NBER WORKING PAPER SERIES

HYSTERESIS AND THE WELFARE EFFECT OF CORRECTIVE POLICIES: THEORY AND EVIDENCE FROM AN ENERGY-SAVING PROGRAM

Francisco Costa
François Gerard

Working Paper 24608
http://www.nber.org/papers/w24608

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2018

We would like to thank four anonymous referees, Michael Best, Gharad Bryan, Robin Burgess, Lucas Davis, Marina de Mello, Ben Faber, Claudio Ferraz, Fred Finan, Greg Fischer, Meredith Fowlie, Jason Garred, Gustavo Gonzaga, Michael Greenstone, Ryan Kellog, Gilat Levy, Jamie McCasland, Guy Michaels, Edward Miguel, Mushfiq Mobarak, Marcelo Moreira, Gerard Padro i Miquel, Joana Naritomi, Emmanuel Saez, Edson Severnini, Reinaldo Souza, Anna Spurlock, Eric Verhoogen, Catherine Wolfram, and seminar participants at ASU, the Behavioural Environmental Economics conference (Toulouse), Cergy-Pontoise, EconCon, the EEA meetings, FGV EPGE, FGV EESP, the Granthan Institute for Climate Change, IADB, IFS, IPEA, LBS, LSE, the NBER Public Economics Meetings, PacDev, PSE, Puc-Rio, UC Berkeley, University of Bristol, University of Cambridge, University of Chicago, UPF, UQAM, USP, UT Knoxville, UCSD, and the World Bank. Stephanie Dinoa, Louise Guillouet, Tiago Lazier, Gustavo Macedo, John Pease, and Beatriz Rache provided outstanding research assistance. We also thank ANEEL, PROCEL, LIGHT, FAME, and Whirlpool for sharing important data. Francois Gerard benefited from the support of an Excellence Scholarship of Wallonie-Bruxelles International. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Hysteresis and the Welfare Effect of Corrective Policies: Theory and Evidence from an Energy-Saving Program
Francisco Costa and François Gerard
NBER Working Paper No. 24608
May 2018
JEL No. D62,H23,Q50

ABSTRACT

A growing body of evidence documents that policies can affect household behaviors persistently, even if they are no longer in place. This paper studies the importance of such "hysteresis" – the failure of an effect to reverse itself as its underlying cause is reversed – for the welfare evaluation of corrective policies. First, we introduce hysteresis into the textbook framework used to derive canonical sufficient statistics formulas for the welfare effect of corrective policies. We then derive new formulas allowing for hysteresis. We show that, under certain conditions, the persistent effect of a short-run (i.e., temporary) policy becomes a new key statistic for evaluating the welfare effect of such a policy, and also of a long-run (i.e., permanent) version of a similar policy. Second, we estimate the persistent effect of a short-run policy, for which we argue that these conditions are met, in a policy-relevant context: residential electricity use in a developing country setting. We estimate that about half of the dramatic short-run reductions in residential electricity use induced by a 9-month-long policy that was imposed on millions of Brazilian households in 2001 persisted for at least 12 years after the policy ended. Finally, we combine our estimates with our framework to illustrate the implications that hysteresis can have for the welfare evaluation of corrective policies.

Francisco Costa
Getulio Vargas Foundation (FGV/EPGE)
Praia de Botafogo, 190 / 1100
Rio de Janeiro, RJ 22250-900
Brazil
francisco.costa@fgv.br

François Gerard
Department of Economics
Columbia University
1022 IAB
420 West 118th Street
New York, NY 10027
and NBER
fgerard@columbia.edu

A data appendix is available at http://www.nber.org/data-appendix/w24608

In the presence of an externality, a policy that changes the behavior of private economic agents may be socially beneficial. For instance, environmental damages from energy generation may justify a tax to reduce energy demand. A canonical result in economics is that such a policy must weigh the welfare gain from addressing the externality and the loss in private surplus due to the change in the behavior at stake. When evaluating this tradeoff, the literature typically assumes away the possibility that a policy may affect choices persistently, even if it is no longer in place. This paper studies the importance of such "hysteresis" – the failure of an effect to reverse itself as its underlying cause is reversed (Dixit, 1989) – for the welfare evaluation of corrective policies.

A growing body of evidence documents the persistent effects of short-run (i.e., temporary) policies on a range of individual or household behaviors (e.g., the adoption of new health or income-generating technologies, energy use, water consumption, smoking, commuting, and voting) and some of the mechanisms behind these effects (e.g., investments in physical capital, information acquisition, and habit formation).[1] This evidence suggests that hysteresis is a relevant phenomenon to consider for evaluating corrective policies. Several papers show that the persistence affects the *cost-effectiveness* of policies in their context. However, the literature lacks a general framework to evaluate the implication of hysteresis for the *welfare* effect of corrective policies.

This paper aims to address that limitation. First, we introduce the possibility of hysteresis in household behaviors into the textbook framework used to derive canonical sufficient statistics formulas for the welfare effect of corrective policies (e.g., for a Pigouvian tax). We then derive new formulas allowing for hysteresis. We show that, under certain conditions, the persistent effect of a short-run policy becomes a new key statistic for evaluating the welfare effect of such a policy, and also of a long-run (i.e., permanent) version of a similar policy. Second, we estimate the persistent effect of a short-run policy, for which we argue that these conditions are met, in a policy-relevant context: residential electricity use in a developing country setting. We find that about half of the dramatic short-run reductions in residential electricity use induced by a 9-month-long policy that was imposed on millions of Brazilian households in 2001 persisted for at least 12 years after the policy ended. Third, we combine our empirical estimates with our framework to illustrate the implications that such hysteresis can have for the welfare evaluation of corrective policies.

We begin by laying out a standard partial-equilibrium framework, in which a representative household must choose its level of a behavior of interest – we use electricity use as running example – in each of two periods. We introduce hysteresis by allowing choices in the first period to affect the marginal utility from electricity use in the second period in two ways. First, the household can make "active" investments (e.g., in physical capital or information acquisition) that may

[1] See, for instance, Reiss and White (2008), Charness and Gneezy (2009), Giné, Karlan and Zinman (2010), Ferraro and Price (2013), Allcott and Rogers (2014), Bryan, Chowdhury and Mobarak (2014), Dupas (2014), Acland and Levy (2015), Ito (2015), Fujiwara, Meng and Vogl (2016), Miller (2016), Brandon et al. (2017), Ito, Ida and Tanaka (2017) and Larcom, Rauch and Willems (2017).

affect its present and future utility from electricity use, as these investments may be persistent. Second, past consumption choices may "passively" affect the household's utility from electricity use, like in experienced-good or habit formation models. We motivate the need for corrective policies in this framework by assuming that electricity use generates negative externalities (e.g., environmental damages). We then derive sufficient statistics formulas for the welfare effect of two types of policies: a *short-run* policy, in which the government restricts electricity use in period 1 only, and a *long-run* policy, in which it restricts electricity use in both periods. We first show that we obtain existing formulas in the literature if we assume away the possibility of hysteresis. Next, we obtain new formulas allowing for hysteresis, and we highlight the bias from failing to take it into account when evaluating the welfare effect of corrective policies.

We first derive formulas under benchmark assumptions that electricity is priced at a constant marginal cost, that the restrictions are implemented through traditional instruments (i.e., quotas or taxes), and that the household fully internalizes the effect of its choices on its future utility. Assuming away hysteresis, the welfare effect of the *short-run* policy is the sum of a welfare gain from addressing the externality in period 1 and of a loss in consumer surplus in that period. Moreover, three statistics are sufficient for evaluating these effects: the change in electricity use, the marginal damage of electricity use, and the slope of the demand curve for electricity. Allowing for hysteresis, however, the household may also use less electricity in period 2 after the policy ended. Price theory implies that the slope of the demand curve in period 1 remains a sufficient statistic for the overall change in consumer surplus, if the policy was *expected to be short run*. In that case, the only additional effect to evaluate is the welfare gain from a possible correction of the externality in period 2. The persistent impact of the policy on electricity use and the marginal damage in period 2 are sufficient for evaluating this effect, as well as the bias from assuming away hysteresis. Our qualifier regarding policy expectations matters for active mechanisms of hysteresis. The household may make investments in period 1 that it would not make if it knew that the policy were to be temporary, and their cost is not captured by the demand curve in period 1. This highlights a limitation of the results from existing studies for our purpose. Studies estimating the persistent effect of a short-run policy rarely specify whether the policy was expected to last longer than it did.

The persistent effect of a short-run policy is also key for evaluating the welfare effect of a *long-run* policy. Assuming away hysteresis, one can evaluate its welfare effect by using the same approach used for the short-run policy but in both periods. In so doing, one would correctly account for the correction of the externality in both periods. However, one would overestimate the loss in consumer surplus in the presence of hysteresis. In that case, part of the reduction in electricity use in period 2 would take place in absence of any policy in that period because of the choices induced by the policy in period 1. Moreover, the associated loss in consumer surplus is already captured by the demand curve in period 1. Because of this, the slope of the demand curve in period 2 should only be used to evaluate the loss in consumer surplus due to the change in electricity use caused by the policy in that period, and not the overall change in electricity use in period 2. The former is the difference between the overall change in period 2 and the persistent change caused by the policy in period 1. This is why the persistent effect of a short-run policy is a key statistic for evaluating the welfare effect and the bias from assuming away hysteresis in the long-run case as well.

These results carry through if we relax our benchmark assumptions, but additional considerations arise.[2] First, if electricity is not priced at a constant marginal cost, the sign of the bias from assuming away hysteresis becomes ambiguous for the short-run policy, as the persistent reduction in electricity use also generates a loss in producer surplus in period 2. The persistent effect of the short-run policy remains a key statistic for evaluating the bias because it allows us to evaluate this new source of welfare loss as well. Second, the welfare implications of hysteresis are unchanged if a corrective policy includes social incentives (e.g., peer pressure), as long as any persistent effect of a short-run policy results from choices made in response to that policy. In contrast, the welfare implications can be very different if such a policy changed social norms persistently, creating incentives to behave in a certain way even after it ended. This qualifier is important because studies that estimate the persistent effect of social incentives rarely assess whether the incentives led to persistent changes in social norms. Third, the welfare effect of corrective policies may include additional terms related to the correction of an "internality" if the household does not fully internalize the future effects of its choices. The benchmark formulas still provide useful lower bounds for the welfare effect in cases where internality and externality issues are likely aligned.[3]

In sum, estimates of the persistent effect of a short-run policy can be informative for the welfare evaluation of short-run and long-run corrective policies. Yet, these estimates must come from a setting in which the short-run policy was expected to be short run, did not lead to persistent changes in social incentives, and helped address (rather than exacerbate) any potential internality issue.

In the second part of the paper, we estimate the long-run impact of a temporary energy-saving program for which we argue that these conditions are met. Our empirical setting is particularly well suited to study hysteresis. The policy, which was implemented during the "2001 Brazilian electricity crisis," led to the largest short-run reductions in household electricity use among temporary energy-saving programs around the world (Meier, 2005). Households may be especially likely to make persistent changes in their physical capital or their consumption habits when required to make drastic changes in behavior of this kind. Moreover, we can study the effect of the policy for more than 10 years after it ended. Our setting is also policy-relevant. Energy use is expected to remain a major source of greenhouse gas emissions in the future, and potential energy savings from

[2] Our results also carry through in models with heterogeneous households or uncertainty (see Section 1.6).

[3] E.g., if the household uses too much electricity from both a private and a social perspective.

residential users have attracted a lot of attention.[4] Yet, because of low price elasticities, the loss in private surplus from inducing large long-run changes in residential energy use may be considered too large despite sizable externalities, at least based on existing welfare formulas.

In the beginning of 2001, electricity generation capacity was severely reduced in some regions of Brazil. An exceptional and temporary drop in the streamflow level of the rivers that serve the hydroelectric power plants in these regions led to historically low water levels in their reservoirs. Brazil heavily relies on hydroelectric generation and low transmission capacity constrained transfers of electricity across regions. To prevent generalized blackouts, the government adopted a temporary energy-saving program in affected areas that was aimed at reducing residential electricity use by 20%. Customers were only given a three-week notice before the policy was implemented between June 2001 and February 2002. Households were assigned individual quotas and were subject to a series of incentives, including social incentives, to consume less than their quota.

We estimate the short- and long-run impacts of the energy-saving program using a difference-in-differences approach that compares distribution utilities subject to the policy to those that were exempt. We use data on average residential electricity use per customer from monthly administrative reports for every distribution utility between 1991 and 2014. We confirm that the policy had a large short-run impact (−23%). We then show that half of that impact persisted in the long run (−12%).[5] Consumption levels partially rebounded after the policy ended, but point estimates are stable from 2005 onward. This can be seen in the raw data in Figure 1. Our estimates are robust to using different specifications and to controlling for a series of confounders, such as changes in electricity prices, demographics, or income levels. Moreover, using synthetic control methods, we find negative long-run impacts for *every* distribution utility subject to the policy. As an order of magnitude, the total reduction in electricity use implied by our estimates from 2003 to 2014 (100.5 billion kWh) is almost twice as large as what the Brazilian government claims to have achieved through all its energy-saving policies over the same period (see procelinfo.com.br).

We complement our findings with evidence based on longitudinal monthly billing data from 2000 to 2005 for 3 million customers of one affected distribution utility. We show that the change in average electricity use at the level of the distribution utility came from large and persistent reductions by most customers throughout the distribution of consumption levels. We then use idiosyncratic variation in quotas across customers to estimate their causal effect on short-run and long-run electricity use levels. We argue that this variation is unlikely to lead to persistent differences in social incentives (or policy expectations) that affect the interpretation of our estimates for the welfare evaluation of corrective policies. For the typical household, the quota was set at 80%

[4]Most scenarios to mitigate climate change involve large reductions in electricity use by buildings (IPCC, 2014).

