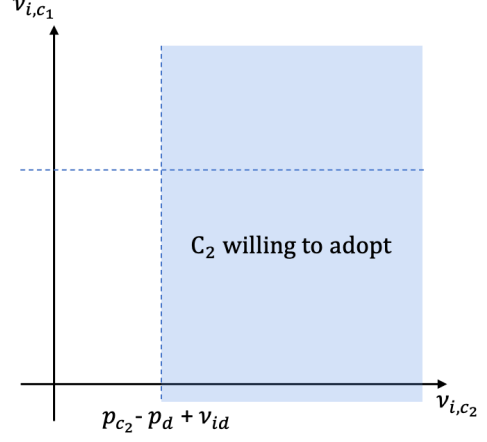
This section presents the main counterfactual outcomes using several additional model specifications; results are presented in Figures 27, 28, 29, and 30.

D.3 Additional Analyses


I also present in Figure 31 the predicted impact on LED quantities and the average discounted externality from simultaneously varying the timing of standards implementation and CFL rebates phase-out.



(a)




(b)




(c)

Figure 1: Consumer Optimization Across Consumption Bundles (D, D) , (C_1, C_1) , and (D, C_2) , Given $v_{id} \geq p_d$




(a)

Figure 2: Consumer Optimization Across Consumption Bundles $(0, 0)$, (C_1, C_1) , and $(0, C_2)$, Given $\nu_{id} < p_d$



(a) Given $\nu_{id} \geq p_d$



(b) Given $\nu_{id} < p_d$

Figure 3: Consumer Optimization Under Optimal Pigouvian Tax

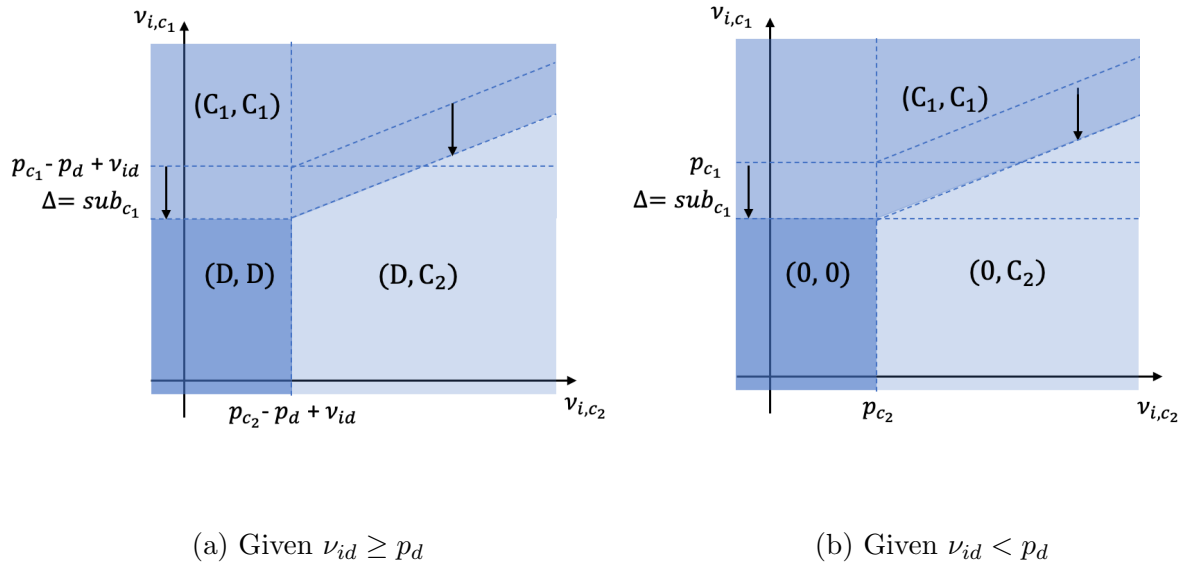


Figure 4: Consumer Optimization Under Subsidy for C_1







Figure 5: Lighting Technologies



(a) Incandescent Bulbs



(b) Halogen Bulbs



(c) CFL Bulbs



(d) LED Bulbs

Figure 6: Median Price Per Bulb, By Technology and Wattage-Equivalent. Source: Nielsen retail scanner data


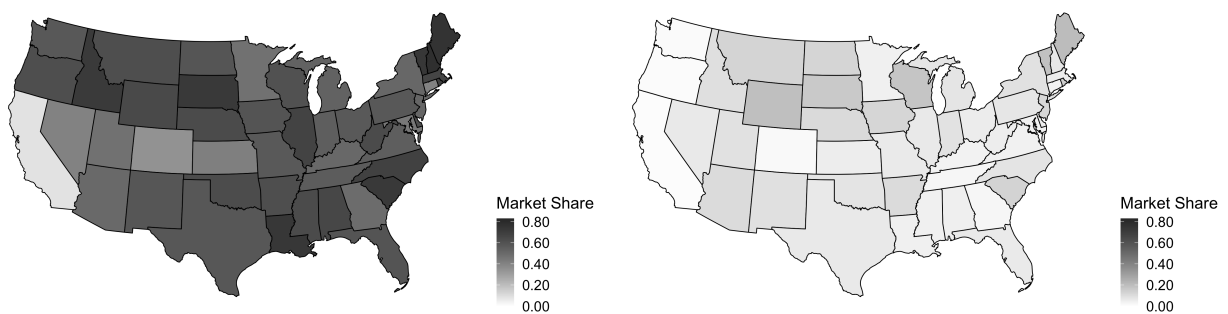



Figure 7: A-shape bulb market share by technology. Source: National Electrical Manufacturers Association (2018)




(a) Incandescent Market Share, Q1 2014

(b) Incandescent Market Share, Q4 2014

Figure 8: Evolution of Incandescent Market Share (2014). Source: Nielsen retail scanner data.




