

CHAPTER 11

## The Tradeoffs between Energy Efficiency, Consumer Preferences, and Economic Growth

*James Broughel*

The traditional economic rationale for government intervention in the economy is market failure.<sup>1</sup> Underlying the market failure concept is the idea that, because of certain market frictions known as transaction costs, beneficial gains from trade are prevented from occurring that would otherwise increase social welfare. The usual cases where these kinds of transaction costs are present are situations involving externalities, asymmetric information, monopoly power, or the provision of public goods. In such situations, the government can potentially be a corrective force to improve welfare.

While these sources of market failure have led to no shortage of regulatory and other policy interventions, in recent years government agencies and academics have begun justifying policy interventions on the grounds that an additional form of market failure exists, known as a “behavioral market failure.”<sup>2</sup> They extend the standard list of “neoclassical market failures” mentioned above to include instances of suboptimal individual decision-making. Behavioral market failures occur because of various cognitive biases afflicting individuals, which result in people making poor decisions that reduce their own welfare.

The concept of behavioral market failure is an outgrowth of research in the field of behavioral economics, a subset of economic theory that

focuses on the intersection of economics and psychology and tends to emphasize the cognitive limitations of human beings. Behavioral economics can be contrasted with neoclassical economics, which assumes a high degree of rationality on the part of economic agents. According to the neoclassical view, consumers can generally be expected to act in a self-interested manner that accords with their own welfare, given the constraints they face in terms of resources and information. Neoclassical economists therefore assume that consumers' preferences can be inferred simply by observing their behavior. If a consumer engages in a particular transaction, that consumer is assumed to be made better off by the transaction. Otherwise, the consumer would have held on to his or her money and the transaction would not have taken place. In this way, consumers "reveal" their preferences through their actions, giving rise to the term *revealed preference*.

Behavioral economists hold a different view. They routinely drop this revealed preference assumption, instead arguing that cognitive bias prevents consumers from advancing their own welfare in many cases. The "true" preferences of individuals and their behavior in the marketplace are often not aligned, even in cases where consumers possess all the relevant information needed to make an informed choice. Consumers may buy unhealthy food, for example, because they act hastily and without regard for long-run consequences. They might come to regret their decision later, and would act differently with hindsight. It is for this reason that behavioral economists claim to be able to bestow a benefit on consumers—making them better off "as judged by themselves"<sup>3</sup>—by preventing certain market transactions from occurring. In other words, by overriding consumer choices that are influenced by biased decision-making, policymakers can potentially improve consumer welfare, as gauged by the consumers' own preferences.

Some neoclassical economists acknowledge that behavioral biases are a real phenomenon, but they downplay their importance. For example, biased behavior could be present in markets, but behavior may not be *systematically* biased in one direction or another. Just as there are people who consume more unhealthy food than is optimal, so too there are other people who take a healthy lifestyle too far. In the

aggregate, perhaps these biases cancel each other out. Or perhaps over time biases are systematically weeded out of the marketplace, because biased behavior is punished by market forces while rational behavior is rewarded.<sup>4</sup>

Some examples of situations in which overriding consumer choice could have potentially welfare-improving results include banning or taxing certain unhealthy foods, banning addictive products such as cigarettes or alcohol, and capping interest rates on or banning certain kinds of payday lending. However, not all behavioral interventions are intended to override consumer choice—that is, to take the rather extreme step of banning the activity in question. Some interventions, such as so-called nudges, are choice-preserving.<sup>5</sup> The idea behind nudges is that choices are presented to consumers in a different arrangement (known as the *choice architecture*), or with a different default rule, such that a consumer is more likely to make the welfare-improving choice. For example, rather than banning sugary soda, government regulators could simply require supermarkets to place soda in a less visible part of the store. In that case consumers could still search for the soda if they truly want it, but the soda wouldn't be as salient to them as it might be if it were placed near the checkout aisle.

Similarly, employees could be opted into an employer-administered retirement plan by default, while still being left the option of refusing the plan if they prefer to make alternative arrangements for their retirement. However, even these seemingly benign nudges aren't entirely free of coercion. In the soda example, a supermarket is still told by the government where it must place soda, and in the retirement account example, an employer is told how it must enroll its employees in retirement plans.

According to the neoclassical view, the choice architecture that nudges focus on changing should matter only slightly in most cases. So long as information relevant to a consumer's choice is available relatively costlessly (e.g., is presented in a not-overly-complex manner), the consumer should be able to process the information and make the decision that best fits his or her circumstances. Empirically, however, the choice architecture presented to consumers does seem to matter: consumer

choices often change if information is presented in a slightly different manner or if default options change. For example, when employees are defaulted into a retirement plan, more of them tend to sign up. Similarly, if signing up to be an organ donor becomes the default option at the DMV, more people tend to do so. In this manner, a nudge can potentially exploit biases and help guide consumers toward a better choice.