[5]Little work investigates the long-run impact of this energy-saving program on residential electricity use. Bardelin (2004) and Maurer, Pereira and Rosenblatt (2005) provide some descriptive evidence on short-run impacts with aggregate data. Pimenta, Notini and Maciel (2009) use time-series techniques.

of its average consumption between May and July 2000. However, for households that moved into their housing unit after May 2000, the quota was set at 80% of the average consumption *in their first 3 billing months*. Because of strong seasonality in electricity use, this rule created large differences in quotas for otherwise comparable customers entirely stemming from differences in moving dates. Using this variation as an instrument for the movers' quotas, we find that 36%-56% of the short-run effect of the quotas during the crisis persisted in the long run. Finally, surveys suggest that a main mechanism of hysteresis was a change in consumption habits.

Last, we return to the framework to illustrate the implications of our estimates for the welfare evaluation of corrective policies. Back-of-the-envelope calculations suggest that the Brazilian government would have been willing to pay US$874 million for the long-run electricity savings of the policy. This figure illustrates the additional welfare gain for the short-run policy. Because electricity is unlikely priced at its marginal cost, the persistent effect of the policy may also imply a long-run loss in producer surplus. Accounting for hysteresis would in fact decrease the welfare effect of the short-run policy with a mark-up above the marginal cost of only 6.5%. In contrast, in the case of a long-run policy, one would unambiguously underestimate the welfare effect by assuming away hysteresis. Using our own estimate of the price elasticity of demand, we calculate that the bias would reach $7.75 billion for a policy lasting until 2014, what corresponds to 7.18% of the discounted sum of all residential electricity bills over the period. This value would be even larger if we were to account for any internality issue, as households were unlikely to be using too little electricity before the crisis or to be harming themselves by using less electricity afterward.

This paper contributes to the literature on the economics of corrective policies, which dates back to Pigou (1920). Recent papers have extended the canonical formulas for the welfare effect of corrective policies in several ways (e.g., Allcott and Taubinsky, 2015; Allcott and Kessler, forthcoming), but have not considered the role of hysteresis systematically.[6] Although derived in the context of our application, our results generalize to other behaviors of interest as well. Our results also generalize beyond a specific model of hysteresis, as is usual with sufficient statistics formulas.

This paper also contributes to the growing literature that estimates the long-run effect of short-run policies in various contexts (see the papers cited in footnote 1). Existing papers usually evaluate the policy implications of their findings in terms of cost-effectiveness. We provide a framework to evaluate the policy implications of hysteresis in terms of welfare, and we apply it to a specific empirical setting. In so doing, we show that estimates of hysteresis can inform the welfare effect of both short-run and long-run policies, and that the welfare implications of hysteresis can be important and are not necessarily positive. We also show that first-order considerations for

[6] In recent theoretical work, Acemoglu et al. (2012) argue that temporary policies promoting greener technologies may have persistent effects on the supply side through directed technical change, and Hintermann and Lange (2013) investigate the dynamic optimal regulation of experience goods that generate environmental externalities.

cost-effectiveness become secondary under a welfare criterion, but that other considerations arise. For instance, from a cost-effectiveness perspective, we must quantify the respective role of mechanisms that do or do not involve monetary costs (e.g., physical investments vs. behavioral changes). Although researchers can sometimes provide evidence of specific mechanisms (e.g., Ito, Ida and Tanaka, 2017; Brandon et al., 2017), they typically cannot reject that other mechanisms are also at play. We show that this issue becomes secondary once we evaluate the implications of hysteresis from a welfare perspective. In contrast, understanding *why* agents make persistent adjustments becomes important, as implications for welfare can differ in important ways if the persistent effect of a short-run policy is affected, e.g., by policy expectations or by persistent changes in social norms.

Additionally, our findings contribute to the empirical literature on corrective policies in the context of energy use. Starting with Reiss and White (2008), several papers have provided evidence of hysteresis in residential electricity use. However, they typically study persistent effects for a shorter period of time (e.g., Ito, 2015), or the effects that they estimate dissipate over time (e.g., Allcott and Rogers, 2014). Our findings that a 9-month-long policy can have sizable and stable impacts for more than a decade shed new light on the importance of hysteresis in this context. Moreover, our setting is particularly interesting. Most of the growth in energy use, including in residential electricity use, comes from the developing world (IPCC, 2014). Yet, the energy-saving potential of households in developing countries that are poorer, own fewer appliances, are more credit constrained, and use less energy in general, is largely unknown.[7] Finally, our conclusion that one may overestimate the social cost of inducing long-run changes in energy use by assuming away hysteresis parallels recent empirical findings indicating that one may overestimate the damages from climate change by failing to take adaptation into account (e.g., Barreca et al., 2016).

The rest of the paper proceeds as follows. Section 1 lays out our conceptual framework. Section 2 presents the empirical setting, and Section 3 presents our main findings. Section 4 provides complementary evidence using microdata, and Section 5 discusses mechanisms. Section 6 illustrates the possible welfare implications of our estimates using our framework. Section 7 concludes.

1 Conceptual framework

We begin by showing how hysteresis affects the welfare evaluation of corrective policies. We build on the textbook price-theory framework used to derive canonical results in the economics of corrective policies (e.g., optimal Pigouvian taxation). This allows us to highlight the role of hysteresis in a familiar setting. Moreover, the results we derive using this framework have the advantage that they generalize beyond a specific model, as is usual with sufficient statistics formulas.

[7] For instance, Davis, Fuchs and Gertler (2014) find that an appliance replacement program in Mexico did not generate the expected energy savings because of households' behavioral responses to the new appliances.

We start by laying out a benchmark framework and by discussing mechanisms of hysteresis it applies to. We then derive sufficient statistics formulas for the welfare effect of corrective policies. In so doing, we show how estimates of the persistent effect of short-run policies can inform the welfare effect of short-run (i.e., temporary) and long-run (i.e., permanent) policies, whether the government has a short-run goal (e.g., curbing energy demand in the face of a temporary shortage) or a long-run goal (e.g., reducing energy demand to mitigate climate change). Next, we consider extensions to the benchmark framework. In each case, we obtain existing formulas for the welfare effect of corrective policies when assuming away hysteresis. We then obtain new formulas allowing for hysteresis, and we highlight the bias from failing to take hysteresis into account.

In the interest of space, we only present the intuition for some of our results here, but Appendix B includes an extended version of this section. Although we use household electricity use as running example throughout the section, note that the theory applies equally well to other household behaviors for which the possibility of hysteresis may be relevant. Moreover, our results are not specific to settings in which households are on the demand side of the market.[8]

1.1 General setup and mechanisms

We make some assumptions to focus on the implications of hysteresis for the welfare evaluation of corrective policies. First, we consider the problem of a representative household deciding how much electricity to use in each of two periods. A two-period setting is sufficient for our purpose, and we allow for heterogenous households in an extension. Second, we assume that electricity use generates externalities in order to motivate the need for corrective policies. The government has a short-run (resp. long-run) goal when the externality applies to the first period only (resp. to both periods). Third, we adopt a partial-equilibrium setup to focus on the household decision, and we use quasi-linear utility functions to assume away income effects and redistributive concerns, as is standard in the literature (e.g., Allcott and Taubinsky, 2015). Fourth, we assume that electricity is priced at a constant marginal cost, such that we can illustrate the role of hysteresis using simple consumer surplus concepts. We introduce a producer surplus in an extension. Fifth, we model the household as fully rational and forward looking; doing so is necessary for slopes of the demand curves to be sufficient statistics for changes in consumer surplus. We later relax these assumptions.

We introduce hysteresis in this rather standard setup by allowing changes in electricity use due to a corrective policy in period 1 to affect the household's propensity to use electricity in period 2, or the energy (in-)efficiency of the mapping from its electricity use to the utility it derives from it.

As we show below, it is useful to distinguish between two categories of mechanisms of hysteresis. First, changes in electricity use in period 1 may involve *active* investments that modify the

[8]E.g., for an application on charitable donation, we would model the household on the supply side of the market and we would obtain the same results (appropriately replacing references to "demand curves" by "supply curves").

household's propensity to use electricity persistently. In our context, this category encompasses physical investments (e.g., in the energy efficiency of electrical appliances) and information acquisition (e.g., learning about energy-efficient behaviors). In both cases, the household's propensity to use electricity may change persistently because investment costs are not easily reverted and the assets created do not depreciate immediately. Second, changes in electricity use in period 1 may have persistent effects on the household's propensity to use electricity because of its history of consumption. Like in experience-good models, the household may learn about the costs and benefits of adopting new behaviors by experimenting with different levels of electricity use in period 1, acquiring new information or developing new habits *passively* (e.g., Bryan, Chowdhury and Mobarak, 2014; Dupas, 2014). Habit formation models, in which the marginal utility of consumption depends on past consumption levels, belong to this category (Becker and Murphy, 1988).[9]

Our framework captures these two categories in a "reduced-form" fashion such that our results are not tied to modeling a particular mechanism. In fact, our results generalize to any such mechanism, as long as slopes of the demand curves remain sufficient statistics for changes in the household welfare from having to deviate from its privately optimal choices. This condition is a natural starting point to focus on the implications of hysteresis, as it is the same condition used to derive the canonical results for optimal Pigouvian taxation. We show that, under this condition, we do not need to know which specific mechanisms cause the persistent effect of a short-run policy for the welfare evaluation of corrective policies, as long as we can quantify the degree of hysteresis.[10]

1.2 Formal presentation of the benchmark framework

We now present the benchmark framework formally. The household derives utility from the services provided by electricity use x_t, with unit price p_t, and from a numeraire c_t, given its income y_t in periods $t = 1,2$. The per-period utility from electricity services is $v_t(x_t, s_t)$ and is strictly increasing and concave in x_t. This function captures the mapping from electricity use to utility, i.e., electricity use provides services that create utility. It implicitly involves usage and electrical appliance choices, e.g., consuming more electricity implies adding appliances or new ways to use these appliances.[11] s_t is the household's "propensity to consume" electricity. It captures the energy (in-)efficiency of the mapping from electricity use to utility, i.e., usage and appliance choices can be more or less energy-efficient. Without loss of generality, high values of s_t imply a high propensity to consume (e.g., energy-inefficient appliances or habits, limited knowledge of energy-efficient behaviors). We make two key assumptions about this variable that are also made in, e.g., Becker and

[9] The optimal consumption path in rational habit formation models may have multiple steady states, and a corrective policy that pushes the household far enough from its prior steady state may lead to hysteresis.

[10] We implicitly assume that there are no (differential) pre-existing distortions on the household's various margins of adjustment. Otherwise, it would be necessary to know which mechanisms cause hysteresis.

[11] It is common in the literature to not specify all the choices involved in such a mapping (see, e.g., Ito, 2015).

Murphy (1988). First, we assume $\frac{\partial v_t}{\partial s_t} \leq 0$: the household derives less utility from a given level of electricity use if its stock of appliances, its habits, or its knowledge are less energy-efficient (this assumption essentially defines s_t). Second, we assume complementarity, $\frac{\partial^2 v_t}{\partial x_t \partial s_t} > 0$: the higher the propensity to consume (e.g., the more electricity the household needs to obtain services), the higher the marginal utility from electricity use (e.g., the greater the benefits from more electricity).

We introduce hysteresis by allowing the propensity to consume to be a function $s_t(s_{t-1}, x_{t-1}, I_t)$. The household can make investments, I_t, to reduce its concurrent propensity to consume, $\frac{\partial s_t}{\partial I_t} \leq 0$, at strictly increasing and convex costs $\kappa_t(I_t)$, which may or may not be monetary (e.g., price of new appliances or time devoted to learning).[12] These investments can affect the future propensity to consume because we assume $\frac{\partial s_t}{\partial s_{t-1}} \geq 0$ (this allows for depreciation). Consumption choices can also affect the household's future propensity to consume directly, $\frac{\partial s_t}{\partial x_{t-1}} \geq 0$. The costs of reducing its propensity to consume through this channel are already captured by assuming $\frac{\partial v_t}{\partial x_t} > 0$. Changes in the future propensity to consume lead to hysteresis because of the complementarity assumption.

The household then solves the following problem:

$$\max_{x_1, x_2, I_1, I_2} V = U_1 + \beta U_2 = y_1 - p_1 x_1 - \kappa_1(I_1) + v_1(x_1, s_1) + \beta\left[y_2 - p_2 x_2 - \kappa_2(I_2) + v_2(x_2, s_2)\right] \tag{1}$$
$$\text{s.t. } s_t = s_t(s_{t-1}, x_{t-1}, I_t) \text{ for } t = 1,2 \text{ and } \{s_0, x_0\} \text{ given,}$$

where V is the household "lifetime" utility and β accounts for discounting and possible differences in the relative length of the two periods. The first-order conditions for this problem are:

$$\frac{\partial v_1(x_1, s_1)}{\partial x_1} + \beta \frac{\partial v_2(x_2, s_2)}{\partial s_2} \frac{\partial s_2}{\partial x_1} = p_1 \quad ; \quad \frac{\partial v_2(x_2, s_2)}{\partial x_2} = p_2 \tag{2}$$
$$\left[\frac{\partial v_1(x_1, s_1)}{\partial s_1} + \beta \frac{\partial v_2(x_2, s_2)}{\partial s_2} \frac{\partial s_2}{\partial s_1}\right] \frac{\partial s_1}{\partial I_1} - \kappa_1'(I_1) = 0 \quad ; \quad \frac{\partial v_2(x_2, s_2)}{\partial s_2} \frac{\partial s_2}{\partial I_2} - \kappa_2'(I_2) = 0 \tag{3}$$

The first-order conditions highlight the costs and benefits of marginal changes in electricity use and investments, respectively. They imply that the household uses more electricity if its concurrent propensity to consume is higher ($\frac{\partial^2 v_t}{\partial x_t \partial s_t} > 0$); that it uses less electricity in period 1 if its future propensity to consume depends on past choices ($\frac{\partial s_2}{\partial x_1} > 0$); and that it invests more in period 1 if investments are persistent ($\frac{\partial s_1}{\partial I_1} < 0$ and $\frac{\partial s_2}{\partial s_1} > 0$). These first-order conditions and baseline electricity prices (p_{10} and p_{20}) determine baseline electricity use and investment levels (x_{10}, x_{20}, I_{10}, and I_{20}).