(a) Halogen Market Share, Q1 2014




(b) Halogen Market Share, Q4 2014

Figure 9: Evolution of Halogen Market Share (2014). Source: Nielsen retail scanner data.




(a) CFL Market Share, Q1 2014




(b) CFL Market Share, Q4 2014

Figure 10: Evolution of CFL Market Share (2014). Source: Nielsen retail scanner data.



(a) LED Market Share, Q1 2015



(b) LED Market Share, Q4 2015

Figure 11: Evolution of LED Market Share (2015). Source: Nielsen retail scanner data.




Figure 12: A-shape bulb market share by technology and bulb type. Source: Nielsen retail scanner data.

Notes: Market shares are computed using total quantities observed in Nielsen data.



Figure 13: A-shape bulb market share by technology and bulb type. Source: Nielsen retail scanner data.

Notes: Market shares are computed using total quantities observed in Nielsen data.

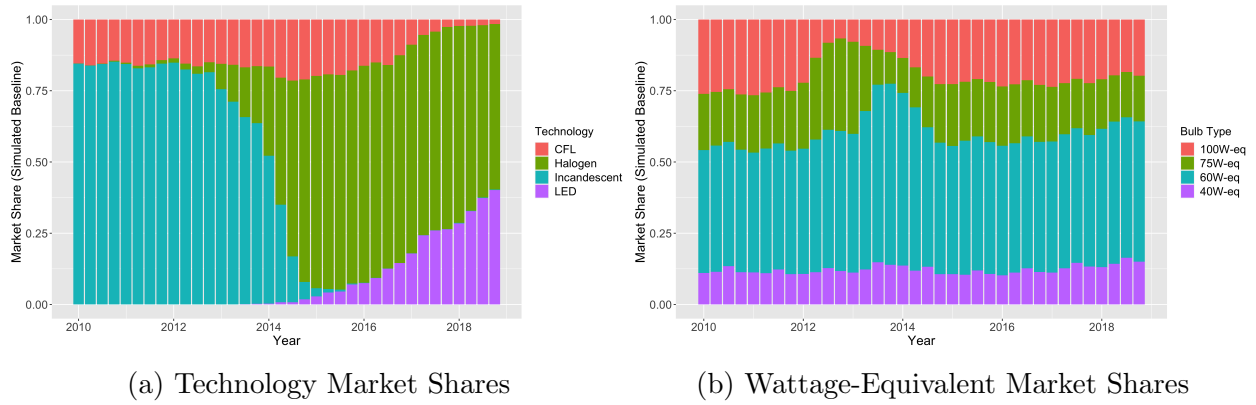


Figure 14: Simulated Baseline Market Shares

Notes: Aggregate shares are simulated using estimated individual choice probabilities, weighted by the county-level distribution of demographics and the number of households in each county.

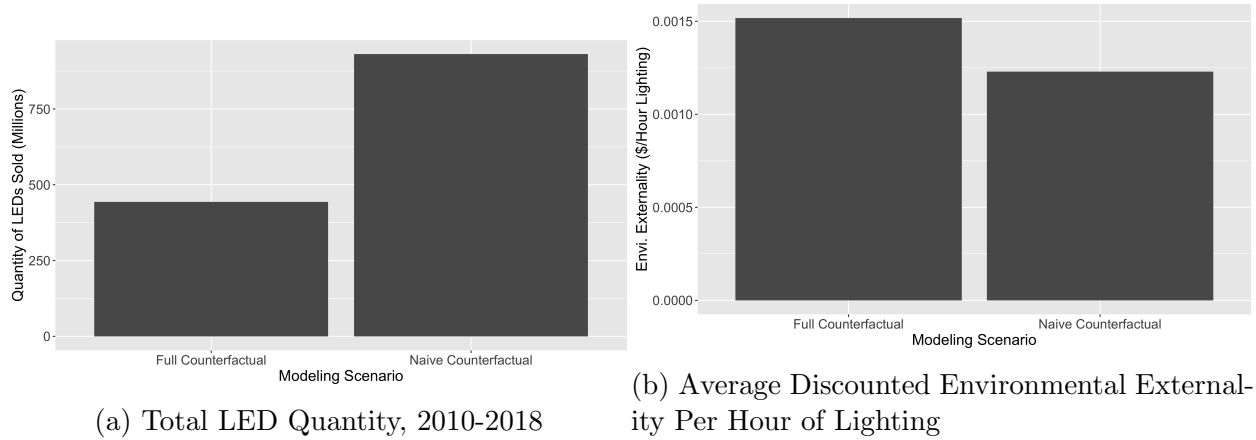


Figure 15: Simulated Baseline Under Full Counterfactual and “Naive” Counterfactual.