As may be evident from table 1, the line between a neoclassical market failure owing to asymmetric information and a behavioral market failure owing to cognitive bias can be a fuzzy one. One could easily argue that the choice architecture creates cognitive costs for consumers, and therefore that improving the choice architecture is just a sophisticated way of reducing transaction costs associated with collecting information. That would suggest that the real issue facing policymakers is how to present information in the best possible way to consumers, such that problems of asymmetric information can be overcome.

In fact, regulators sometimes assert both information problems and behavioral bias when they justify regulatory interventions. For example, the Department of Energy (DOE), the Environmental Protection Agency, and the Department of Transportation have done so when setting energy and fuel efficiency standards.<sup>6</sup> Energy efficiency regulations, which are the focus of this chapter, are not nudges because they do not rearrange the choice architecture presented to consumers. Rather, they are product bans, which simply remove certain classes of devices from the marketplace. Products that fail to meet the minimum standards set by the DOE are not legally permitted in the marketplace, and are therefore removed as an option for consumers.

Neoclassical economists generally view product bans as imposing costs on consumers, since they result in fewer options for consumers in the marketplace. In fact, these energy interventions raise up-front costs for consumers; equipment is made more expensive because it has to be remodeled to use less energy. However, at the same time a more energy-efficient device will generate a stream of financial benefits in the future through lower utility bills or (with an automobile) fewer stops at the gas station. A neoclassical economist would typically argue that consumers are in the best position to weigh these trade-offs between

**Table 1. The Neoclassical View vs. the Behavioral View**

	Neoclassical view	Behavioral view
Human beings . . .	. . . are assumed to be rational for modeling purposes.	. . . are viewed as sometimes systematically irrational.
Consumer and business decisions	. . . are individually optimal from the perspective of the agent, given budget and information constraints.	. . . can be systematically biased.
Market transactions	. . . reveal consumer preferences, absent information constraints.	. . . may or may not reveal consumer preferences.
Market failures	. . . result from transaction costs preventing mutually beneficial exchanges.	. . . result from biased decision-making—i.e., cognitive transaction costs—in addition to neoclassical market failures.
Individual behavior	. . . is individually optimal, assuming full information, but may not be socially optimal.	. . . may or may not be optimal from the perspective of the individual or society.
Choice architecture	. . . is not important so long as information is relatively inexpensive to obtain and process.	. . . can influence choices, owing to behavioral bias, even with freely available information.
Product bans	. . . reduce consumer choice. (Absent asymmetric information between buyer and seller, the consumer is made worse off: an unambiguous cost.)	. . . can improve consumer welfare: a potential benefit.
Behavioral market failures	. . . are impossible, because rational, self-interested behavior is assumed.	. . . are possible owing to any one of countless cognitive biases identified by researchers.

up-front costs and future savings. After all, consumers know their own preferences and their relevant financial constraints better than anyone else. If a consumer chooses to purchase a relatively less energy-efficient appliance, according to the neoclassical economist, it must be because that's the option best suited for his or her needs.

Behavioral economists don't see it this way. They believe consumers are often incapable of making the choice that corresponds with their own preferences—for example, because they are myopic, or because they are too present-oriented to fully appreciate benefits that will come to them in the future. Regulatory agencies therefore sometimes claim that eliminating energy-inefficient products from the marketplace (i.e., banning products that some consumers would otherwise purchase) confers a benefit on consumers. Neoclassical economists, on the other hand, view the removal of a product from the marketplace—except in cases where consumers lack critical information about the product—as an unambiguous cost to the consumers who would have opted to purchase the banned product in the absence of a regulation.

This chapter walks through the theory underlying behavioral market failures and provides a brief overview of the “energy efficiency gap,” which is a potential behavioral market failure often asserted to exist in markets for energy-consuming devices. It turns out that the behavioral market failure concept fits squarely within the mainstream market failure theory in economics, suggesting that behavioral market failures need to be taken seriously. However, as we will see, traditional market failure theory is not fully satisfying. Moreover, there are significant knowledge problems facing policymakers when they try to correct behavioral market failures, which make these problems considerably more challenging to address than neoclassical market failures. For example, how should an analyst attempt to identify instances when consumers' actions deviate from their “true” preferences? Neoclassical market failures can be identified through revealed preferences, but if consumers' actions are not a credible reflection of their preferences, then what is? If an analyst can never be certain about what a consumer wants, neither can the analyst ever be certain that a behavioral market failure is present. Furthermore, so many behavioral biases have been

identified in the academic literature that one could argue that almost any decision a consumer makes is biased by cherry-picking a bias that suits the situation, which suggests there is a danger that this kind of rationale for regulating will be abused.

This chapter also provides an example of a regulation where behavioral market failure forms an implicit rationale for the regulation. The example comes from the Department of Energy and involves energy efficiency standards recently set for ceiling fans. It turns out that 80 percent of the benefits the DOE claimed for the ceiling fan regulation, which amounts to \$16.5 billion (2015 dollars), stem from overriding consumer choices. This estimate is highly dependent on an array of assumptions, including assumptions about a product's usage over its lifetime, quality remaining the same between more- and less-efficient devices, and the relevant consumer or business discount rate. Without these benefits, the regulation fails a cost-benefit test, according to the DOE's own analysis.