[12] The two representations are equivalent with quasi-linear utility. That we model investments as a continuous variable is inconsequential; we only need the aggregate demand curve for electricity to be continuous. Note that our model is consistent with Dubin and McFadden (1984), where investing in a more energy-efficient technology increases the marginal productivity of electricity. In their model, this can have ambiguous predictions for the household's marginal utility from electricity use. We assume that it decreases the marginal utility of electricity use ($\frac{\partial^2 v_t}{\partial x_t \partial s_t} \frac{\partial s_t}{\partial I_t} \leq 0$) in order to introduce the possibility of hysteresis in the direction that is consistent with our empirical evidence. However, our sufficient statistics formulas would still hold if the persistent effect was instead an increase in electricity use.

1.3 Social welfare and corrective policies

We motivate the need for corrective policies in this framework by assuming that social welfare, W, differs from private welfare or the household's lifetime utility, V. In particular, there are externalities from electricity use, $E_t(x_t)$, which we assume to be negative without loss of generality, such that welfare can be written $W = V - E_1(x_1) - \beta E_2(x_2)$, with $E_t(x_t)$ positive and increasing in x_t.

We then evaluate the welfare effect of two types of policies. First, we consider a short-run policy, in which the government restricts electricity use to $\overline{x_1} < x_{10}$ in period 1 only. The welfare effect of this policy is $\Delta W^{SR} = W(\overline{x_1}, x_2(\overline{x_1}), I_1(\overline{x_1}), I_2(\overline{x_1})) - W(x_{10}, x_{20}, I_{10}, I_{20})$, where SR stands for "short run." We denote by $x_2(\overline{x_1})$, $I_1(\overline{x_1})$, and $I_2(\overline{x_1})$, privately optimal levels of these variables given the short-run policy and the price p_{20}. Second, we consider a long-run policy, in which the government restricts electricity use to $\overline{x_1} < x_{10}$ in period 1 and to $\overline{x_2} < x_{20}$ in period 2. The welfare effect of this policy is $\Delta W^{LR} = W(\overline{x_1}, \overline{x_2}, I_1(\overline{x_1}, \overline{x_2}), I_2(\overline{x_1}, \overline{x_2})) - W(x_{10}, x_{20}, I_{10}, I_{20})$, where LR stands for "long run." We denote by $I_1(\overline{x_1}, \overline{x_2})$ and $I_2(\overline{x_1}, \overline{x_2})$, privately optimal levels of these variables given the long-run policy. Finally, we assume that these restrictions on electricity use are implemented through traditional instruments, i.e., through quotas or taxes (we consider a policy that includes social incentives in an extension). The welfare gain of both policies then arises from the correction of the externalities and the welfare loss from the reduction in the household private welfare from having to deviate from its privately optimal choices regarding electricity use.

Our framework allows the government to have short-run or long-run goals for these policies. Specifically, it has a short-run (resp. long-run) goal if the slope of the externality function – the marginal damage $MD_t(x_t)$ – is assumed to be nil in period 2 (resp. to be positive in both periods).

1.4 Welfare effect assuming away hysteresis

We now show how a researcher *who assumes away hysteresis* would evaluate the welfare effect of short-run and long-run corrective policies using a standard price-theory approach. In this case, the household's propensity to consume is assumed to not depend on past consumption choices, and investments fully depreciate from one period to the next; that is, we have $s_t(s_0, I_t)$ for a given s_0.

1.4.1 Short-run policy

The researcher can recover the welfare effect of the short-run policy by tracing it along the path from x_{10} to $\overline{x_1}$, letting I_1 adjust endogenously, as variables in period 2 are unaffected by assumption:

$$\Delta W_{NoH}^{SR} = \int_{x_{10}}^{\overline{x_1}} \frac{dU_1(x_1, I_1)}{dx_1} dx_1 - \int_{x_{10}}^{\overline{x_1}} \frac{dE_1(x_1)}{dx_1} dx_1 = \int_{x_{10}}^{\overline{x_1}} [p_{1,NoH}(x_1) - p_{10}]\, dx_1 - \int_{x_{10}}^{\overline{x_1}} MD_1(x_1) dx_1 \quad (4)$$

where NoH stands for "no hysteresis." Equation (4) is the textbook formula for the welfare effect of a corrective policy. The welfare gain from correcting the externality corresponds to the area under the marginal damage function $MD_1(x_1)$ between the consumption levels with and without the policy in period 1. The welfare loss from reducing electricity use below the privately optimal level is captured by a Harberger triangle. It corresponds to the area under the demand curve $p_{1,NoH}(x_1)$ and above the baseline price p_{10} between the consumption levels with and without the policy in period 1, or the triangle $A_1B_1C_1$ in Figure 2a for the case of a linear demand curve.

The key insight of price theory is that the slope of the demand curve is a sufficient statistic for the change in the household private welfare (i.e., the consumer surplus) because it accounts for changes in its utility due to changes in x_1 and to changes in other choices variables following changes in x_1 (e.g., I_1). As a result, it does not matter *how* the household reduced its electricity use. The household optimally chooses among all its margins of adjustments, and any loss in its utility is reflected in the demand curve. Therefore, if p_{10} and x_{10} are known, two empirical objects are sufficient for evaluating the welfare effect of the short-run policy. The slope of the demand curve $p_{1,NoH}(x_1)$ can be estimated using exogenous changes in electricity prices in period 1. The marginal damage function $MD_1(x_1)$ can be estimated using the various approaches in the literature.

1.4.2 Long-run policy

Since choices in the two periods are assumed to be independent, the researcher can recover the welfare effect of the long-run policy by separately tracing it along the paths from x_{10} to $\overline{x_1}$ and from x_{20} to $\overline{x_2}$, letting I_1 and I_2 adjust endogenously:

$$\begin{aligned}\Delta W^{LR}_{NoH} &= \int_{x_{10}}^{\overline{x_1}} \frac{dU_1(x_1,I_1)}{dx_1}dx_1 + \beta \int_{x_{20}}^{\overline{x_2}} \frac{dU_2(x_2,I_2)}{dx_2}dx_2 - \int_{x_{10}}^{\overline{x_1}} \frac{dE_1(x_1)}{dx_1}dx_1 - \beta \int_{x_{20}}^{\overline{x_2}} \frac{dE_2(x_2)}{dx_2}dx_2 \qquad (5)\\ &= \int_{x_{10}}^{\overline{x_1}} [p_{1,NoH}(x_1) - p_{10}]\,dx_1 + \beta \int_{x_{20}}^{\overline{x_2}} [p_{2,NoH}(x_2) - p_{20}]\,dx_2 - \int_{x_{10}}^{\overline{x_1}} MD_1(x_1)dx_1 - \beta \int_{x_{20}}^{\overline{x_2}} MD_2(x_2)dx_2.\end{aligned}$$

Equation (5) is a two-period version of the textbook formula in equation (4). The total loss in private welfare corresponds to the (discounted) sum of Harberger triangles in both periods, or the areas $A_1B_1C_1$ and $A_2B_2C_2$ in Figures 2a and 2b, respectively. If p_{20} and x_{20} are also known, equation (5) shows that three additional objects are needed to evaluate the welfare effect of the long-run policy: the discount rate β, the marginal damage function $MD_2(x_2)$, and the slope of the demand curve $p_{2,NoH}(x_2)$, which could be estimated using exogenous changes in electricity prices once in period 2. Finally, it is straightforward from equations (4) and (5) that a government with a short-run (resp. long-run) goal will want to implement a short-run (resp. long-run) policy.

1.5 Welfare effect allowing for hysteresis

We now show how a researcher *who allows for hysteresis* would evaluate the welfare effect of the two policies using price theory. In this case, the propensity to consume takes the form $s_t\left(s_{t-1}, x_{t-1}, I_t\right)$.

1.5.1 Short-run policy

The researcher can still recover the welfare effect of the short-run policy by tracing it along the path from x_{10} to $\overline{x_1}$. However, one must now recognize that x_2 and I_2 may also adjust endogenously:[13]

$$\begin{aligned}\Delta W^{SR} &= \int_{x_{10}}^{\overline{x_1}} \frac{dV(x_1,x_2,I_1,I_2)}{dx_1}dx_1 - \int_{x_{10}}^{\overline{x_1}} \frac{dE_1(x_1)}{dx_1}dx_1 - \beta \int_{x_{10}}^{\overline{x_1}} \frac{dE_2(x_2)}{dx_2}\frac{dx_2}{dx_1}dx_1 \\ &= \int_{x_{10}}^{\overline{x_1}} \left[p_1(x_1) - p_{10}\right]dx_1 - \int_{x_{10}}^{\overline{x_1}} MD_1(x_1)dx_1 - \beta \int_{x_{20}}^{x_2(\overline{x_1})} MD_2(x_2)dx_2. \end{aligned} \tag{6}$$

The same insight of price theory implies that the slope of the demand curve in period 1 remains a sufficient statistic for the change in private welfare due to the short-run policy, including from changes in x_2 and I_2. Moreover, the slope of the demand curve $p_1(x_1)$ could be estimated in the same way as a researcher would estimate the slope of the demand curve $p_{1,NoH}$ when assuming away hysteresis. The only novelty in equation (6) is that the researcher must now consider a potential effect on the externality in period 2 because of hysteresis in electricity use. As Figure 2b illustrates, the demand curve in period 2, $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$, may have shifted inward given the change in the household's propensity to consume due to its choices in period 1. This shift leads to a privately optimal level of electricity use in period 2, $x_2(\overline{x_1})$, that is lower at price p_{20}.

The bias from assuming away hysteresis when evaluating the welfare effect of a short-run policy therefore comes from neglecting the possible correction of the externality in period 2:

$$Bias^{SR} = \Delta W^{SR}_{NoH} - \Delta W^{SR} = \beta \int_{x_{20}}^{x_2(\overline{x_1})} MD_2(x_2)dx_2 \leq 0. \tag{7}$$

One would underestimate the welfare gain of the policy, unless the government has only a short-run goal ($MD_2(x_2) = 0$, $\forall x_2 \leq x_{20}$). The magnitude of the bias depends on the marginal damage function in the long run, $MD_2(x_2)$, and on the long-run effect of the short-run policy $x_{20} - x_2(\overline{x_1})$.

The long-run effect of the policy is thus the new key empirical object to estimate. More precisely, it corresponds to the long-run effect of a policy that was *expected to be short run*. This qualifier is important because studies that estimate the persistent effect of a short-run policy rarely specify whether the policy was expected to last longer than it did. Policy expectations matter because the household may make investments in period 1 that it would not make if it knew that the

[13] Ito, Ida and Tanaka (2017) find persistent effects for 3 months after a policy ended, and use a formula similar to the one derived here to evaluate the welfare implications of that specific short-run policy (see their Web Appendix).

policy were to be only temporary. These "extra" investments, which may contribute to the persistent effect of the policy in period 2, are ex-post suboptimal for the household and imply losses in private welfare that are not accounted for by the demand curve in period 1.

Formally, we can illustrate this point by considering a policy in which the government first announces electricity use restrictions $\overline{x_1} < x_{10}$ and $\overline{x_2} < x_{20}$, but then cancels the policy at the start of period 2. Let $x_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2}))$ and $I_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2}))$ be the household's choices in period 2 given its choices in period 1 when planning for a long-run policy, $\overline{x_1}$ and $I_1(\overline{x_1}, \overline{x_2})$. We can recover the welfare effect of such a policy that was "not expected to be short run" (*NotExp*) by tracing it along the path from x_{10} to $\overline{x_1}$. However, we must then also account for the extra investments that the household may make in period 1 in anticipation of the period-2 policy:

$$\Delta W^{SR}_{NotExp} = \int_{x_{10}}^{\overline{x_1}} [p_1(x_1) - p_{10}]\, dx_1 + \int_{I_1(\overline{x_1})}^{I_1(\overline{x_1}, \overline{x_2})} \frac{dV(\overline{x_1}, x_2, I_1, I_2)}{dI_1} dI_1 - \int_{x_{10}}^{\overline{x_1}} MD_1(x_1) dx_1 - \beta \int_{x_{20}}^{x_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2}))} MD_2(x_2) dx_2$$

Compared with equation (6), two differences are clear. First, the correction of the externality in period 2 includes the reduction in electricity use due to the extra investments. As we have $x_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2})) < x_2(\overline{x_1})$, this implies that one would overestimate the welfare gain of an actual (i.e., known to be) short-run policy by evaluating equation (6) using estimates of the long-run effect of a policy that was not expected to be short run. Second, the new integral above captures the loss in private welfare from these extra investments. One would thus overestimate the welfare effect of a policy that was not expected to be short run by failing to take these investments into account. In fact, the sign of the bias from assuming away hysteresis is ambiguous for such a policy: $Bias^{SR}_{NotExp} = \beta \int_{x_{20}}^{x_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2}))} MD_2(x_2) dx_2 - \int_{I_1(\overline{x_1})}^{I_1(\overline{x_1}, \overline{x_2})} \frac{dV(\overline{x_1}, x_2, I_1 I_2)}{dI_1} dI_1$. It depends on the sizes of the welfare gain from the correction of the externality in period 2 and of the welfare loss from the extra investments.[14] This latter term is challenging to estimate because one would need to measure the difference in investment choices with and without the wrong policy expectations for all possible types of investments separately, as well as the slope of the demand curve for all these investments. The persistent effect of a policy that was not expected to be short run is thus much less informative.