Notes: In the “naive counterfactual,” the market size and distribution of consumer heterogeneity are held fixed at Q1 2010 levels. The “full counterfactual” includes endogenous updating of the market size and the market-level distribution of consumer heterogeneity, based on the history of past lighting purchases in each county market. The “full counterfactual” is consistent with the counterfactual simulations I run for alternative policy environments.

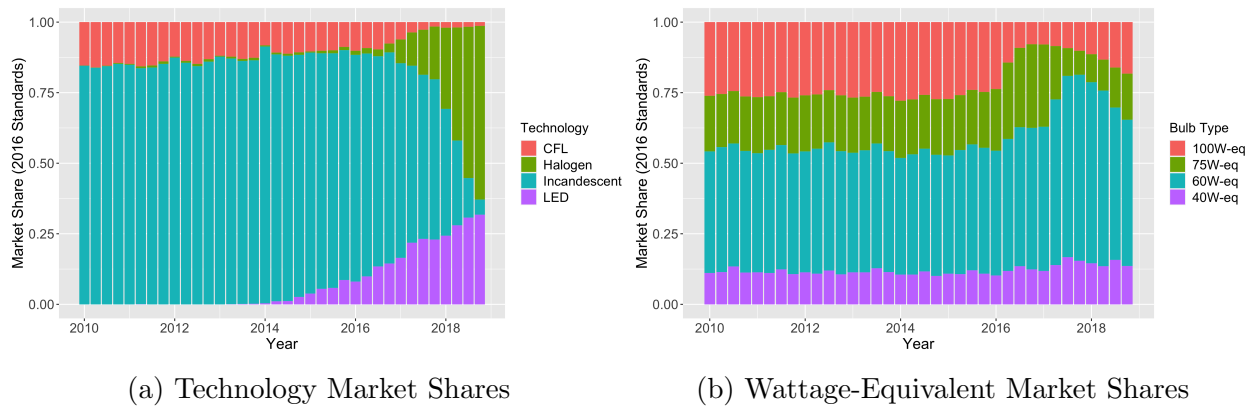


Figure 16: Simulated Counterfactual Market Shares, for Standards Implemented 2016-2018
Notes: Aggregate shares are again simulated using estimated individual choice probabilities, weighted by the county-level distribution of demographics and the number of households in each county. For this counterfactual policy simulation, I also vary the product set to be consistent with efficiency standards implemented in 2016 (or in 2015 for California).




Figure 17: Total Quantity of LEDs Sold, by Timing of Standards Implementation
Notes: Total quantities of LEDs sold over the study period are simulated from market-level individual choice probabilities, weighted by county-level demographics and the number of households in each county. Panel (a) presents the difference in total LED quantities relative to a baseline of standards implemented beginning in 2012, the actual policy implementation.




Figure 18: Average Discounted Environmental Externality Per Hour of Lighting, by Timing of Standards Implementation

Notes: The average discounted externality per hour of lighting sold during the study period is simulated from individual choice probabilities by technology and wattage-equivalent. The externality is calculated assuming the wattages presented in Table 1 and lifetimes of 1,200 hours for incandescents, 2,000 hours for halogens, 8,000 hours for CFLs (adjusted for on-off switching), and 25,000 hours for LEDs. I assume 1.13 lb CO_2 per kWh of electricity consumed across the U.S. (EPA, 2017) and apply social cost of carbon values from IWG (2015), assuming 2.5% discount rate. Panel (a) presents the difference in the average discounted externality relative to a baseline of standards implemented beginning in 2012, the actual policy implementation.

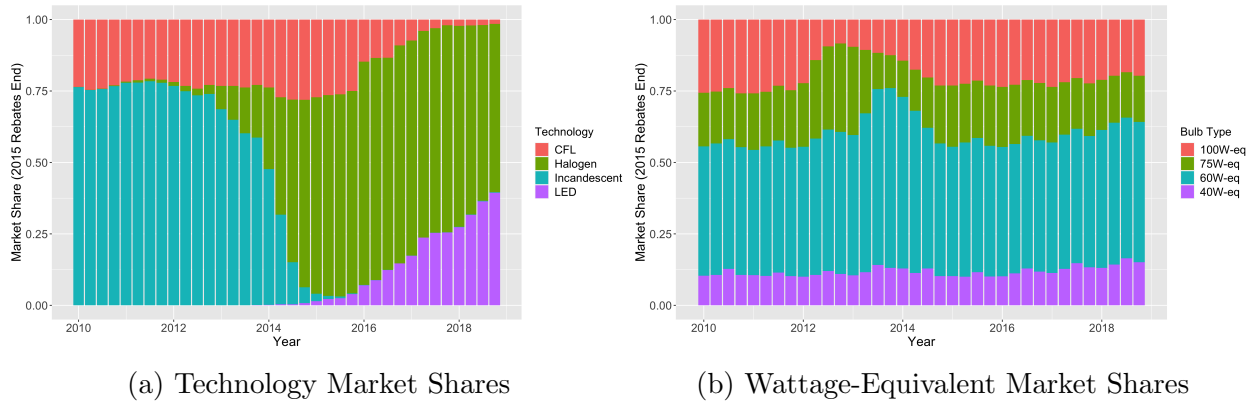


Figure 19: Simulated Counterfactual Market Shares, for CFL Subsidies Ended After 2015

Notes: Aggregate shares are again simulated using estimated individual choice probabilities, weighted by the county-level distribution of demographics and the number of households in each county. For this counterfactual policy simulation, I also impose a \$1 rebate per CFL bulb until the fourth quarter of 2015.