The chapter concludes with discussion of an underexplored rationale for these energy efficiency regulations: externalities imposed on future generations as a result of slower economic growth. Although evidence is at best ambiguous about whether an actual behavioral market failure exists to justify these regulations, there may be a basis for them that does not aim to correct consumer irrationality but instead aims to boost growth. Whether energy efficiency regulations are the best means to promote economic growth is far from clear. However, exploring this issue could prove useful in other contexts, since there are some significant gaps in the theory of market failure as it now stands. Walking through the logic of how present-day consumers impose externalities on future citizens will make the limitations of current market failure theory more clear.

## **Behavioral Market Failures**

Behavioral market failures are similar in many respects to traditional market failures in that a behavioral market failure can be illustrated using the same static supply-and-demand framework commonly used to illustrate neoclassical market failures. Figure 1 shows a case where


at every unit of production, total willingness to pay for a good is less than the private internal marginal benefit that consumers derive from consuming it. Under laissez-faire conditions, consumers demand quantity  $F$  of the good and pay a price of  $D$ , but consumers are demanding too little relative to what is internally optimal, thereby generating a deadweight loss represented by the triangle  $BCE$ . At each unit of production between quantities  $F$  and  $G$ , the internal marginal benefit to the consumer is greater than what the consumer is willing to pay for the resource. An optimal allocation of resources would therefore have production occur at the quantity  $G$ , with consumers paying a price of  $A$ . However, some behavioral bias is preventing the consumers' marginal willingness to pay from aligning with their internal marginal benefit.

This kind of situation has been referred to as an *internality* because of the divergence between willingness to pay and the internal benefit consumers enjoy.<sup>7</sup> In the case of an internality, there is a deviation between unbiased and biased, or rational and irrational, willingness to pay. The scenario presented in figure 1 is similar to the neoclassical case of a public good or a positive externality. Too little of a good is produced relative to what is socially optimal, so the social benefits of more production exceed private benefits at the margin under laissez-faire.

In the neoclassical case, a similar diagram might be used to describe the market for lawn improvements. Well-maintained lawns and gardens increase the desirability and value of property in a neighborhood, but most people do not fully take external benefits to their neighbors into account when deciding how much time and effort to spend on lawn care. Hence, people may invest too little in lawn maintenance from a social point of view. Similarly, in the behavioral case, people might demand too little healthy food or spend too few resources on exercise equipment or on a gym membership because they are myopic or put too little weight on the long-run benefits of a healthier lifestyle.<sup>8</sup> In either case, the market fails because socially or internally beneficial transactions should be taking place but are prevented from occurring, either because of transaction cost frictions such as externalities or as a result of behavioral bias, which can be thought of as a kind of cognitive transaction cost.



**Figure 1. A Behavioral Market Failure**



### The Energy Efficiency Gap

Figure 1 might also describe a potential behavioral market failure in markets for energy-consuming appliances. If the  $x$  axis represents the quantity of investment in energy-efficient devices, investing quantity F would represent underinvestment relative to what is internally optimal for the consumer.<sup>9</sup> Such underinvestment could be explained by a lack of foresight or by some other behavioral problem that has led the consumer to undervalue future financial savings owing from using a less-energy-intensive device. Underinvestment of this kind provides a potential rationale for government intervention in this market to force people to conserve energy.

A large empirical literature explores the existence of an *energy efficiency gap*,<sup>10</sup> which is sometimes also referred to as the *energy paradox*<sup>11</sup>—phrases that reflect the idea that consumers may underinvest in energy efficiency. This literature traces its origins to a 1979 study that found very high implied consumer discount rates in markets for air conditioners.<sup>12</sup> Some individual discount rates ran as high as 90 percent annually, with an average rate of about 25 percent found in the study, leading the author to posit that consumers might be putting too little weight on future energy savings.

The evolution of the literature on a potential undervaluation problem in markets for energy-consuming products coincided with the explosion of research on behavioral economics, which identified countless biases that afflict individual decision-making. Many of the identified biases relate to situations that involve trade-offs that occur across time, suggesting that decisions could be problematic if they involve up-front costs that produce streams of benefits in the future.

For example, hyperbolic discounting is an anomaly involving discount rates that decline over time.<sup>13</sup> This phenomenon and others could explain why consumers might have difficulty optimizing consumption across their lifetimes,<sup>14</sup> and it is closely connected to the idea that consumers sometimes display “time inconsistent” behavior, in that the optimal choice today might not look optimal when reflecting at some point in the future. Not surprisingly, these issues are often associated with self-indulgent behavior or self-control problems. Examples of biases that could contribute to time inconsistent preferences include inertia, inattention, and procrastination. These biases explain why consumers might come to regret decisions that seem optimal in the moment.