This discussion highlights the importance of studying *why* agents make persistent changes following short-run policies when evaluating the welfare implications of hysteresis. Note also that this issue only arises when hysteresis is due to costly investments. This highlights the usefulness of distinguishing between the two categories of mechanisms of hysteresis in our framework.

[14] The welfare effect of this policy could also be larger than that of an actual short-run policy if the gain from the additional correction of the externality outweighs the loss from the extra investments. However, by the *targeting principle* (Sandmo, 1975), a policy in which the government announces and actually implements electricity use restrictions $\overline{x_1}$ and $x_2(\overline{x_1}, I_1(\overline{x_1}, \overline{x_2}))$ would lead to the same changes in electricity use for smaller losses in private welfare.

1.5.2 Long-run policy

Next, we show that the persistent effect of a short-run policy also informs the size of the bias from assuming away hysteresis in the case of a long-run policy. The researcher can recover the welfare effect of such a policy by tracing it along any path from (x_{10}, x_{20}) to $(\overline{x_1}, \overline{x_2})$. Yet, with hysteresis, one cannot consider the paths in the two periods separately. Following the path that first changes x_1 to $\overline{x_1}$ and then x_2 to $\overline{x_2}$ given $\overline{x_1}$, which is natural in our intertemporal setting, one obtains:

$$\begin{aligned}
\Delta W^{LR} &= \int_{x_{10}}^{\overline{x_1}} \frac{dV(x_1,x_2,I_1,I_2)}{dx_1} dx_1 + \int_{x_2(\overline{x_1})}^{\overline{x_2}} \frac{dV(\overline{x_1},x_2,I_1,I_2)}{dx_2} dx_2 - \int_{x_{10}}^{\overline{x_1}} \frac{dE_1(x_1)}{dx_1} dx_1 - \beta \int_{x_{20}}^{\overline{x_2}} \frac{dE_2(x_2)}{dx_2} dx_2 \\
&= \int_{x_{10}}^{\overline{x_1}} [p_1(x_1) - p_{10}]\, dx_1 + \beta \int_{x_2(\overline{x_1})}^{\overline{x_2}} [p_2(x_2|\overline{x_1}, I_1(\overline{x_1})) - p_{20}]\, dx_2 + \int_{I_1(\overline{x_1})}^{I_1(\overline{x_1},\overline{x_2})} \frac{dV(\overline{x_1},\overline{x_2},I_1,I_2)}{dI_1} dI_1 \\
&\quad - \int_{x_{10}}^{\overline{x_1}} MD_1(x_1) dx_1 - \beta \int_{x_{20}}^{\overline{x_2}} MD_2(x_2) dx_2. \qquad (8)
\end{aligned}$$

The welfare gain from correcting the externalities is identical to that in equation (5), so the difference, which we divide into three parts in equation (8), comes from the loss in private welfare. The first integral captures the loss from the period-1 policy, which can still be measured using the demand curve $p_1(x_1)$. This term accounts for changes in private welfare due to changes in x_1 and in the other choice variables to $x_2(\overline{x_1})$, $I_1(\overline{x_1})$, and $I_2(\overline{x_1})$. The second integral captures the loss from the period-2 policy if the household had not anticipated the policy while in period 1. It is the Harberger triangle under the demand curve $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$ and above the baseline price p_{20} between the consumption levels with and without the period-2 policy. This term accounts for changes in private welfare from the additional changes in x_2 to $\overline{x_2}$ and in I_2 to $I_2(\overline{x_1}, \overline{x_2}, I_1(\overline{x_1}))$. Finally, the household could have anticipated the period-2 policy and chosen a different level of investment in period 1. The impact on private welfare is captured by the third integral and corresponds to the change in private welfare from moving the period-1 investment levels from $I_1(\overline{x_1})$ to $I_1(\overline{x_1}, \overline{x_2})$, holding electricity use constant.[15] By revealed preferences, this term is positive because the only reason to make such investments is to mitigate the loss in private welfare due to the period-2 policy.

Equation (8) highlights three sources of bias from assuming away hysteresis:

$$\begin{aligned}
Bias^{LR} &= DWL^{LR}_{NoH} - DWL^{LR} \qquad (9) \\
&= \beta \left[\int_{x_{20}}^{\overline{x_2}} [p_{2,NoH}(x_2) - p_{20}]\, dx_2 - \int_{x_2(\overline{x_1})}^{\overline{x_2}} [p_2(x_2|\overline{x_1}, I_1(\overline{x_1})) - p_{20}]\, dx_2 \right] - \int_{I_1(\overline{x_1})}^{I_1(\overline{x_1},\overline{x_2})} \frac{dV(\overline{x_1},\overline{x_2},I_1,I_2)}{dI_1} dI_1.
\end{aligned}$$

First, part of the reduction in electricity use in period 2, from x_{20} to $x_2(\overline{x_1})$, is due to the period-1 policy. The associated change in private welfare is already accounted for in the period-1 demand curve and should not be double counted. This is why the second integral in equation (8) is only

[15] The period-2 investment levels endogenously change from $I_2(\overline{x_1}, \overline{x_2}, I_1(\overline{x_1}))$ to $I_2(\overline{x_1}, \overline{x_2})$.

taken from $x_2(\overline{x_1})$ to $\overline{x_2}$. In Figure 2b, it corresponds to the triangle $A_2D_2E_2$ under the demand curve $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$ with base $x_2(\overline{x_1}) - \overline{x_2}$. The first source of bias from assuming away hysteresis corresponds to the trapezoid $D_2B_2C_2E_2$, the difference with the triangle $A_2B_2C_2$ under the demand curve $p_{2,NoH}(x_2)$ with base $x_{20} - \overline{x_2}$. The persistent effect of the short-run policy, $x_{20} - x_2(\overline{x_1})$, which again corresponds to the persistent effect of a policy that was expected to be short run,[16] is thus a key empirical object for evaluating this source of bias.

Second, the change in the period-2 propensity to consume due to the period-1 policy may have changed the slope of the period-2 demand curve. Whether this leads to a bias depends on the variation used to estimate the slopes of the demand curves $p_{2,NoH}(x_2)$ and $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$. A researcher who allows for hysteresis would estimate the slope of $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$ using exogenous price variation in period 2 following a short-run policy. A researcher who assumes away hysteresis could estimate the slope of $p_{2,NoH}(x_2)$ using the same variation. There would be no second source of bias in this case. However, one may also wrongly believe that the slopes of the demand curves in periods 1 and 2 are identical and use an estimate of the slope of $p_{1,NoH}(x_1)$ for $p_{2,NoH}(x_2)$ as well. The welfare loss would be smaller (resp. larger) in that case, and the bias larger (resp. smaller), if the demand curve became more elastic (resp. inelastic) following the short-run policy, e.g., if households started paying more attention to their electricity use, like in Jessoe and Rapson (2014).

Third, a researcher who assumes away hysteresis would fail to recognize that the household could make different investment choices in period 1 in order to mitigate its overall loss in private welfare. This highlights once again the usefulness of distinguishing between the two mechanisms of hysteresis in our framework. The third source of bias is challenging to estimate. First, one would need to measure the difference in period-1 investments, for all possible types of investments, in the case of a long-run policy and in the case of a short-run policy. Second, one would need to evaluate the change in private welfare resulting from such changes in investments, given electricity use levels $\overline{x_1}$ and $\overline{x_2}$.[17] However, given the sign of this third source of bias, we can abstract from it and still draw meaningful welfare conclusions. In so doing, we evaluate informative bounds: a lower bound for the overall bias and the welfare effect, and an upper bound for the loss in private welfare.

We focus on the first source of bias in our application, which is a natural first step. When we evaluate the size of the bias, we thus assume that a researcher assuming away hysteresis would use the same estimate of the slope of the period-2 demand curve as a researcher allowing for hysteresis. We also acknowledge that our estimates of the welfare effect and the bias may constitute a lower

[16]The second integral in equation (8) may be taken over an interval that is too small if one uses instead an estimate of the persistent effect of a policy that was not expected to be short run. We can also write formulas in terms of such an estimate (see Appendix B). But, the bounds they provide would overestimate rather than underestimate the welfare effect and the size of the bias, and upper bounds are less informative given the limited evidence that such a bias exists.

[17]This last step is conceptually challenging: it is unclear how to estimate demand curves for investments, holding fixed both present and future electricity use levels.

bound if hysteresis is partly due to active investments in the household's propensity to consume.

Finally, because the correction of the externalities is identical in equations (5) and (8), the bias from assuming away hysteresis is the same whether the government has a short-run or a long-run goal. It is also straightforward from equations (6) and (8) that, even with hysteresis, a government with a short-run (resp. long-run) goal will want to implement a short-run (resp. long-run) policy.

1.6 Extensions

We now present extensions to the benchmark framework and show that the long-run effect of a short-run policy remains informative for the welfare evaluation of corrective policies. We only discuss the main intuitions for our results here and present all the derivations in Appendix B.

1.6.1 Adding heterogeneity or uncertainty

First, we can add heterogeneity or uncertainty in our framework and the main results carry through.

A. Heterogeneity. We add heterogeneity by assuming that J households face a problem similar to that of the representative household, indexing all variables and functions by $j = 1, ..., J$. We then consider policies aimed at aggregate changes in electricity use (i.e., externalities depend on aggregate levels of electricity use). Specifically, we consider policies implementing these restrictions efficiently, for example, with tradable quotas or taxes. Such policies equalize the marginal costs of reducing electricity use across households, therefore minimizing the loss in private welfare with heterogenous households. Moreover, changes in private welfare can be traced along aggregate demand curves for electricity when electricity use is allocated efficiently. Otherwise, allocative inefficiencies arise, complicating the welfare evaluation whether or not one allows for hysteresis.

The formulas for the welfare effect of these policies are simply aggregate versions of equations (4), (5), (6), and (8). The new key empirical object to estimate for evaluating the bias from neglecting the correction of the externality in period 2 for the short-run policy, and the first source of bias from assuming away hysteresis for the long-run policy, becomes the aggregate long-run effect of the short-run policy. The only real difference comes from new gains from anticipating the period-2 policy. Mechanisms of hysteresis provide gains from trade across households with heterogeneity in investment costs or in the effect of past choices on their propensity to consume. These mechanisms can further mitigate the loss in private welfare from the long-run policy. One would thus further underestimate the overall bias by abstracting from the third source of bias.

B. Uncertainty. We add uncertainty by considering a "rational inattention" version (e.g., Sallee, 2014) of the benchmark framework. We assume that the household is uncertain about the cost of investing in its propensity to consume, which is multiplied by a random variable $\theta \in [\underline{\theta}, \overline{\theta}]$ with known distribution. Before period 1, the household can choose how much attention effort

$e \in [0,1]$ to devote to learning the value of θ, with strictly increasing and convex costs $\psi(e)$. With probability e, uncertainty is then resolved for both periods; otherwise, the household makes choices under uncertainty. This allows uncertainty resolution to be a mechanism of hysteresis: the household may adjust its attention effort in response to a corrective policy in period 1, and changes in acquired information may affect the household's choices persistently. In contrast, a researcher who assumes away hysteresis must assume that any resolution of uncertainty in period 1 has no impact in period 2 or that the uncertainty must be independently resolved in each period.

The welfare effect of corrective policies must now account for attention costs.[18] However, the slopes of the demand curves for electricity remain sufficient statistics for the loss in private welfare. The demand curves simply factor in losses in private welfare from changes in an additional variable, e, following changes in electricity use. The welfare formulas are thus similar to those in equations (4), (5), (6), and (8), and the long-run effect of the short-run policy remains the new key empirical object to estimate for evaluating the bias from assuming away hysteresis for both policies. The only real difference from the benchmark framework relates again to the third source of bias for the long-run policy. Indeed, the household could have also adjusted its attention effort in anticipation of the period-2 policy. Doing so would further mitigate the loss in its private welfare.

1.6.2 Adding a producer surplus

Second, we relax the assumption that electricity is priced at constant marginal costs.[19] The household's problem is unchanged, but we now assume that electricity is provided by a representative price-taking firm that maximizes profits in each period, with strictly increasing and convex costs $c_1(x_1)$ and $c_2(x_2)$. In that case, baseline prices and quantities are such that the first-order conditions of both the household and the firm are satisfied. We also consider an alternative model that better fits the regulatory framework in our application: marginal costs are constant, but the firm is allowed to charge a fixed mark-up, for instance, to cover some fixed costs.

Assuming away hysteresis, the formulas for the welfare effect of the two corrective policies are still textbook formulas. However, because welfare includes the firm's profits, the loss in private welfare includes not only a loss in consumer surplus but also a loss in producer surplus. One must thus estimate additional empirical object in order to evaluate the loss in producer surplus: the slopes of the supply curves (first model) or the size of the mark-ups (alternative model).

Allowing for hysteresis, the welfare effect of the short-run policy includes the same additional welfare gain as in the benchmark framework from the correction of the externality in period 2. However, hysteresis in electricity use also implies an additional loss in producer surplus from the unrealized profits on the electricity that the firm does not sell in period 2 because of the policy. As

[18] Alternative policies reducing the cost of information acquisition may be useful in this model.

[19] Hysteresis still takes place on the demand side, and not on the firm side like in, e.g., Acemoglu et al. (2012).

a result, the sign of the bias from assuming away hysteresis is ambiguous for the short-run policy, and the bias in fact changes signs if the government has only a short-run goal. Yet, the additional sources of welfare gain and welfare loss are both functions of the long-run effect of the short-run policy. It thus remains the new key empirical object to estimate for evaluating this bias. The additional source of welfare loss also depends on the slope of the supply and demand curves (first model; the price adjusts in period 2) or on the size of the mark-up (alternative model) in period 2.