Figure 20: Total Quantity of LEDs Sold, by Phase-Out of CFL Rebates

Notes: Total LED quantities are simulated as in Figure 17, here with \$1 rebates per CFL bulb through some phase-out year. Panel (a) presents total quantities of LEDs relative to a baseline of the rebates phased out after 2012.




Figure 21: Average Discounted Environmental Externality Per Hour of Lighting, by Phase-Out of CFL Rebates

Notes: The average discounted externality per hour of lighting is simulated as in Figure 18, here with \$1 rebates per CFL bulb through some phase-out year. Panel (a) presents the average discounted externality relative to a baseline of the rebates phased out after 2012.




Figure 22: Trade-Off Between LED Quantities and Average Environmental Externality




Figure 23: Alternative Rates of LED Price Decline Equivalent to Counterfactual Policy Timing

Notes: Median observed LED prices decline by approximately 6.5% per quarter. I simulate market outcomes under alternative rates of LED price declines by multiplying each observed LED price by the ratio of the counterfactual rate of price decline to the observed rate of price decline, adjusted for the relevant period. I then compare total LEDs sold under each of these alternative price trajectories to the total LEDs sold under alternative policy timings.

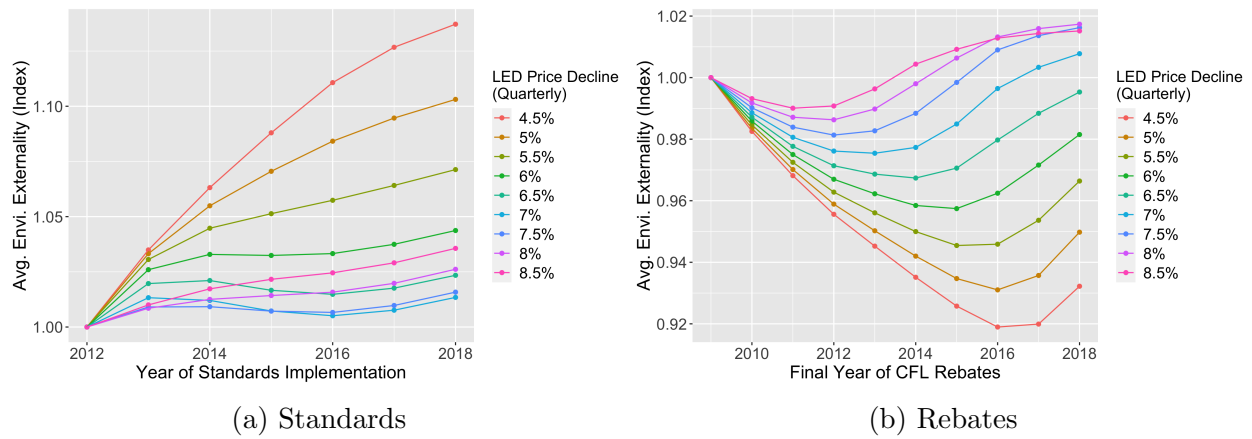


Figure 24: Sensitivity of Policy Timing to Alternate Rates of LED Price Decline
Notes: Counterfactual simulations use alternative rates of LED quarterly price declines and alternative policy timings. For each LED price trajectory, the average discounted externality is normalized at 1 for the policy scenario that implements standards beginning in 2012 (panel a) or phases out CFL rebates after 2009 (panel b).

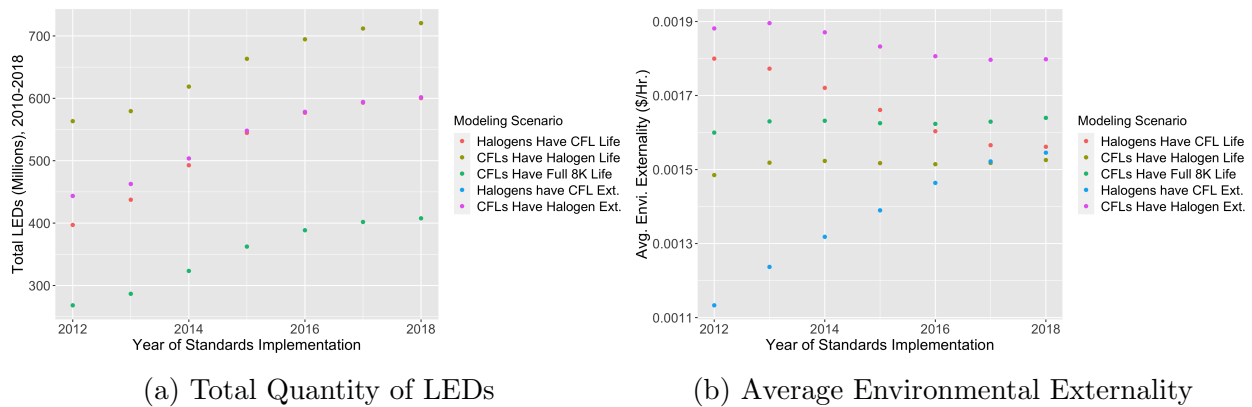




Figure 25: Sensitivity of Standards Timing to Alternate Technology Lifetimes or Externality Values
Notes: Counterfactual scenarios vary key technology parameters (lifetime and externality) and standards timing, holding fixed estimated demand parameters.