Although the idea that consumers are not the perfectly rational automata found in economic models is hardly controversial, the empirical evidence on the existence of an energy efficiency gap is somewhat mixed. New York University professor Hunt Allcott and University of Chicago economist Michael Greenstone reviewed the large literature on the energy efficiency gap,<sup>15</sup> concluding that there is only limited empirical support for it. Similarly, Yale economist Kenneth Gillingham and Resources for the Future economist Karen Palmer find evidence that the extent of the energy efficiency gap may be overestimated, although they acknowledge the true size of the energy efficiency gap is unknown.<sup>16</sup> More recently, a study by economists Todd Gerarden, Richard Newell, and Robert Stavins expresses more confidence that behavioral explanations play a significant role in explaining the energy efficiency gap, but these authors also suggest other factors, such as measurement and modeling errors in the studies evaluating these markets, are important.<sup>17</sup>

One challenge facing researchers studying the energy efficiency gap is that in order to measure the extent to which a bias is present, an

analyst must make certain modeling assumptions and then calibrate the model with data to determine the extent to which real-world behavior deviates from some theoretical optimum. It is always possible—and indeed likely—that relevant variables are left out of the model and that the data used are imperfect. In other words, there is almost always something a neoclassical economist could point at to justify why the behavior of consumers is rational and the model being used to measure bias is imperfect.<sup>18</sup>

For instance, Hunt Allcott and economist Nathan Wozny estimate the extent to which automobile prices adjust in response to changes in gasoline prices,<sup>19</sup> since full adjustment would suggest a high degree of rationality on the part of car consumers. This kind of study inevitably makes certain assumptions about consumers' expectations about the path of future energy prices, as well as about the discount rate appropriate for a consumer to use to discount savings on future gasoline purchases. Allcott and Wozny conclude that automobile prices adjust in response to energy price increases, but probably not fully. They acknowledge the considerable uncertainty in their estimates, but it's also worth acknowledging that any effort to set policy on the basis of this study or similar studies would be highly premature.

Further complicating matters is the fact that many purchasers of automobiles are profit-oriented businesses. With their armies of lawyers and accountants, do corporations suffer from behavioral bias as well? At the very least, behavioral bias seems less likely on the part of managers in corporations than with individual consumers. While it's not surprising that businesses, through advertising or other means, might seek to exploit the behavioral biases of their customers, it also seems likely that businesses themselves would be savvy enough to identify such manipulations when the businesses are on the consumer side of transactions.

Despite the uncertainty surrounding the extent or even the existence of an energy efficiency gap, federal regulatory agencies have justified policies that cost billions of dollars on the grounds that an energy efficiency gap is present. Not surprisingly, many critics have emerged to condemn these regulations. Some worry about a cascade of potential dangers that could result if the traditional assumption of rationality

that underpins standard economic analysis is discarded.<sup>20</sup> Others argue that even if behavioral bias is present in some markets, the same biases that afflict individual decision-making are also likely to be present in government.<sup>21</sup> This does not imply that regulators can never produce beneficial solutions if behavioral market failures occur. However, information problems, the poor incentives regulators often face, and the possibility that regulators will create traditional market failures or government failures in the process of correcting behavioral market failures are legitimate reasons to be wary of behavioral interventions, as they are of other kinds of regulatory interventions.<sup>22</sup> Furthermore, the evidence supporting many behavioral biases is often weaker than is commonly acknowledged, due to problems with the underlying psychological experiments used in academic research. Many of these studies fail attempts at replication, for example.<sup>23</sup>

At the end of the day, regulators do not know what consumers' true preferences are, which means regulators face a possibly insurmountable knowledge problem. Additionally, the list of behavioral biases identified by researchers is so long that almost any decision could be asserted to be either rational or irrational. Any claim that consumers are made "better off as judged by themselves" by overriding their choices is uncertain at best, because consumer preferences are always going to remain to some extent unobservable.<sup>24</sup>

## **An Example**

While the theory that underlies behavioral market failure appears to follow a reasonable line of logic, demonstrating biased behavior empirically is an altogether different task. How can an outside observer ever know what other people's "true" preferences are? One could simply ask them, perhaps, but a whole host of problems emerge from "stated preference" studies that rely on questionnaires or surveys.<sup>25</sup> Such analyses may be even more unreliable than revealed-preference studies—for example, because there are fewer consequences to being wrong on a survey compared to being wrong in the real world. Survey respondents also often respond in a way that they think will please those conducting the survey, a phenomenon known as social desirability bias.<sup>26</sup>

Yet, despite these knowledge problems, government agencies have finalized dozens of energy efficiency standards based on explicit or implicit behavioral economics reasoning. In fact, some of the largest, most impactful regulations, known as economically significant regulations, fall into this category. Economically significant regulations are those expected to impose an impact of at least \$100 million in a single year, and according to the Office of Information and Regulatory Affairs in Washington, DC, nine such regulations cleared its review process in 2016, three in 2015, and eight in 2014.<sup>27</sup> One regulation, which set energy conservation standard for ceiling fans, was published in the *Federal Register* on January 19, 2017—the day before Donald Trump took office.<sup>28</sup> This regulation is notable because in four years the Trump administration did not allow any new economically significant energy efficiency standards to clear the regulatory review process (although several leftover Obama-era rules were finalized in the last days of the Trump administration).