For the long-run policy, the correction of the externality and the loss in producer surplus are the same whether or not hysteresis is assumed away. Therefore, the bias from assuming away hysteresis comes from the same three sources of bias as in the benchmark framework, the first of which is still a function of the long-run effect of the short-run policy.[20]

1.6.3 Considering policies that include "social incentives"

Third, we show that the welfare implications of hysteresis are unchanged if corrective policies include social incentives (Allcott and Mullainathan, 2010), as long as any persistent effect of a short-run policy still results from choices made in response to that policy. In contrast, we show that the implications of hysteresis can be very different if a policy persistently changes social norms, creating incentives for households to behave in a certain way even after the policy ended.

We begin by assuming that the corrective policies implement the electricity use restrictions at least partly through a social incentive component, which we model in a way similar to that in, e.g., Allcott and Kessler (forthcoming). The policies create a social price τ_t and social norm N_t in the periods in which they are in place. The norm captures an ideal electricity use level and the social price captures a social marginal incentive to reduce electricity use, e.g., a feeling of shame for every unit consumed above the norm or a feeling of pride for every unit reduced below the norm. The welfare effect of the corrective policies then includes an additional term that corresponds to the direct effect of the social incentives on the household's utility (besides any change in behavior). This term appears in the welfare formulas whether or not one assumes away hysteresis, so the welfare implications of hysteresis remain the same as those in the benchmark framework.

Let's assume instead that the household's propensity to consume does not depend on past choices – we have $s_t(s_0, I_t)$ – but that the short-run policy creates a persistent social incentive in period 2 (N_2, τ_2), even though the policy itself ended. This is now the reason the household uses less electricity in period 2. Assuming away hysteresis is equivalent to assuming away the creation of the persistent social incentive in this model, so the formula for the welfare effect of the short-run policy is the same as that in the benchmark framework. However, the formula differs in two ways when we allow for hysteresis. The additional welfare gain from the correction of the externality

[20] A complication arises in the first model because of the price adjustment in period 2 following the short-run policy, such that one must also estimate the slope of the period-2 supply curve to evaluate the first source of bias in that model.

in period 2 is unchanged, but there is an additional loss in private welfare from the reduction in electricity use in period 2. It corresponds to a Harberger triangle under the period-2 demand curve and above the baseline price between consumption levels with and without the short-run policy, as if the reduction in electricity use was due to a Pigouvian tax. The long-run effect of the short-run policy is thus a key empirical object to estimate for evaluating these two sources of bias. There is also the direct effect of the persistent social incentive, which can be positive or negative, as the household may get positive or negative social utility from its electricity use in period 2. In practice, quantifying this latter component is challenging; doing so requires eliciting the household's willingness to pay for the policy (Allcott and Kessler, forthcoming). In any case, the sign of the bias from assuming away hysteresis is now ambiguous. This highlights again the importance of studying *why* agents make persistent changes for evaluating the welfare implications of hysteresis.

1.6.4 Allowing for present bias or myopic beliefs

The slopes of the demand curves are sufficient statistics for evaluating changes in private welfare in the benchmark framework because the household maximizes the same private welfare that is included in the social welfare function. This rests on relatively standard, but nevertheless strong, assumptions. Thus, we explore natural deviations from these assumptions, that is, models in which the household is present-biased or has myopic beliefs about the persistent effect of its choices.

A. Present bias. We assume that the household makes decisions in period 1 by maximizing its perceived lifetime utility $\widetilde{V}$, in which it discounts the future by $\alpha\beta$ instead of β. $\alpha \in [0,1)$ captures a bias toward the present. A present-biased household consumes more electricity in period 1 and invests less in reducing its propensity to consume. Conditional on period-1 choices, however, the household makes privately optimal choices in period 2, that it planned to make while in period 1.[21]

In this model, the government could evaluate private welfare based on the household's "experienced" utility V (as suggested by, e.g., Farhi and Gabaix, 2017), in which it discounts the future without a bias like in equation (1), or based on the household's "decision" utility from its period-1 perspective $\widetilde{V}$. Assuming away hysteresis, the present bias does not affect choices, so the formulas for the welfare effect of the corrective policies are similar to those in the benchmark framework. The terms related to period 2 are simply scaled down by α with the decision utility criterion.

Allowing for hysteresis, the household's choices are affected by its present bias. With the experienced utility criterion, it is useful to rewrite welfare as $W = \widetilde{V} + (V - \widetilde{V}) - E_1(x_1) - \beta E_2(x_2)$, where the difference between decision and experienced utility $(V - \widetilde{V})$ captures an *internality*. The household consumes too much electricity and underinvests in period 1 for its own private welfare.

[21] In a model with three periods, the household consumes too much electricity and underinvests in its propensity to consume in period 2, even from its own period-1 perspective. As a result, the corrective policies further increase welfare under both welfare criteria by correcting an internality in period 2 in the right direction (see below).

As a result, if the perceived losses in private welfare from the corrective policies are still captured by the same terms as in equations (6) and (8), these terms now misestimate the actual loss in private welfare because the policies also correct this internality. Formally, there is an additional term in the formulas for both short-run and long-run policies: an *average marginal bias* (Allcott and Taubinsky, 2015). It corresponds to the area between the period-1 demand curve and the function that captures the household actual utility from x_1, given the levels of the other choice variables, over the interval from $\overline{x_1}$ to x_{10}. In practice, it is difficult to estimate this function,[22] but the average marginal bias often can be signed. For instance, it is positive in our case because the policies correct the internality in the right direction, thus increasing the welfare effect and the bias.

With the decision utility criterion, the social welfare function uses the same private welfare as maximized by the household, $\widetilde{V}$, so the formulas are similar to those in equations (6) and (8). The terms in period 2 and the bias from assuming away hysteresis are simply scaled down by α.

B. Myopic beliefs. We assume that the household maximizes its perceived lifetime utility in period 1, assuming away any effect of its choices on its future propensity to consume, but makes decisions in period 2 based on its actual propensity to consume. This myopic household consumes more electricity and invests less in reducing its propensity to consume in period 1. It also fails to predict its propensity to consume and thus its choices in period 2.

The government could again evaluate private welfare based on the household's experienced utility, V, or on its decision utility from its period-1 perspective, $\widetilde{V}$. The latter case could make sense in some cases, e.g., if the impact of its choices on its future utility is harmful for the household, like with addictions. Assuming away hysteresis, the myopic beliefs have no effect on the household choices, and the decision and experienced utility criteria coincide (the household also assumes away hysteresis). So the formulas for the welfare effect of the corrective policies are the same as in equations (4) and (5). Allowing for hysteresis, the third source of bias for the long-run policy in the benchmark framework is absent because the myopic household does not anticipate the period-2 policy. However, the formulas differ between the two welfare criteria in other respects.

With the experienced utility criterion, the corrective policies generate an extra source of welfare gain, like in the previous model, because of the correction of an internality in period 1. It can be decomposed into two components. First, the household misperceives its gains from reducing electricity use in period 1, holding fixed its plans regarding period-2 choices (because $\frac{\partial v_t}{\partial s_t} \leq 0$). Second, it does not anticipate its gains from reducing its electricity use once in period 2 below the level it had planned to use (because $\frac{\partial^2 v_t}{\partial x_t \partial s_t} > 0$). The first component remains difficult to estimate. The second component corresponds to the area between the demand curve and the baseline price

[22] It is not captured by the demand curve that we observe if the household is "debiased." The household may adjust the level of other choice variables for any given level of x_1 in that case, and these changes will be reflected into such a demand curve. Allcott and Taubinsky (2015) avoid this complication because their model has a single choice variable.

in period 2, and between the expected and actual levels of electricity use following the short-run policy. The latter component is thus a function of the long-run effect of the short-run policy if we assume that the household expected to consume what it would have consumed absent the policy.[23]

With the decision utility criterion, there is also an extra term in the welfare formulas, but it is a source of *welfare loss* because of an internality in period 2 in this case. First, under this criterion, the household *overestimates* its gains from reducing electricity use in period 2 following the short-run policy. This first component corresponds to the area between the household actual utility for x_2 and the baseline price in period 2 between the actual levels of electricity use with and without the short-run policy. It thus depends on the long-run effect of the short-run policy. Second, the perceived loss in private welfare from the period-2 policy, which is measured under the period-2 demand curve, underestimates the actual loss in private welfare. The two components are difficult to evaluate because they depend on the household's utility for x_2 ignoring the change in its propensity to consume. Yet, these components reduce the welfare effect of the corrective policies, such that the sign of the bias from assuming away hysteresis becomes ambiguous.

In sum, the bias from assuming away hysteresis is positive and may even be larger in these models using the household's experienced utility, which would make sense in many applications. However, using the decision utility in some prior period could make sense with, e.g., addictions.

This section has shown how hysteresis modifies the textbook formulas for the welfare effect of corrective policies. It has also shown that estimates of the persistent effect of a short-run policy can be informative for the welfare evaluation of both short-run and long-run corrective policies if the short-run policy was expected to be short run, if it did not lead to persistent changes in social incentives, and if it did not persistently change behaviors in a way that would harm the household. In the remainder of this paper, we estimate the long-run effect of a specific short-run policy, for which we argue that these conditions are met, and then combine our estimates with the theory to illustrate the welfare implications of hysteresis in a policy-relevant context.

2 Background and data for the empirical application

We study the persistent effect of a 9-month-long energy-saving program implemented in response to an exceptional shortage of electricity supply in specific areas of Brazil in 2001. Our empirical setting is particularly well suited for studying hysteresis. The policy was known to be temporary. It was a bold policy that led to the largest short-run reductions in household electricity use among temporary energy-saving programs around the world (Meier, 2005). Households may be more

[23] In other words, the household is myopic regarding the effect of changes in period-1 choices away from baseline choice levels, but it is not wrong about its future propensity to consume at baseline. In a model with three periods, the period-2 demand curve would further underestimate its gains from reducing electricity use in period 2.

likely to make persistent changes in their physical capital or their habits when required to make drastic changes in behavior. Moreover, we can estimate the long-run effect of the policy for up to 12 years after it ended. It would be difficult for a controlled experiment to share these features. One limitation is that the policy did not rely on a single fixed instrument to achieve its short-run effect.[24] Nevertheless, estimating the causal effect of the policy as a whole and of one of its key component (consumption quotas) allows us to illustrate the welfare implications of hysteresis.

2.1 Background information

We start with an overview of the electricity distribution system in Brazil, the causes of the 2001 electricity crisis, and the energy-saving program. Appendices C to E provide more details.

2.1.1 Electricity distribution in Brazil

The major national electricity system in Brazil is divided into four subsystems: North (6.5% of total load in 2000), Northeast (14.5%), Southeast/Midwest (62%), and South (17%). Almost all households had access to the electricity grid in the South and the Southeast/Midwest in 2000. In contrast, the electricity grid was not fully developed in the two other subsystems at the time, but it developed rapidly in the following years, thanks to strong support from the federal government (e.g., the program "Luz Para Todos"). The Brazilian electricity system almost exclusively relies on hydrological resources. In 2000, hydropower was responsible for 81% of the production capacity and 94% of the electricity generated in the country (ONS, 2011). More than 60 local monopolies (distribution utilities) distribute electricity to end consumers, and housing units are typically metered and billed separately every month. Finally, electricity theft – illegal connections to the grid – is a serious concern in Brazil, amounting to 15% of the total load for some distribution utilities.

Electricity prices are regulated by a federal agency (Agência Nacional de Energia Elétrica, ANEEL). The main residential tariff is a flat unit price per kilowatt hour (kWh). An alternative tariff for low-income and small consumers offers percentage discounts on the main tariff depending on the quantity consumed. Price changes typically modify the main tariff and therefore imply a proportional change in every marginal price. The regulatory framework is a price-cap mechanism. Every 4 or 5 years, prices are *revised* to guarantee the economic viability of distribution utilities. However, demand risk falls on the utilities between revision years. Yearly price *adjustments* only factor in changes in non-manageable costs, such as energy costs (ANEEL, 2005).[25]

[24] Any policy that aims at changes in behavior on such a large scale would likely include a set of policy instruments.

[25] The price-cap mechanism encourages distribution utilities to address electricity theft. Price revisions and adjustments occur at different times for different utilities. In June 2001, the main tariff was US$.23/kWh in Rio de Janeiro, for instance. Marginal prices in the alternative tariff were US$.08 (up to 30 kWh), US$.14 (up to 100 kWh), US$.20 (up to 140 kWh), and US$.23 (above 140 kWh). Monetary values are in US$ from 2012 throughout the paper.

2.1.2 History of the 2001 electricity crisis

The 2001 electricity crisis was due to supply factors. It was caused by a historically low streamflow in the rivers that serve hydropower plants in specific areas of the country, combined with infrastructure constraints on generation and transmission capacity. Figure 3a displays the evolution of the hydro-reservoirs' water level in the Southeast/Midwest and in the South. We focus on the two largest subsystems because of their similar development stage at the time (see next section). Water levels follow a seasonal pattern in the Southeast/Midwest, with heavy rain replenishing the reservoirs downstream during the austral summer. Levels were low in the two subsystems by 2000. In March 2001, in the Southeast/Midwest, levels reached their lowest point in 40 years (for the season) because of exceptionally low summer rainfall, and thus low flow into the reservoirs (see Figure 3b). In contrast, generous rain dissipated any risk of shortages in the South as the flow into the reservoirs was relatively high (see Figure 3c). The South could not transfer its excess supply to the other subsystems because of limited transmission capacity across subsystems.[26] Moreover, while growth in demand had never outpaced growth in projected demand in the years prior to 2001 (see Appendix C), it outpaced growth in generation capacity. This was a nationwide issue, and experts later concluded that the South would have faced a similar crisis if it had experienced hydrological conditions similar to those faced in the Southeast/Midwest (Kelman, 2001).