(a) Total Quantity of LEDs




(b) Average Environmental Externality

Figure 26: Sensitivity of Rebates Timing to Alternate Technology Lifetimes or Externality Values


Notes: Counterfactual scenarios vary key technology parameters (lifetime and externality) and rebates timing, holding fixed estimated demand parameters.




(a) Price Specified in Logs




(b) Democratic Voter Interacted with Efficient Technology



(c) California Counties Dropped




(d) No Operating Cost Variable




(e) CFL Product Set Changed Based on Standards Timing

Figure 27: Impact of Standards Timing on Total LED Quantities, Under Alternative Model Specifications


Notes: As robustness tests, I simulate total quantities of LEDs purchased over the study period as a function of the timing of standards implementation, for several different demand model specifications.




(a) Price Specified in Logs




(b) Democratic Voter Interacted with Efficient Technology



(c) California Counties Dropped




(d) No Operating Cost Variable




(e) CFL Product Set Changed Based on Standards Timing

Figure 28: Impact of Standards Timing on Average Discounted Externality, Under Alternative Model Specifications


Notes: As robustness tests, I simulate the average environmental externality per hour of lighting purchased over the study period as a function of the timing of standards implementation, for several different demand model specifications.




(a) Price Specified in Logs



(b) Democratic Voter Interacted with Efficient Technology



(c) California Counties Dropped



(d) No Operating Cost Variable

Figure 29: Impact of Rebates Timing on Total LED Quantities, Under Alternative Model Specifications

Notes: As robustness tests, I simulate total quantities of LEDs purchased over the study period as a function of the final year of CFL rebates, for several different demand model specifications.