Consequently, this economically significant regulation, issued at the tail end of the Obama administration, is one of the most recent major-impact energy efficiency rules. In the technical support document associated with the rulemaking, the DOE estimated that consumers will see incremental installed cost increases of \$2.5–\$4.4 billion because of this rule. Meanwhile, consumer operating cost savings were estimated to be \$7.0–\$16.5 billion for the years 2020–2049. The combined net present value of all estimated costs and benefits, including environmental benefits, for the standard ranged from \$8.5 billion to \$16.3 billion (in 2015 dollars), depending on the discount rate used, according to the agency.<sup>29</sup>

The operating-cost savings reflect a prediction by the DOE that it will save consumers and businesses billions on lower utility bills. But it is worth asking why consumers are not taking into account of these savings on their own or if there are factors aside from energy efficiency that are more important to consumers when they make their final purchase decision.

Without the energy efficiency regulation, less-efficient products would almost certainly be available in the marketplace. Ceiling fan purchasers would have a wider array of choices among more efficient

and less efficient fans. Perhaps they would also have a wider array of choices across other product attributes. The efficiency standard limits their range of choice, and in essence forces them to choose between a more efficient fan and an alternative product (for example, an air conditioner), or to purchase no product at all.

From a neoclassical perspective, when the option of a less efficient ceiling fan is removed from a consumer's choice set, an unambiguous cost to the consumer occurs, unless the consumer is acting with less than full information. There may well be other benefits associated with a product ban aside from energy savings (for example, benefits to the environment), but banning products that consumers would otherwise purchase is not a benefit *to them* in regulatory impact analysis.<sup>30</sup>

This is not the perspective of the DOE, however. The DOE estimates that the combined benefits of reducing carbon dioxide and nitrous oxide emissions are valued at \$4.2 billion (at a 3 percent discount rate), which means that 80 percent of the claimed benefits of the regulation are related to "operating cost savings" to consumers—the value of reduced energy costs to consumers. If the regulation is evaluated solely on the basis of the environmental benefits,<sup>31</sup> it fails a cost-benefit test according to the DOE's own numbers (see table 2 and figure 2).

Implicit in the DOE's decision to include these operating cost savings as benefits is the assumption that a market failure must be present, which the agency is correcting with the regulation. If no market failure were present, then it would not be possible for the agency to improve the allocation of resources and produce corresponding benefits for the public. Two possible examples of market failures that could be present are the possibility that consumers are acting with less than perfect information when purchasing ceiling fans or that there is a behavioral market failure present. The DOE offers a few justifications for the regulation, including:

The economics literature provides a wide-ranging discussion of how consumers trade off upfront costs and energy savings in the absence of government intervention. Much of this literature attempts to explain why consumers appear to un-

dervalue energy efficiency improvements. There is evidence that consumers undervalue future energy savings (or appear to do so) as a result of (1) a lack of information; (2) a lack of sufficient salience of the long-term or aggregate benefits; (3) a lack of sufficient savings to warrant delaying or altering purchases; (4) excessive focus on the short term, in the form of inconsistent weighting of future energy cost savings relative to available returns on other investments; (5) computational or other difficulties associated with the evaluation of relevant tradeoffs; and (6) a divergence in incentives (for example, between renters and owners, or builders and purchasers).<sup>32</sup>

Notably, the DOE offers no citations to support these assertions. There is a kitchen-sink-like quality to the list, as if the DOE was searching for any justification it could find for the regulation. Even if the DOE had provided more direct evidence, the relevant literature—as discussed earlier—is somewhat ambivalent about the extent of a market failure in markets for energy-consuming goods. Furthermore, it’s not clear that the academic literature has much, if anything, to say about the ceiling fan market in particular. Many consumers of ceiling fans are also businesses, which one would generally expect to be quite sophisticated in their purchasing decisions.


**Table 2. Benefits of Ceiling Fan Regulation**

Benefit	Billions of 2015 dollars
Consumer operating cost savings	16.5
Carbon dioxide reduction	3.8
Nitrous oxide reduction	0.4

Note: CO<sub>2</sub> reduction is calculated using the estimated mean social cost of carbon. All estimates calculated using a 3% discount rate.

Source: US Department of Energy, “Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment: Ceiling Fans,” 2016, 1-2.

**Figure 2. Benefits of Ceiling Fan Regulation**



Source: US Department of Energy, “Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment: Ceiling Fans,” 2016, 1-2.

Leaving aside the question of whether a market failure is present in the ceiling fan market, the operating cost savings calculations are based on a host of assumptions, some or all of which could turn out to be incorrect. For example, the DOE estimated the average lifetime of ceiling fans to be 13.8 years,<sup>33</sup> and assumed that a representative high-speed, small-diameter fan would operate 12 hours per day in active mode (with a range of about 6–18 hours), and that a large-diameter fan would also operate 12 hours per day.<sup>34</sup> The DOE also relied on predictions of the path of future energy prices, something that is notoriously difficult.