By April, it was clear that electricity use had to decrease to avoid generalized blackouts. The government announced that an incentive-based energy-saving program would start in June, although details remained unclear (*O Globo*, April 23, 2001).[27] Rolling blackouts were part of a plan B that never became necessary. The government program started on June 4, and, from the start, it was expected to last until February 2002 (the end of the next rainy season; *Veja*, July 19, 2001). The goal of the program was to reduce electricity use by 20% in the Southeast/Midwest. The program also applied in the North and the Northeast, which faced a similar situation, but not in the South. Mation and Ferraz (2011) provide evidence that the crisis, the energy-saving program, and its differential implementation across subsystems were unanticipated.[28] As expected, the crisis ended in February 2002, but according to a specialized periodical, "people were giving signals that they learned how to avoid wasting electricity" (*Energia Elétrica*, March 15, 2002).

Minimum consumption levels are also charged, and local taxes increase what customers eventually pay.

[26] After 2001, the government heavily invested in transmission capacity to reduce the risk of electricity shortages.

[27] This was despite a first set of national policies in early April. Among these measures were the giveaway of efficient lightbulbs in low-income neighborhoods, a 15% reduction in electricity consumption in federal public buildings, the import of energy from Argentina, and the construction of new thermoelectric facilities (*Veja*, April 5, 2001).

[28] For instance, President Cardoso's approval rates dropped differentially in areas subject to the energy-saving program after the program's announcement. Figure 1 also shows that households in the different subsystems did not anticipate the energy-saving program by changing their electricity use differentially prior to June 2001.

2.1.3 Incentives of the energy-saving program

The energy-saving program included individually assigned quotas and a set of incentives for residential customers to consume less than their quota. The rules were frequently repeated in the media and on electricity bills, but they were relatively complicated, nonlinear, and changed more than once.[29] Therefore, even though we show that households responded to their quotas, the relevant contribution of each of the policy's components to its overall impact is unclear.

Every customer was assigned a *quota* at the start of the crisis. The typical quota was equal to 80% of a *baseline* corresponding to the customer's average consumption from May to July 2000.[30] Quotas for customers who moved into their housing unit after May 2000 were set at 80% of the average consumption in their first 3 billing months. Because of the strong seasonality in electricity use, this assignment rule generated variation in quotas entirely because of differences in moving dates. We exploit this variation to provide evidence of the short- and long-run effects of the quotas.

The incentives included "carrots and sticks." Customers exceeding their quota were charged a *fine* per kilowatt hour consumed above 200 kWh equal to 50% of the usual tariff; it was increased to 200% of the usual tariff for customers consuming above 500 kWh. Fines thus targeted larger consumers. Customers who repeatedly exceeded their quotas were also under the *threat of power cuts* of 3 to 6 days. Customers consuming less than their quota were eligible for a *bonus* per kilowatt hour reduced below their quota. The bonus was guaranteed for customers consuming less than 100 kWh but was conditional on the revenue from the fines otherwise.[31] Therefore, the quotas effectively created a household-specific incentive schedule.

In practice, the implementation of these incentives was not smooth. First, the program was so successful that fines did not raise enough money to pay out non-guaranteed bonuses. In September 2001, the government then introduced a new guaranteed bonus for customers with quotas below 225 kWh. Second, distribution utilities did not have enough staff to implement power cuts, so they were limited to customers who repeatedly consumed far above their quota. Power cuts were even prohibited in Rio de Janeiro (Lei Municipal 3266/2001), a city for which we have microdata.

[29] Firms and the public sector also faced incentives in this period. Mation and Ferraz (2011) look at impacts on firms. We do not consider firms because the nature of their response to temporary corrective incentives may be very different and because changes in the industrial composition of the economy complicates the study of long-term effects. That firms were subject to some incentives is an issue for our purpose only to the extent that it indirectly affected households' electricity use. However, we control for employment or income effects in our empirical analysis.

[30] Quotas for small consumers were set at 100% of the baseline or 100 kWh, whichever was smaller. Quotas were revised upward in December 2001 and January 2002 because the situation was improving ("quota" in the rest of the paper refers to the initial quota). Customers received a letter with their quota prior to their first affected billing cycle (see Appendix E, where we also display the exact mapping between quotas and baseline consumption levels).

[31] The remaining funds from the fines, after paying the guaranteed bonuses, would then be divided among other complying customers. The bonus per kilowatt hour reduced below the quota was equal to at most 200% of the tariff. Fines and bonuses were added to electricity bills, which could not be negative, limiting the payment of bonuses. In Appendix E, we illustrate how these incentives modified the cost of electricity.

Finally, an information and conservation appeal ("social incentives") campaign was carried out in collaboration with distribution utilities and media outlets. Daily television reports compared achievements to policy targets. Media reports and electricity bills included energy conservation advice. They also included stories of exemplary behavior and appeals to social preferences and patriotism. This campaign reached the whole country because the main media outlets are national in Brazil. It may thus explain some spillover effects to unaffected areas. It may also explain part of the larger impact for households subject to the energy-saving program if households internalized the information and the social incentives differently because they were subject to the policy. However, our analysis of the causal effect of the quotas shows that the program's long-run effect cannot be rationalized by the information campaign or by some persistent change in social incentives.

Other factors would not have affected electricity use differentially between the Southeast/Midwest and the South. Tariff changes followed the usual regulatory framework during and after the crisis.[32] Accordingly, we find no differential change in the Southeast/Midwest if we replicate Figure 1 for electricity prices instead of electricity use (see Appendix A). Some policies were implemented nationally and are thus part of our counterfactual. Taxes on energy-efficient lightbulbs and appliances were reduced, and taxes on electric showers, water heaters, and incandescent lightbulbs were temporarily increased (Decreto 3827, May 21, 2001). Efficiency standards for domestic appliances were adopted (Lei 10295, October 17, 2001), but were only implemented in later years. Finally, the rainfall pattern of the 2000-2001 summer was an outlier, so there was no rational reason for customers to differentially update their beliefs about the risk of future shortages.[33] Even in that case, it is unclear whether households would have consumed less electricity in response, since the policy used grandfathering to assign quotas. Our analysis of the causal effect of the quotas confirms that the program's long-run effect cannot be rationalized by differences in policy expectations.

2.2 Data

Next, we describe the main datasets used in our analysis. Appendix F provides more details.

A. Utility-level administrative data. We use longitudinal data at the level of distribution utilities to estimate the overall impact of the energy-saving program. ANEEL provided us with monthly administrative data from mandatory reports of distribution utilities on total electricity consumption,

[32] This is with the exception of a 2.9% extraordinary increase for distribution utilities subject to the energy-saving program (*Camara de Gestão da Crise de Energia, Resolução 91*, December 21, 2001). This small price change is unlikely to drive any of our results. Moreover, we control for electricity tariffs in our empirical analysis.

[33] Accordingly, when the government created an *insurance fund* to prevent such crises, it chose to finance it through a nationwide tariff increase ($.51 per 100 kWh; Camara de Gestão da Crise de Energia, Resolução 115). Moreover, in the South, reservoir levels were also low in 2000, and they started to follow the same seasonal pattern as in the Southeast/Midwest after the crisis because of the increase in transmission capacity (see Figure 3a). Finally, the country had previously experienced smaller weather-induced shortages in all subsystems (Maurer, Pereira and Rosenblatt, 2005). These shortages had led to isolated blackouts and not to the implementation of any incentive-based program.

total revenues, and total number of customers by category (e.g., residential) from 1991 to 2014. Our main outcome, average residential electricity use, is equal to the total residential consumption divided by the number of residential customers. We also compiled data from every tariff regulation from 1996 to 2014. As a result, we have a balanced panel of average residential electricity use and residential electricity tariff at the monthly level for all distribution utilities. We also use decennial census data (2000 and 2010) and yearly data on population (1992-2014), GDP (1999-2012), formal employment (1996-2014), and average temperature (1996-2014), which are available at the municipality level and can be matched to distribution utilities based on their concession area.[34]

B. Household-level billing data. We use longitudinal data at the customer level for one distribution utility subject to the energy-saving program to evaluate the robustness of our results, to go beyond average effects, and to estimate the short- and long-run impacts of variation in the quotas. We obtained billing data from January 2000 to December 2005 for the universe of low-voltage customers of LIGHT, the distribution utility serving Rio de Janeiro and 31 surrounding municipalities in the Southeast. The data include about three million residential customers in 2000. They detail every bill component and include metering and billing dates, meter location, and the quantity consumed in each month. Customers are uniquely identified over time unless they move.

C. Household-level survey data. We use the last two rounds of the Survey of Appliances and Utilization Habits (PPH surveys) conducted by the Electrical Energy Saving Program of the Brazilian government (PROCEL) to shed light on the mechanisms of hysteresis. Representative samples of customers from several utilities were surveyed before and after the crisis (1998-1999 and 2004-2005) about their household characteristics, appliance ownership, and consumption habits. Moreover, in the second round, customers of distribution utilities that had been subject to the energy-saving program were asked about their behavior and experience during and after the crisis. We use the microdata from the 8 utilities in the Southeast/Midwest that were included in both rounds.[35]

3 The overall effect of the energy-saving program

This section presents our main results. We estimate the average effect of the energy-saving program by using our panel of distribution utilities in a difference-in-differences strategy that compares distribution utilities in the Southeast/Midwest and in the South over time. We begin by providing descriptive statistics to support our approach. We then present three sets of results. First, we consider a flexible specification in which we estimate separate difference-in-differences coefficients

[34] Census, GDP, and demographic data are from the Brazilian Institute of Geography and Statistics (IBGE). Formal employment data are from the Labor Ministry (RAIS data). Temperature data are from Matsuura and Willmott (2012).

[35] The name of these distribution utilities is masked in the data because of confidentiality concerns, so we cannot match the PPH surveys to the ANEEL or LIGHT data. We were told, however, that LIGHT is one of these utilities.

at a relatively high frequency to show the dynamics of the treatment effect. Second, we consider a sparser specification to summarize our results and evaluate their robustness. Third, we estimate the treatment effect for each distribution utility separately, using synthetic control methods.

3.1 Descriptive statistics supporting our approach

Table 1 provides descriptive statistics to support our identification assumption of a common trend in average residential electricity use for distribution utilities in the Southeast/Midwest and in the South. It also shows that this assumption is unlikely to hold, especially in the long run, for distribution utilities in the North and the Northeast. Columns (1)-(5) compare the mean and range of relevant variables in 2000 (pre-crisis) across distribution utilities in the four subsystems and for LIGHT separately. Columns (6)-(8) present the differential trend in these variables between 2000 and 2010, comparing distribution utilities in each subsystem to distribution utilities in the South.[36] Some of the variables discussed below are presented in a similar table in Appendix A.

Before comparing subsystems, we note that average residential electricity use was low in Brazil (below 200 kWh/month in 2000) compared to richer countries (903 kWh/month in the US in 2012; see www.eia.gov/tools/faqs), as households were less likely to own major domestic appliances.

We do not include distribution utilities from the North and the Northeast in our analysis for two main reasons.[37] First, nearly all households had electricity prior to the crisis in the South, but not in the North or in the Northeast. Moreover, the gap decreased by 8%-11% in the following decade. A common-trend assumption is unlikely to hold when customer bases evolve so differently. Second, households in the North and the Northeast were poorer and less likely to own major domestic appliances in 2000. For instance, there is no common support for the median household income and the ownership rate of refrigerators – a major source of electricity use – between utilities in the Northeast and in the South. Strong poverty alleviation, as experienced in Brazil in the following decade, can have very different effects on electricity use when initial ownership rates are so different (Wolfram, Shelef and Gertler, 2012). Accordingly, the ownership rate of refrigerators increased by 15%-26% in the North and the Northeast compared to the South between 2000 and 2010. This is unlikely a consequence of the crisis and thus violates a common-trend assumption.[38]

[36] We use yearly data from the most recent census years $t \in \{2000, 2010\}$ and the following specification:

$$log\left(y_{i,t}\right) = \alpha_i + \gamma \cdot \mathbb{1}\{t = 2010\} + \delta \cdot \mathbb{1}\{t = 2010\,\&\,\text{TreatRegion}_i = 1\} + \nu_{i,t}.$$

where a_i is a fixed effect for distribution utility i, and $TreatRegion_i$ indicates a distribution utility from a subsystem subject to the energy-saving program. $\nu_{i,t}$ is an error term clustered by distribution utility. Columns (6)-(8) report estimates of $\widehat{\delta}$ for samples including distribution utilities from one of the other three subsystems and from the South.

[37] There is also a data limitation. Many customers in the North are served by isolated electricity systems. Our data do not differentiate consumption from "isolated" and "connected" customers, and the formers were not subject to the energy-saving program. The policy also started later (August) and ended earlier (December) in the North.

[38] In contrast, the share of households with a TV was comparable in 2000 and did not increase more in the North

In contrast, distribution utilities in the South constitute a more credible counterfactual for those in the Southeast/Midwest. First, Figure 1 shows that distribution utilities in the two subsystems shared a common trend in average residential electricity use prior to the crisis (we explicitly test this below). Second, nearly all households had electricity in 2000 in both subsystems, and customer bases evolved at a similar rate in the following decade. Moreover, this does not hide any differential trend in population size, urbanization, household size, dwelling size, or dwelling characteristics. Third, initial ownership rates of major domestic appliances (refrigerator, washing machine, TV, air conditioner) and electricity tariffs, as well as and their subsequent growth, were comparable.

Distribution utilities in the two subsystems differ in some respects. For instance, average residential electricity use was higher in 2000 in the Southeast/Midwest, where median income levels were higher and average temperatures are always higher. Median income levels also grew slower in the Southeast/Midwest in the following decade, although labor market outcomes (employment, formal employment, and farm employment) did not evolve differentially (we attribute the differential trend in average residential electricity use to the impact of the energy-saving program below). Importantly, however, the range of initial values and of changes in these values in subsequent years overlapped for all these variables and others (see Appendices G-J). Therefore, we can control for relevant variables without purely relying on parametric assumptions.

In sum, the data appear to support our common-trend assumption, at least conditionally on controlling for some relevant factors. Additionally, as explained in Section 2, the timing of the crisis and the differential treatment between subsystems were entirely due to supply factors that are likely exogenous to potential changes in households' electricity use.