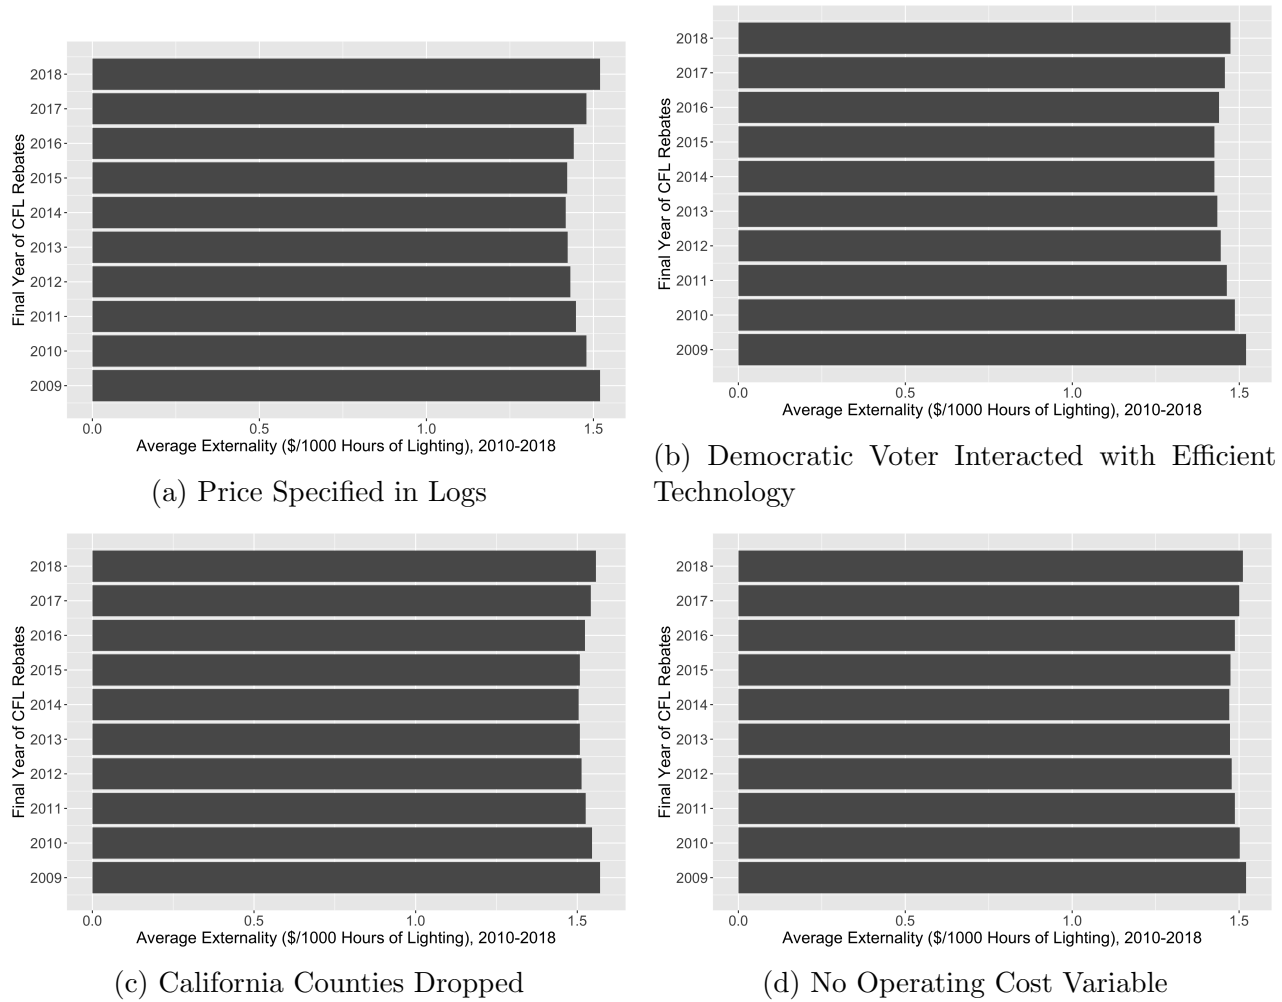


Figure 30: Impact of Rebates Timing on Average Discounted Externality, Under Alternative Model Specifications

Notes: As robustness tests, I simulate the average externality per hour of lighting purchased over the study period as a function of the final year of CFL rebates, for several different demand model specifications.

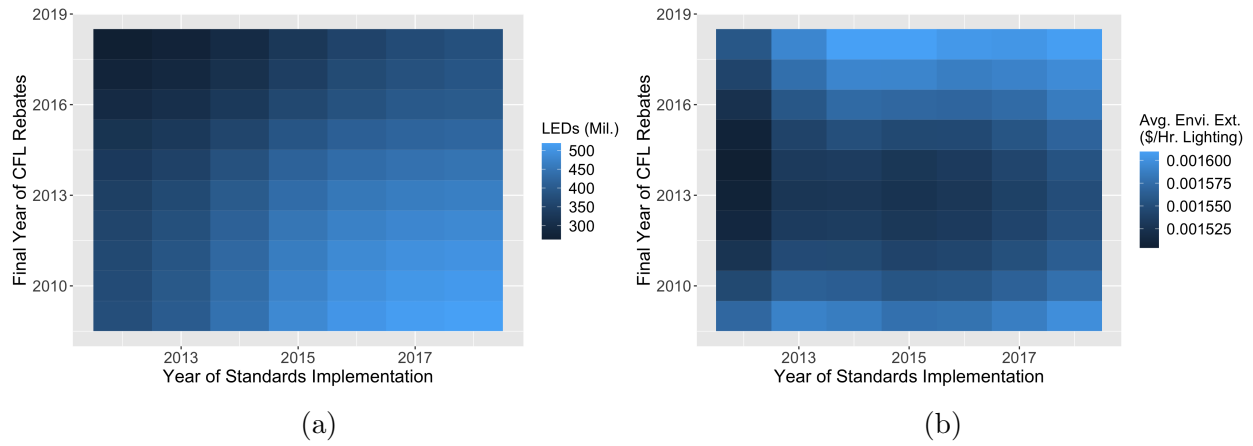


Figure 31: Impact of Altering Timing of Both Standards and Rebates

Notes: Using my baseline model specification (for which estimated demand parameters are presented in Column 1 of Table 5), I simultaneously vary the timing of standards implementation and the phase-out of CFL rebates. I then calculate the predicted quantity of LEDs sold and average externality for lighting purchased during the study period.

Bulb Type	Incandescent	Halogen	CFL	LED
450 lumens	40W	29W	10W	5.5W
800 lumens	60W	43W	13W	8.5W
1,100 lumens	75W	53W	18W	11W
1,600 lumens	100W	72W	23W	14.5W

Table 1: Typical Wattages for General Purpose Bulbs by Technology

Technology	Incand.	Halogen	CFL	LED
Share 40W-eq	0.115	0.118	0.122	0.210
Share 60W-eq	0.494	0.399	0.521	0.677
Share 75W-eq	0.222	0.230	0.141	0.0512
Share 100W-eq	0.169	0.254	0.217	0.0625
Share Private Label	0.232	0.200	0.197	0.208
Mean Price	\$0.54	\$1.81	\$3.44	\$7.35
Median Price	\$0.48	\$1.74	\$3.39	\$5.91
Std. Dev. Price	\$0.28	\$0.53	\$1.30	\$5.09
Min Price	\$0.10	\$0.11	\$0.10	\$0.23
Max Price	\$5.00	\$6.99	\$17.99	\$52.99
Mean Share	0.0993	0.0349	0.0101	0.0210
Median Share	0.0623	0.0674	0.0168	0.00552
Std. Dev. Share	0.107	0.0785	0.0209	0.0397

Table 2: Descriptive Statistics of Product Characteristics

Notes: I calculate summary statistics for the sample of data used in demand estimation. The share of bulb characteristics (wattage-equivalent and brand) by technology is calculated using the total quantities of bulbs sold. Summary statistics for prices and shares by technology are calculated using product-county-quarter observations.