While it could be argued that the DOE made a good-faith effort to estimate these factors, and backed up its estimates with data that were available, it nonetheless seems likely one could find alternative assumptions and alternative data to justify a shorter lifespan for these products and less intensive use.<sup>35</sup> This could easily change the cost-benefit calculus.

Even assuming the DOE’s assumptions in these areas are correct, the agency still estimated that almost 30 percent of consumers of standard ceiling fans would experience net costs,<sup>36</sup> and that the rule would



result in an average installed cost of standard fans of \$124.95,<sup>37</sup> which is no small sum. The DOE's assumptions about operating cost savings are also highly dependent on the several discount rates used in analysis. The agency estimates national benefits and costs using social discount rates, assumed to be 3 or 7 percent, while individual or household-specific discount rates are used to calculate payback periods—that is, the amount of time it takes consumers to recoup the up-front costs of more-energy-efficient devices.

The DOE is right that for social, or aggregate, purposes, the relevant discount rate is the social discount rate, not an individual's personal discount rate. Here, DOE analysts are being careful to distinguish impacts from an individual perspective, such as payback periods, from impacts from a social perspective, such as cumulative benefits and costs to society. It remains an open question whether the social discount rates used in the analysis are correct (an issue that is beyond the scope of this chapter). However, there are reasons to believe that the DOE's assessment of individual-level discount rates may be off the mark. The DOE calculated household-specific discount rates in the ceiling fan analysis by calculating a weighted average of various debt and equity instruments available to households across different income groups. These weighted average rates vary from 3.49 percent (for the highest-income earners) to 5.08 percent (for the second-lowest-income group), with an average rate of 4.43 percent.<sup>38</sup> The DOE also calculated the weighted average cost of capital discount rates for businesses in the office, retail, lodging, and food-service industries. These rates ranged from 4.9 percent to 6.0 percent, with a mean rate of 5.0 percent.<sup>39</sup> Tables 3 and 4 display the various discount rates used by the DOE in its payback period analysis, broken down across income and business sectors.

The DOE's approach, which is to calculate weighted averages of various debt and equity interest rates available to households and businesses, where weights are based on the share of each source of financing for each group, is defensible in that, according to standard economic theory, a rational agent's own personal discount rate—that is, the rate at which the agent trades present for future consumption—should equal the market interest rate the agent faces, in equilibrium.<sup>40</sup>

**Table 3. Department of Energy Average Real Effective Discount Rate, by Income Group**

Percentile of income	Discount rate (%)
1st to 20th	4.88
21st to 40th	5.08
41st to 60th	4.67
61st to 80th	3.95
81st to 90th	3.68
91st to 100th	3.49
Overall average	4.43

Source: US Department of Energy, "Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment: Ceiling Fans," 2016, 8–23, 8–27.

**Table 4. Department of Energy Real Weighted Average Cost of Capital for Various Sectors**

Sector	Mean Discount rate (%)
Office	5.1%
Retail	5.0%
Lodging	6.0%
Food service	4.9%
Other	5.0%
Average discount rate	5.0%

Source: US Department of Energy, "Technical Support Document: Energy Efficiency Program for Consumer Products and Commercial and Industrial Equipment: Ceiling Fans," 2016, 8-29.

In this sense, the DOE's basic approach—of identifying market rates consumers and businesses face and using these as a proxy for the consumer's or business's personal discount rate—is not unreasonable.<sup>41</sup> A key question is whether the particular rates and their corresponding weights that are used are actually representative of the market conditions consumers and businesses face.

Many consumers face interest rates far higher than those calculated by the DOE, especially those who are credit constrained. For instance, the real interest rates on credit cards the DOE uses range from 9.87 to 11.95 percent,<sup>42</sup> which seem low given that consumers routinely face nominal rates closer to 20 percent in the marketplace. Low-income earners are also likely to have higher discount rates than high-income earners,<sup>43</sup> which is evident in table 3 and suggests that these rules may have important distributional effects. Many consumers don't even have credit. For example, the Consumer Financial Protection Bureau reported in 2015 that 26 million Americans, or roughly 1 in 10 adults, have no credit history.<sup>44</sup> The DOE acknowledges that credit constraints are a potential problem, but the agency doesn't seem to recognize that for some individuals the relevant interest rate is infinity—they can't borrow at any rate.

While the DOE's broad approach to calculating consumer and business discount rates follows a certain logic that is (at least on the face of it) reasonable, from a neoclassical perspective, all this analysis of pay-back periods may turn out to be unnecessary. If one accepts revealed preference as a principle of economic analysis, then whatever discount rate the consumer uses must be the correct discount rate. Only he or she knows the market conditions and interest rates relevant to his or her personal situation. The DOE's approach would make sense if it can be demonstrated that consumers and businesses have some informational problem or constraint that prevents them from calculating these rates effectively for themselves, such that the DOE can calculate their cost of capital better than they can. Absent such evidence, DOE's pay-back period analysis rests on shaky foundations.