3.2 Flexible specification and graphical presentation of the main results

We first estimate the overall short- and long-run effects of the energy-saving program through a flexible difference-in-differences specification to best display our main result graphically. We divide our monthly observations into yearly periods before 2001 and after 2002; we divide 2001 and 2002 into three periods: pre-crisis (early 2001), crisis (June 2001–February 2002), and post-crisis (rest of 2002). We then regress the logarithm of average residential electricity use $Y_{i,t}$, for distribution utility i in calendar month t (e.g., June 2001) using the following specification:

$$Y_{i,t} = \alpha_i + \theta_{r,m} + \gamma_p + \delta_p \cdot Treat_i \cdot \mathbb{1}\left\{t \in TimePeriod_p\right\} + X_{i,t}\beta + \nu_{i,t}. \tag{10}$$

The coefficients α_i, γ_p, and $\theta_{r,m}$ are fixed effects for each distribution utility, each period, and each month m per subsystem r (e.g., June in the Southeast/Midwest; they absorb seasonality ef-

and the Northeast between 2000 and 2010. Note that ownership rates of air conditioners were relatively low in 2000.

fects), respectively. The coefficients δ_p estimate a differential effect in each period for distribution utilities in the Southeast/Midwest ($Treat_i = 1$), which were subject to the energy-saving program. Estimates of δ_p during and after the crisis capture the average treatment effect of the policy on the treated under a common-trend assumption. Estimates of δ_p in the preceding periods test for a common-trend prior to the crisis. We also reinforce the common-trend assumption by controlling for variables that may be correlated with other factors affecting electricity use. Specifically, we use the following variables (in log) available at the distribution utility level in each year ($X_{i,t}$): main residential electricity tariff (1996-2014), population size (1996-2014), GDP (1999-2012), formal employment levels (1996-2014), and average temperature (1996-2014). Our preferred specification includes all these controls and is thus restricted to a balanced panel of utilities between 1999 and 2012. Finally, we cluster error terms, $\nu_{i,t}$, at the level of the distribution utility.[39]

Figure 4a displays the estimated $\widehat{\delta_p}$'s with their 95% confidence intervals. The pre-crisis period (early 2001) is the reference period. Point estimates are close to 0 prior to the crisis, supporting our common-trend assumption. Average residential electricity use dropped by -.26 log points (-23%) in the Southeast/Midwest compared to the South when the energy-saving program came into effect. This is the first quasi-experimental estimate of the short-run effect for residential customers, but it is well known that electricity use was successfully reduced at the time. Our main empirical result is that about half of that effect persisted in the long run. Consumption levels partially rebounded after the crisis, but point estimates are stable from 2005 onward at about -.12 log points (-12%).

Our estimates in Figure 4a suggests that it is sufficient to divide our observations into four time periods to describe our results: a pre-crisis period, a crisis period, a rebound post-crisis period until 2005, and a stable post-crisis period after 2005. Below, we thus adopt a sparser specification with these four periods to summarize and evaluate the robustness of our results.

3.3 Summary of the main results and robustness checks

Table 2 displays the results from estimating variants of the following specification:

$$\begin{aligned} Y_{i,t} = {} & \alpha_i + \theta_{r,m} + \gamma_t + \beta_S \cdot Treat_i \cdot Crisis_t \\ & + \beta_R \cdot Treat_i \cdot PostCrisisUntil2005_t + \beta_L \cdot Treat_i \cdot PostCrisisAfter2005_t + X_{i,t}\beta + \nu_{i,t}. \end{aligned} \tag{11}$$

The coefficients α_i and $\theta_{r,m}$ are the same fixed effects as above; $X_{i,t}$ is the same vector of covariates; and the coefficients γ_t are now fixed effects for each calendar month t (e.g., June 2001). The coefficients β_S, β_R, and β_L capture the average treatment effect of the policy on the treated in the

[39]Our independence assumption seems reasonable. The two subsystems cover large and heterogeneous areas providing electricity to more than 100 million individuals. Moreover, electricity sector policy, including policies related to energy efficiency, is centralized at the federal level in Brazil.

short run (crisis months), in the rebound period (March 2002-December 2005), and in the long run (from January 2006 onward). The average difference in the months leading to the crisis is absorbed by the fixed effects α_i. Finally, we cluster error terms, ν_{it}, at the level of the distribution utility.

Column (1) displays the estimated $\widehat{\beta_S}$, $\widehat{\beta_R}$, and $\widehat{\beta_L}$ for a specification that uses the same sample and set of covariates used in Figure 4a. Accordingly, our estimates are consistent with the estimates in Figure 4a: average effects of -.26 log points (-23%) during the crisis, -.14 log points (-13%) in the rebound period, and -.12 log points (-12%) in the stable post-crisis period after 2005.

Columns (2)-(8) then display similar results from estimating variants of this specification. In column (2), we estimate a specification in levels instead of logarithms, for both the outcome and the covariates. The percentage effects correspond to 43 kWh in the short run, 25 kWh in the rebound period, and 21 kWh in the long run. Column (3) shows that results are essentially identical to those in column (1) when we improve the comparability of distribution utilities in the two subsystems by excluding "outliers" from the Southeast/Midwest. In particular, we exclude the 8 distribution utilities with levels of key variables in 2000 (average residential electricity use, main residential electricity tariff, median household income) that fall outside the range of values observed for these variables in the South in the same year (see Table 1). Column (4) shows that results are slightly larger when we weight distribution utilities by their customer base in 2000. For instance, the long-run effect now amounts to -.15 log points (-14%). Results are also slightly noisier (but still significant) because customer bases are unevenly distributed across distribution utilities. Column (5) shows that results are not driven by a particular season (e.g., the summer when consumption levels are higher), as they are unchanged if we only include winter months. Column (7) shows that the results are essentially identical if we use all the years for which we can measure our outcome (1991-2014), even though we cannot include our set of covariates ($X_{i,t}$) in this specification.[40]

Finally, in columns (6) and (8), we use another specification to test whether the energy-saving program caused not only level shifts but also trend breaks (Greenstone and Hanna, 2014):

$$\begin{aligned} Y_{i,t} = \alpha_i + \theta_{r,m} + \gamma_t + \delta_B \cdot Treat_i \cdot BeforeCrisis_t \cdot t_B + \beta_S \cdot Treat_i \cdot Crisis_t + \delta_S \cdot Treat_i \cdot Crisis_t \cdot t_S \\ + \beta_R \cdot Treat_i \cdot PostCrisisUntil2005_t + \delta_R \cdot Treat_i \cdot PostCrisisUntil2005_t \cdot t_R \\ + \beta_L \cdot Treat_i \cdot PostCrisisAfter2005_t + \delta_L \cdot Treat_i \cdot PostCrisisAfter2005_t \cdot t_L + X_{i,t}\beta + \nu_{i,t} \quad (12) \end{aligned}$$

The coefficients δ_B, δ_S, δ_R, and δ_L test for a differential trend in treated distribution utilities before

[40] The number of distribution utilities is slightly smaller because a few distribution utilities were split or created after 1991. In Appendix A, we also show that the estimated long-term effect is robust to controlling for covariates that are not available at the yearly level but are available in the 2000 and the 2010 censuses (median household income, population size, share of households living in urban areas, average household size, average dwelling size, share of dwellings with a bathroom, and employment rate). Moreover, we graphically show that long-term changes in consumption levels are systematically lower for distribution utilities in the Southeast/Midwest than in the South for given baseline levels or long-term changes in all the variables in Table 1 (see Appendices I and J).

the crisis, during the crisis, in the rebound period, and in the stable period post-crisis, respectively. The linear trends (t_B, t_S, t_R, t_L) are normalized to the start of their respective period, such that the coefficients β_S, β_R, and β_L estimate the average difference-in-differences at the start of each period (the difference pre-crisis is still absorbed by the fixed effects α_i). The coefficients β_S, β_R, and β_L then capture level shifts, and the coefficients δ_S, δ_R, and δ_L capture trend breaks. Additionally, the coefficient δ_B tests for a common-trend between treated and control distribution utilities pre-crisis.

Columns (6) and (8) display results that compare to those in columns (1) and (7), respectively. We find no evidence of a differential trend pre-crisis, which supports again our common-trend assumption. During the crisis, the level shift is slightly less negative than the average effect because we estimate a negative trend break. The treatment effect was thus not fading away during the crisis. In the rebound period, the level shift is slightly more negative than the average effect because we find a positive trend break, like in Figure 4a. Finally, we find no evidence of any differential trend after 2005, showing again that our estimated long-run effect is stable. The last row in columns (6) and (8) also provides estimates of the 10-year effect of the short-run policy based on the estimated level shift and trend break in the stable period post-crisis. The 10-year effect is slightly smaller (resp. larger) than the average long-run effect in column (6) (resp. in column 8) because $\widehat{\delta_L}$ is positive (resp. negative but rounded to 0), although it is not statistically different from 0.

In sum, the estimates presented in Figure 4a are quite robust. The short-run policy led to large reductions in electricity use in the short run, and about half of the short-run effect persisted in the long run. Moreover, the long-run effect did not slowly fade away but instead appears permanent.

3.4 Synthetic control estimates

We end this section by estimating separate treatment effects for each distribution utility using synthetic control methods (Abadie, Diamond and Hainmueller, 2010). This serves two purposes. First, it shows that the estimated effects in Figure 4a and Table 2 are not driven by a few distribution utilities. Second, it shows that the treatment effects are also large for LIGHT, the distribution utility for which we have microdata. This motivates our analysis of these data in the next section.

Our synthetic control estimator is the difference between the logarithm of average residential electricity use (demeaned and seasonally adjusted for this exercise) in a treated distribution utility and in a weighted average of the control distribution utilities. The vector of weights W is chosen to minimize $\| Y_{d0} - Y_{c0}W \| = \sqrt{(Y_{d0} - Y_{c0}W)' V (Y_{d0} - Y_{c0}W)}$, where Y_{d0} and Y_{c0} are vectors containing the values of the outcome in pre-crisis periods in the treated and control distribution utilities, respectively. An optimal choice of V minimizes the mean squared error of the estimator.

Figure 4b displays our results. Monthly estimates are averaged into time periods similar to those used in Figure 4a, but we use all years from 1991 to 2014 because we do not control for any covariate. Darker lines display the estimates for the treatment distribution utilities. Lighter lines display placebo estimates in which we compare each control distribution utility to a weighted average of the other ones. The synthetic controls closely match the trends pre-crisis. The estimated short-run effect is large for all treated distribution utilities at between $-.18$ and $-.40$ log points (-16% and -33%). The long-run effect is negative for all these utilities. The average long-run effect in Figure 4a and Table 2 is thus not driven by outliers (we show this in the raw data in Appendix H). The median and the average of these effects are equal to -.14 log points (-13%) in 2014. In contrast, the median and the average of our placebo estimates are close to 0 in all years.[41]

The evidence in this section shows the importance of hysteresis in our context. A 9-month intervention led to a persistent reduction in average residential electricity use by about 12% until at least 12 years after it ended and potentially well beyond that given that the long-run effect appears to be stable. To give an order of magnitude, this long-run effect amounts to a total reduction of 100.5 billion kWh in the Southeast/Midwest between 2003 and 2014.[42]

4 Evidence using the household-level billing data

Figure 4b also shows that the long-run effect for LIGHT is comparable to the average long-run effect that we estimate using all distribution utilities (slightly larger at -.16 log points or -15%). This supports the relevance of the analysis in this section, in which we delve into the microdata from LIGHT to go beyond the average effect of the overall energy-saving program. In a first step, we analyze the anatomy of the persistent reduction in average residential electricity use. Among other things, we show that changes in electricity use during the crisis were also persistent at the household level, and not just at the level of the distribution utility. Some of these short-run changes, however, may have been unrelated to the incentives of the energy-saving program. In a second step, we thus estimate the causal effect of a specific type of policy variation, namely variation in consumption quotas across households, on short-run and long-run electricity use levels.

4.1 Anatomy of the reduction in average residential electricity use

Figure 5 presents our analysis of the anatomy of the reduction in average residential electricity use.

A. Validating the aggregate utility-level data and ruling out composition effects. Figure 5a shows that the changes in average residential electricity use during and after the crisis are similar in the aggregate LIGHT data used in the previous section and in the microdata. It displays the

[41] We show results from a similar exercise for the main electricity tariff in Appendix A. In that case, the range of estimated effects completely overlaps and centers around 0 for distribution utilities in the two subsystems.

[42] This is calculated as: $\sum_{t=2003}^{2014} .12 \cdot \frac{X_t}{1-.12}$, where X_t is the yearly residential consumption in the Southeast/Midwest.

average residential electricity use based on the two datasets in each month between 2001 and 2005 (last year in the microdata), compared to the same month in 2000 (first year in the microdata) to show changes and net out seasonality effects. The two graphs are similar, which indicates that the aggregate data reported to the regulator is consistent with the metering data, a finding that supports our use of the utility-level data to study actual changes in electricity use in the previous section.

A remaining limitation of utility-level data is that they could be mixing changes in household electricity use and composition changes in the customer base of distribution utilities (Levinson, 2014).[43] Figure 5a, however, shows that changes in average residential electricity use are also similar if we restrict attention to the balanced panel of all customers metered regularly between 2000 and 2005. Composition effects, which are absent from the balanced panel by construction, are thus unlikely to drive our results, at least until 2005 (Figure 4a shows that the average long-run effect is stable after 2005). This also rules out any important role of "electricity theft" responses at the extensive margin, that is households canceling their legal connection and instead stealing electricity. Such households are excluded from the balanced panel by construction.[44]

B. Changes in the distribution of electricity use levels. Figure 5b shows that the short-run and the long-run changes in average residential electricity use came from sizable reductions at every level of consumption. It displays kernel densities for average monthly electricity use before, during, and after the crisis for the same balanced panel of customers used in Figure 5a. The density during the crisis is first-order stochastically dominated by the other ones. Densities 1 year and 4 years after the crisis are similar and fall exactly between the crisis and pre-crisis densities.