Instrument	Mean	Std. Dev.	Min.	Max.	Obs.
Rare earth prices	\$132.63	\$133.65	\$27.13	\$416.8	9
Semiconductor prices	\$80.83	\$2.90	\$76.10	\$86.53	36
Fluorescent ballast prices	\$194.56	\$1.38	\$191.70	\$195.80	14
Retail wages	\$429.10	\$44.77	\$327.18	\$627.70	90,429
Commercial real estate	\$96.87	\$33.47	\$32.10	\$184.80	1,836
Diesel prices	\$3.33	\$0.63	\$1.97	\$4.33	288

Table 3: Descriptive Statistics for Cost Shifters

	(1) OLS	(2) IV Logit	(3) IV Logit	(4) OLS	(5) IV Logit	(6) IV Logit
Average Price	-0.231 (0.00391)	-0.674 (0.0222)		-0.241 (0.00436)	-0.705 (0.0248)	
Log Average Price			-2.119 (0.0696)			-2.093 (0.0743)
CFL × 100	-1.381 (0.0278)	0.995 (0.0916)	1.100 (0.0924)	-1.343 (0.0287)	1.117 (0.100)	1.067 (0.0983)
CFL × 40	-1.812 (0.0280)	-0.552 (0.0729)	-0.435 (0.0747)	-2.017 (0.0240)	-0.448 (0.0805)	-0.463 (0.0794)
CFL × 60	-0.295 (0.0210)	1.263 (0.0732)	1.466 (0.0788)	-0.135 (0.0208)	1.360 (0.0802)	1.438 (0.0839)
CFL × 75	-1.269 (0.0234)	0.455 (0.0849)	0.535 (0.0841)	-1.994 (0.0283)	0.574 (0.0937)	0.503 (0.0896)
Halogen × 100	-3.903 (0.0766)	0.913 (0.0795)	0.660 (0.0823)	0.125 (0.0589)	0.969 (0.0802)	0.649 (0.0825)
Halogen × 40	-3.112 (0.0316)	0.365 (0.0500)	0.181 (0.0477)	-3.089 (0.0319)	0.419 (0.0523)	0.169 (0.0487)
Halogen × 60	0.936 (0.0366)	1.734 (0.0590)	1.552 (0.0583)	0.953 (0.0369)	1.788 (0.0607)	1.540 (0.0591)
Halogen × 75	-3.876 (0.0487)	1.069 (0.0636)	0.836 (0.0638)	0.307 (0.0427)	1.122 (0.0648)	0.825 (0.0643)
Incandescent × 100		-1.256 (0.0875)	-3.291 (0.124)		-1.238 (0.0883)	-3.270 (0.126)
Incandescent × 40	-0.282 (0.0321)	-1.385 (0.0531)	-3.041 (0.0695)	-0.276 (0.0322)	-1.362 (0.0547)	-3.027 (0.0707)
Incandescent × 60	1.323 (0.0449)	-0.332 (0.0707)	-2.267 (0.0879)	-0.607 (0.0612)	-0.313 (0.0722)	-2.249 (0.0896)
Incandescent × 75	-1.189 (0.0625)	-0.899 (0.0730)	-2.805 (0.0967)	0.705 (0.0547)	-0.879 (0.0740)	-2.787 (0.0984)
LED × 100	0.194 (0.400)	24.04 (1.522)	-1.603 (0.433)	2.757 (0.496)	2.623 (0.137)	2.108 (0.124)
LED × 40	-2.097 (0.207)	13.08 (0.833)	-1.609 (0.315)	-1.234 (0.148)	1.060 (0.0636)	1.207 (0.0695)
LED × 60	-0.303 (0.352)	21.24 (1.273)	-1.804 (0.361)	-3.265 (0.686)	2.559 (0.0599)	2.606 (0.0612)
LED × 75	0.321 (0.410)	28.39 (1.512)	-1.473 (0.556)	-0.984 (0.223)	1.749 (0.124)	1.355 (0.111)
Private Label	-1.380 (0.0157)	-1.572 (0.0164)	-1.680 (0.0181)	-1.382 (0.0157)	-1.578 (0.0165)	-1.676 (0.0186)
Operating Costs	-0.0128 (0.00603)	-0.00522 (0.00666)	0.0232 (0.00794)	-0.0128 (0.00604)	-0.00497 (0.00677)	0.0227 (0.00794)
Before/After Stds.	Yes	Yes	Yes	Yes	Yes	Yes
Time Trends x LEDs	Yes	Yes	Yes	No	No	No
Time Dummies x LEDs	No	No	No	Yes	Yes	Yes
F-stat	-	270.51	435.06	-	291.04	341.23
N	560098	560098	560098	560098	560098	560098

Standard errors in parentheses (clustered at county level)