More concretely, if revealed preference holds, then those consumers who opt to buy a particular device, and subsequently cannot buy

it because the device is banned by the DOE, suffer an unambiguous cost from the regulation. No payback period analysis need even be done, since such consumers are never paid back on net for their loss: the payback period is infinite. This is not to say that a regulation isn't worth promulgating. All policies create winners and losers, and the consumers that would have opted to purchase a product that is banned by regulators are simply some of the losers from the policy in question. This is not how the DOE markets its regulations, however. Instead, the agency routinely claims that consumers benefit from its appliance and equipment standards.<sup>45</sup>

There is also the possibility that consumers lose because product quality is impaired by energy-efficiency standards. One retrospective study identified a number of areas where appliance efficiency regulations tend to impose greater costs and lower cost savings on consumers than is typically estimated in the DOE's ex ante analyses. These areas include reduced product life and reliability, greater energy usage than anticipated, and additional operation and maintenance costs.<sup>46</sup>

One reason energy savings can be less than expected is that consumers sometimes respond to energy efficiency improvements by using a device more intensively, a phenomenon known as the *rebound effect*.<sup>47</sup> For example, drivers tend to drive more miles when the miles-per-gallon increases on their vehicle. The same general principle holds for appliances. While the DOE does acknowledge rebound effects and calculates some scenarios where a rebound effect occurs, it does not adjust operating-cost savings calculations due to a rebound effect in its core calculations for its ceiling fan rule.<sup>48</sup> In extreme cases, however, these unintended consequences have been known to fully offset the aim of reducing energy consumption.

Making a device more energy efficient can also complicate its mechanics, making it more prone to breaking—and shortening a product's life span adds to the likelihood that energy savings will not materialize. The DOE is supposed to account for the possibility of quality changes as a result of its regulations, according to various statutory mandates.<sup>49</sup> However, it is a change in relative quality, not just absolute quality, that should ideally be considered—that is, how quality has changed

or stayed the same relative to how quality would have improved in absence of a regulation. Such changes can be exceedingly hard to measure. All told, the uncertainty facing regulators is considerable when it comes to energy efficiency standards.

## **Is There an Economic Growth Rationale for These Regulations?**

Underlying revealed preference methods is the idea that consumers are generally most capable of deciding for themselves what is in their own interests. But what is in the individual's interest and what is in society-at-large's broader interest are not always the same thing. Market failures mean that individual and social interests often deviate, even under neoclassical assumptions.

However, standard portrayals of market failure, as illustrated in figure 1, often lack a temporal component. In other words, analysts tend to think of and explain market failures in a static context, without expressly considering a time element. They also fail to distinguish between social benefits and costs that come in the form of consumption from those that come in the form of investment. This is odd because the consumption choices of consumers today, even seemingly unimportant choices, can impact people in the future, not just those alive today. If I decide to purchase something as simple as an ice-cream cone for \$5, that decision can have a future impact if, had I not bought the ice cream, I would have invested some of the \$5 instead.<sup>50</sup>

Consider the hypothetical example of an energy efficiency regulation for an appliance that would result in savings of \$1 million by reducing energy use. Perhaps these savings come at the expense of a loss of consumer utility valued by present consumers at \$2 million—maybe because appliance functionality is impaired in some way. From the standpoint of the appliance's purchasers, this is clearly a bad deal. They value the lost product quality more than the financial savings from lower utility bills. They would be made worse off by this regulation, were it enacted.

But consider the same example from the standpoint of future consumers. If some fraction of the million dollars in financial savings were

invested and reinvested continually, it would grow into far more than \$1 million or even \$2 million in the future. The short-run reduction in consumer utility as a result of lower-performing appliances seems trivial from this perspective, because it is a temporary loss. Utility, after all, can't be invested in an account as money can. Meanwhile, the compounding gains stemming from increased investment could potentially have far-reaching longer-run consequences.

In theory, it seems that it could be worthwhile to override consumer choices in the present in order to increase investment that can benefit individuals in the future. This is not to say that present consumers are made better off by having had their choices overridden. Unlike the behavioral approach, this rationale for energy efficiency regulations does not claim that overriding consumer choices is a benefit—it is clearly a cost—but all regulations override choice to some extent, so the fact that consumer choices are overridden is not a sufficient reason to preclude regulation.

In fact, if future consumers could participate in present-day markets, one might expect they would be willing to pay consumers in the present to accept a lower-functioning appliance. They would likely be willing to pay up to the future value of whatever the investment benefit will be worth to them in the future. In our example, they might be willing to pay consumers \$2.1 million to accept some functionality impairment in their appliance. This is a small price for future consumers to pay if it leads to more investment in the economy, which could grow into huge sums owing to the power of compound interest.

If such compensation took place, then everyone would be made better off by the change without making anyone else worse off, a situation known as a Pareto improvement. Consumers today would be better off because, though they would have a worse device, they would also have some extra money that more than compensates them for it. Consumers in the future would be better off because of the increased investment that will have boosted economic growth and raised their incomes. Even without compensation, the winners would gain by more than the losers would lose, a situation which is known in economics as a Kaldor-Hicks improvement (or a potential Pareto improvement),

and is a principle that underlies cost-benefit analysis.

In this example, time acts as a transaction cost, similar to how traditional externalities, asymmetric information, and poor decision-making result from physical and mental transaction costs. Time creates a market failure because time prevents mutually beneficial exchanges from occurring. The market failure associated with time is unique in that usually with traditional neoclassical market failures, harmed third parties are alive to lobby on behalf of their own interests. This is not the case with future generations.

One potential problem with this logic is that almost any consumption expenditure would seem to impose an externality on people in the future. After all, some of the resources consumed might have been invested instead. So there needs to be some limiting principle that would prevent the present generation from having to save all its income for the sake of the future. In fact, there *is* such a limiting principle. Future generations would be unlikely to want to pay current consumers to invest any more of present income than corresponds with the consumption-maximizing “Golden Rule” rate of economic growth. This is the rate of economic growth that maximizes consumption across generations. Beyond this point, any additional returns to investment would be eaten up by the maintenance of depreciating capital, and so, on net, the rate of return on investment is negative for society.

An interesting aspect of these kinds of intertemporal externalities is that they can override traditional market failures. For example, let’s say that a polluting power plant reduces air quality in a particular city. The residents of the city, if they could organize, might be willing to pay the plant to reduce its emissions such that it would be profitable for the firm to do so. However, it is costly for residents to organize, owing to transaction costs, and so the exchange doesn’t take place. Such an exchange, if it took place, would increase social welfare within the current time period, because both the residents and the power plant would be better off. However, there is still the future to consider. If the exchange simply increased consumer utility—say health—but at the expense of capital accumulation and economic growth, people in the future might well be willing to compensate present citizens to accept

more pollution. On balance, social welfare might be improved if present citizens accepted more risk, in exchange for faster economic growth.

This example highlights how the efficient solution can change depending on whether time is considered in the analysis. If one focuses only on the present moment, it might appear that the efficient solution is for the power plant to pollute far less. But if one takes a perspective that accounts for the future as well, it may well be that the efficient outcome is for the power plant to produce more energy, and by extension, to increase present pollution somewhat. The irony, of course, is that regulators and academics seem to suffer from their own version of present bias, at least when it comes to the theory of market failure that is the basis for many regulatory interventions.

With energy efficiency regulations, the growth rationale for regulations is a moot point if financial savings never materialize. But if the energy savings are indeed real, any policy that on balance increases investment may actually increase social welfare, at least so long as the economy is operating below the consumption-maximizing Golden Rule rate of economic growth. This logic also extends to other behavioral interventions—such as default opt-in retirement accounts—that potentially increase savings and investment.

## **Conclusion and Policy Reform**

The theory underlying behavioral market failures is not so different from the theory that explains market failures of other kinds; however, the knowledge problem facing regulators is likely to be harder to overcome with behavioral market failures. Moreover, the traditional theory of market failure is itself incomplete, because it doesn't fully consider how benefits and costs accrue to individuals over the course of time.

Even the growth rationale for energy efficiency regulations should have caveats attached to it. Almost any regulation that increases capital formation on balance could be justified on these grounds. And while there are sound economic reasons to believe that the market underperforms in this regard (most economic growth models predict that society will not achieve the consumption-maximizing rate of growth), it is far from obvious whether the best way to promote capital formation is



through energy efficiency regulations, or through regulations at all. After all, if subpar growth is really such a big problem, why not encourage more investment in the marketplace, either by private firms or perhaps by the government itself—for example, through tax policy or the creation of a sovereign wealth fund? Energy efficiency regulations are unlikely to be the best method available to boost growth, or even to be a particularly good method.

All told, a dose of humility is likely in order. It may well be that the best option—even when behavioral market failures are present—is for policymakers to defer to the decisions of consumers. Identifying bias in the real world is likely to be extremely difficult, and the knowledge problem associated with proving the existence of behavioral market failures is severe. Even choice-preserving interventions, such as nudges, could well backfire. If a regulator switches a default option for retirement plans, how does the regulator know which default corresponds with employees' "true" preferences? Most likely some will be made better off and some worse off. On balance, the change could well be welfare-reducing, but how can the analyst ever know for sure?

Behavioral economists often argue that their interventions make consumers "better off as judged by themselves." At the end of the day, however, the knowledge constraints facing regulators make the behavioral rationale for promoting energy efficiency less than fully convincing. Fashionable new ideas and theories in academia are being used to justify regulatory interventions that have been around for decades.<sup>51</sup> Upon closer inspection, however, these new rationales end up looking a lot like the old ones: the same old paternalism that has existed throughout the ages.