Table 4: Demand Estimates from OLS and IV Logit Specifications

Notes: I estimate demand parameters for OLS and IV logit specifications without accounting for consumer heterogeneity. In IV specifications (Columns 2, 3, 5, and 6), instruments include manufacturer cost shifters, retailer cost shifters, and differentiation instruments. In all specifications, I interact product category with an indicator for whether the relevant efficiency standards had taken effect, for all incandescent, halogen, and CFL products. In Columns 1-3, I interact LED product categories with linear and quadratic time trends; in Columns 4-6, I instead interact LED product categories with period-specific time dummies. Column 2 most closely matches my baseline specification that allows for consumer heterogeneity.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Average Price	-1.070	-0.971	-1.687	-0.790	-1.071	-1.370	
Log Average Price							-2.938
CFL \times 100	-6.773	0.523	-0.768	-3.030	-7.545	-12.371	-6.101
CFL \times 40	-8.413	-1.034	-2.640	-4.60	-9.184	-14.424	-7.642
CFL \times 60	-6.540	0.843	-0.945	-2.780	-7.311	-11.937	-5.644
CFL \times 75	-7.379	-0.013	-1.434	-3.582	-8.150	-12.808	-6.720
Halogen \times 100	1.360	1.394	3.380	1.095	1.359	-1.779	0.794
Halogen \times 40	0.964	0.879	2.446	0.561	0.965	-0.668	0.477
Halogen \times 60	2.280	2.237	3.938	1.924	2.280	-0.298	1.803
Halogen \times 75	1.561	1.549	3.313	1.250	1.561	-1.130	1.023
Incandescent \times 100	-1.408	-1.182	0.138	-1.232	-1.412	0.748	-4.259
Incandescent \times 40	-1.238	-1.190	-0.255	-1.313	-1.239	0.219	-3.620
Incandescent \times 60	-0.315	-0.196	0.847	-0.284	-0.317	1.637	-3.057
Incandescent \times 75	-0.925	-0.770	0.409	-0.854	-0.927	0.899	-3.622
LED \times 100	48.518	38.990	78.799	28.191	48.546	-3.846	5.123
LED \times 40	18.707	20.209	40.074	12.641	18.817	-7.253	-5.513
LED \times 60	34.841	32.778	64.400	22.914	34.952	-6.164	-3.177
LED \times 75	53.568	45.812	89.663	33.000	53.641	-4.773	3.375
Private Label	-1.736	-1.720	-2.021	-1.621	-1.737	-1.853	-1.827
Operating Costs	0.0310	0.007	-0.054	-9.72e-05	0.031	0.053	0.0643
CFL, LED \times Human	10.789					16.289	9.694
CFL, LED \times College		4.446					
CFL, LED \times Democratic			9.188				
CFL, LED \times Support RPS				5.285			
CFL \times Human					11.569		
LED \times Human					10.690		
Before/After Stds.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time Trends x LEDs	Yes	Yes	Yes	Yes	Yes	No	Yes
Time Dummies x LEDs	No	No	No	No	No	Yes	No
<i>N</i>	560098	560098	560098	560098	560098	560098	560096

Standard errors are bootstrapped and clustered at county level (in progress)

Table 5: Demand Estimates from Logit Specifications with Consumer Heterogeneity

Notes: I estimate demand parameters for specifications that account for consumer heterogeneity using observed distributions of demographics. Columns 1-4 vary the demographic characteristics included in the model (expressed belief in human-caused climate change, college educated, Demographic voter, expressed support for state Renewable Portfolio Standard policies). Column 5 allows for separate interactions between demographics and CFL and LED bulbs, respectively. Column 6 uses period-specific time dummies interacted with LED product categories, instead of linear and quadratic time trends. Column 7 specifies prices per bulb in logs rather than in levels.

	Price Level		Price Log	
Own-Price Elasticity	Mean	Median	Mean	Median
Incandescent	-0.52	-0.46	-2.64	-2.75
Halogen	-1.80	-1.65	-2.72	-2.82
CFL	-3.54	-3.51	-2.82	-2.87
LED	-7.68	-6.18	-2.79	-2.90

Table 6: Estimated Own-Price Elasticities by Technology

Notes: I calculate mean and median own-price elasticities by technology using the specifications in Column 1 of Table 5 (prices in levels) and Column 7 of Table 5 (prices in logs).

	(1) No CA	(2) No Oper. Cost	(3) Alt. Lifetimes
Average Price	-1.070	-0.971	-1.687
CFL \times 100	-6.773	0.523	-0.768
CFL \times 40	-8.413	-1.034	-2.640
CFL \times 60	-6.540	0.843	-0.945
CFL \times 75	-7.379	-0.013	-1.434
Halogen \times 100	1.360	1.394	3.380
Halogen \times 40	0.964	0.879	2.446
Halogen \times 60	2.280	2.237	3.938
Halogen \times 75	1.561	1.549	3.313
Incandescent \times 100	-1.408	-1.182	0.138
Incandescent \times 40	-1.238	-1.190	-0.255
Incandescent \times 60	-0.315	-0.196	0.847
Incandescent \times 75	-0.925	-0.770	0.409
LED \times 100	48.518	38.990	78.799
LED \times 40	18.707	20.209	40.074
LED \times 60	34.841	32.778	64.400
LED \times 75	53.568	45.812	89.663
Private Label	-1.736	-1.720	-2.021
Operating Costs	0.0310		-0.054
CFL, LED \times Human	10.789	10.457	11.001
Before/After Stds.	Yes	Yes	Yes
Time Trends x LEDs	Yes	Yes	Yes
California Incl.	No	Yes	Yes
Alt. Lifetimes	No	No	Yes
<i>N</i>	530893	560098	560098

Standard errors are bootstrapped and clustered at county level (in progress)

Table 7: Demand Estimates from Logit Specifications with Consumer Heterogeneity

Notes: As robustness tests, I estimate demand parameters for additional specifications that account for consumer heterogeneity using observed distributions of demographics. Column 1 drops all California counties, since the timing of efficiency standards was different in California relative to the rest of the U.S. Column 2 drops the operating cost variable. Column 3 uses alternative bulb failure rates by technology, varying the parameters of the Weibull survival distribution from U.S. Department of Energy (2016).