C. Distribution of changes at given quota levels (and thus baseline consumption levels). Figure 5c shows that, at a given quota level (and thus at a given baseline consumption level), most customers severely reduced their electricity use. It displays the distribution of changes in electricity use during and after the crisis compared to before the crisis for a balanced panel of customers with a quota around 300 kWh. We consider customers with similar quotas such that they faced similar incentives during the crisis. We also consider customers with a relatively high baseline consumption (around $300/.8 = 375$ kWh) because their economic incentives were simpler; they essentially faced only fines for exceeding their quota. However, we show similar patterns for other consumption levels in Appendix L. During the crisis, 98% of these customers reduced electricity use below

[43] There is no evidence of meaningful migration across regions between 2000 and 2010. de Oliveira and de Oliveira (2011) document that the Southeast experienced a net out-migration during the period. However, magnitudes are very small at no more than 0.2% of the Southeast population and 0.5% of the South population.

[44] There are no good data on electricity theft. Distribution utilities are supposed to report yearly information on distribution losses to the regulator, but many did not provide this information prior to 2000. In Appendix K, we use yearly reports for 24 distribution utilities in the Southeast/Midwest and in the South from 1998 to 2008. The data are very noisy, and, if anything, point estimates suggest that non-technical losses (a measure of theft) decreased compared to 2000. We also use microdata from the Brazilian Household Expenditure Survey (POF 1996/97, 2002/03, and 2008/09) and find no differential trend in the share of households that do not pay for electricity.

their pre-crisis level (in the same months), and the median customer reduced consumption by 34%. Their consumption level was thus much lower than their quota: 91% reduced consumption below their quota and the median customer reduced consumption by 22% below its quota (see Appendix A). This suggests that social incentives played a role at the time, as these customers had no other incentive to further reduce consumption once complying with their quota. Four years after the crisis, 78% were still using less electricity than before the crisis, and the median customer was using 22% less electricity. These figures are only slightly smaller in the overall balanced panel.[45]

Together, Figures 5b and 5c also provide evidence against electricity theft responses at the intensive margin. Establishing illegal connections to the grid to obtain part of one's electricity use is more likely to take place among poorer and thus smaller consumers in Brazil. However, reductions in electricity use were not concentrated among small consumers. Moreover, if electricity theft took place among relatively large consumers, it would likely be concentrated within a small group. We know that this was not the case because median responses were so large.

D. Persistence of household-level changes in electricity use during the crisis. Finally, Figure 5d displays the correlation between short-run changes in electricity use during the crisis and long-run changes 4 years after the crisis, compared to before the crisis, for the panel of customers with the same quota levels used in Figure 5c (we show similar patterns for other quota levels in Appendix L). The strong correlation (.69) provides household-level evidence that changes in electricity use that took place during the crisis were very persistent.

4.2 Causal effect of quotas on short-run and long-run electricity use levels

Next, we estimate the causal effect of quasi-experimental variation in quotas among LIGHT customers on their electricity use, both in the short run (crisis) and in the long run (until 2005). This exercise is useful for two reasons. First, the correlation in Figure 5d may overstate the household-level hysteresis in electricity use caused by the energy-saving program, as part of the short-run changes in electricity use may have been due to other factors, such as mean reversion. Second, because we use idiosyncratic variation in quotas among comparable customers, the hysteresis in electricity use that we estimate here is unlikely due to persistent differences across households in social incentives (or in policy expectations or in the information that they were exposed to during the crisis.[46] As discussed in Section 1, these could affect the interpretation of our estimates of

[45]In the overall balanced panel, the statistics comparing consumption after vs. before the crisis are 92%, 31%, 69%, and 16.5%, respectively. Customers with higher baseline consumption levels made larger proportional reductions in electricity use (see Appendix L), but this could be driven by mean reversion. We find no bunching at the quota (see Appendix A), which is consistent with other work on nonlinear electricity price schedules (Borenstein, 2009).

[46]The households were comparable at baseline and were exposed to the same conservation appeals during the crisis, so they likely faced similar social incentives once the policy ended (if any); future policies would have unlikely replicated the same idiosyncratic variation in incentives; and they were exposed to the same information campaign.

hysteresis above for the welfare evaluation of corrective policies.

4.2.1 Source of quasi-experimental variation in quotas

We use the quota assignment rules among customers who moved into their housing unit ("movers") in the months preceding the crisis. The quota was set at 80% of their average consumption between May and July 2000 for the typical households, but at 80% of the average consumption in their first 3 billing months for those households that moved into their housing unit after May 2000. Because of strong seasonality in electricity use in Rio de Janeiro, these rules created wide variation in quotas among movers entirely because of differences in moving dates. Figure 6a shows the seasonality in electricity use for LIGHT customers. It displays the average consumption in each month, between March 2000 and February 2001, for the overall balanced panel of customers in Figure 5 and for the top quartile of the distribution in this sample.[47] Consumption levels are higher in the austral summer and lower in the austral winter. The difference reaches more than 60% in the top quartile of the distribution, showing that the seasonality is particularly strong among large consumers.

Figure 6a also shows the impact of the seasonality in electricity use on the movers' quotas. It displays the average quota of customers who moved into their housing unit in the same month, as measured by the date of their first electricity bill. The sample is restricted to large consumers given the stronger seasonality in their electricity use, that is movers whose average consumption prior to the crisis (March-May 2001) was in the top quartile of the distribution. Customers who moved into their housing unit before May 2000, whose quota was based on their consumption in the typical baseline period, had similar average quotas. In contrast, the quota of customers who moved in after May 2000, which was based on their consumption in their first 3 billing months, closely follows the seasonality. As a result, Figure 6b shows that the quota distribution for summer movers first-order stochastically dominates the distribution for winter movers; at the median, the difference reaches 32%, which is equivalent to offering a non-binding quota to summer movers.

4.2.2 Empirical strategy

We exploit the variation from differences in moving dates to estimate the causal effect of the quotas on electricity use during and after the crisis. We use the following specification in different periods for household i that moved in neighborhood n in week t (e.g., first week of May 2000):

$$Y_{i,n,t} = \alpha_n + \beta \cdot \overline{Quota}_t^{-i} + \varepsilon_{i,n,t} \tag{13}$$

[47] More than 25% of households in the concession area of LIGHT owned an air conditioner in 2000 (see Table 1).

where $Y_{i,n,t}$ is the logarithm of the quota assigned to household i ("first-stage" specification) or the logarithm of average monthly electricity use in a given period ("reduced-form" specification). α_n is a neighborhood fixed effect. The variable $\overline{Quota}_t^{-i}$ captures the variation in quotas from differences in moving dates; it is the logarithm of the average of the quota of all movers (excluding i) who received their first bill in the same week as household i (we cluster standard errors by moving week). Below, we present evidence that supports our exclusion restriction that the average quota of same-week movers does not directly affect households' consumption levels.

Our sample includes customers who moved in from March 2000 – to make sure that they did not receive any bill in earlier months – until February 2001 – to have 3 month of electricity use prior to the crisis for every mover. We also restrict attention to movers whose average consumption prior to the crisis (March-May 2001) was in the top quartile of the distribution. Figure 6 shows that the first stage is likely strong in this sample. Additionally, because we are interested in long-term effects, we focus on movers metered regularly until the end of 2005. Estimates of short-run effects are similar without this restriction (see Appendix A). Finally, we present results for both mean and median effects.[48] This is because economic incentives were simpler for the median household in each moving week. Figure 6b shows that the median quota was large enough (even for winter movers) such that they were essentially faced with only fines for exceeding their quota. The monotonicity assumption is thus likely satisfied at the median. Figure 6b also shows that a few customers had quotas small enough (even for summer movers) such that they could be eligible for bonuses. Given the incentive schedule of the energy-saving program, the monotonicity assumption may be violated for some of these movers.[49] We avoid this issue by looking at median effects.

4.2.3 Results

Table 3 presents our results. Column (1) shows that the average quota of same-week movers is positively correlated with movers' own quota (first stage), which is consistent with the patterns in Figure 6. The coefficient is .79 and .66 for the mean effect and the median effect, respectively. Column (2) shows that the average quota of same-week movers does not predict differences in consumption levels before the crisis, however. This evidence supports our exclusion restriction.

The next columns show that the average quota of same-week movers is positively correlated with consumption levels during and after the crisis (we consider the same months between July and December in each year), which we interpret as evidence of a causal effect of the quota. In column (3), we estimate that a 10% increase in the quota of same-week movers increased electricity use

[48]Quantile regressions do not include neighborhood effects because estimators would be inconsistent, and standard errors are obtained from 1,000 bootstrap replications.

[49]For instance, it is unclear whether a customer will want to use more electricity if given a quota of 80 kWh vs. 60 kWh. In both cases, the customer does not face penalties for exceeding the quota. Moreover, it receives bonuses as soon as it consumes less than 80 kWh in the first case, but only if it consumes less than 60 kWh in the second case.

during the crisis by 1.37% for the mean and 1.15% for the median. The effect is precise, but it is not very large. This finding is consistent with the fact that customers tended to reduce consumption well below their quota (see discussion of Figure 5c above). Our estimates increase to 1.38% for the mean and 1.53% for the median when we control for pre-crisis consumption levels to absorb part of the household heterogeneity (we do so in the next columns as well). The effect is smaller in 2002 (0.87% for the mean and for the median), a finding that is consistent with the rebound in electricity use levels in 2002 in Figure 4a, but it remains positive and significant. The effect becomes smaller and noisier between 2003 and 2005, but it is significant in most cases and relatively stable. There is thus no evidence that the effect would have decreased after 2005. The loss in precision is not surprising, as idiosyncratic shocks likely weaken the moderate link between quotas and consumption levels over time. In column (9), we thus estimate the long-run effect in a specification that pools the years between 2003 and 2005 (adding year fixed effects). The long-run effect is more precise in this case: a 10% increase in the quota of same-week movers leads to an increase in electricity use of 0.49% for the mean and 0.86% for the median.

Finally, to support our causal interpretation and our exclusion restriction, we present placebo estimates in Appendix A. We use customers who moved into their housing units in similar months but in the years after the crisis (not in 2000-2001). The average consumption of same-week movers in their first 3 billing months does not predict differences in consumption levels in these samples.

In sum, the degree of hysteresis – the ratio of the long-run to the short-run effect – implied by our estimates in Table 3 ($\frac{.049}{.138} = .355$ and $\frac{.086}{.153} = .562$) is broadly comparable to the degree of hysteresis implied by our estimates in Table 2 ($\frac{-.123}{-.264} = .466$) using utility-level data.

5 Mechanisms of hysteresis

Before using our results to illustrate the welfare implications of hysteresis, we provide here some qualitative evidence for the mechanisms of hysteresis using the microdata of the PPH surveys.

5.1 Main sources of electricity use prior to the crisis

We begin by documenting the main sources of household electricity use in the Southeast/Midwest prior to the crisis. Panel A in Table 4 displays the average quantity per household for each major domestic appliance from the 1998-1999 PPH surveys conducted in a representative sample of 6,482 customers of 8 distribution utilities in that subsystem. It also displays the average electricity use for each appliance, using additional information on its average utilization rate (from the survey) and on the kWh consumption of the average model calculated by PROCEL (see Appendix A).

The main sources of electricity use in our context substantially differ from those in the US.

Electric showers – electric heating devices placed in a shower head – consume the most electricity. These are popular technologies for water heating in Brazil and other developing countries that have a low fixed cost but consume a lot of electricity. Next come refrigerators, lights, TVs, and freezers, which are important sources of electricity use when households own relatively few and inefficient appliances.[50] In contrast, air conditioners only account for a small share of electricity use given that their penetration rates (and usage rates) are much lower than in the US.

5.2 Mechanisms behind the estimated long-run effect

Panel B in Table 4 documents the main mechanisms of hysteresis for the energy-saving program, as reported by the 4,579 customers of the same distribution utilities surveyed in the 2004-2005 PPH surveys. For each major domestic appliance, and conditional on owning the appliance prior to the crisis, households in the Southeast/Midwest were asked to choose the statement that best described their situation: (1) we use it as much as before the crisis; (2) we use it less than before the crisis; (3) we disconnected or disposed of the appliance during or after the crisis; (4) we substituted a more energy-efficient model during or after the crisis. Panel B tabulates their answers.

For every appliance, the share of households reporting that they use it less than before the crisis is large (e.g., 39% for electric showers, which can be set in a colder mode; 41% for lights; and 21% for freezers) and larger than the share of households reporting that they stopped utilizing it or that they substituted a more energy-efficient model. The one exception is for refrigerators: most households reported using them as much as before the crisis, which is in fact reassuring for the overall quality of households' responses, as we don't expect much flexibility in refrigerators' utilization. A sizable share of households reported disconnecting or disposing of their freezers (17%; households can rely on the smaller freezer unit in their refrigerator). Finally, very few households reported replacing appliances with more energy-efficient models, except for lights (8%). In subsequent survey questions, households reported substituting fluorescent for incandescent lightbulbs during the crisis, and continuing to use fluorescent lightbulbs after the crisis (see Appendix A).

In sum, the main reported mechanism of hysteresis is a change in utilization habits. The apparent lack of sizable investments in new physical capital is also consistent with the fact that manufacturers of appliances expected ex-ante and reported ex-post losses due to the electricity crisis (*Folha de São Paulo*, June 5, 2001 and March 6, 2002). In 2001, Brazil was also the country with the highest real interest rate in the *World Development Indicators* (44.65% vs. an OECD average of 8.34%), limiting households' ability to finance physical investments. The evidence in Panel B remains suggestive, however, and does not imply that other mechanisms did not play a role. Households were not allowed to choose more than one answer for each appliance, and this question was

[50] It was common at the time, and remains relatively common, in many countries for households to have smaller refrigerators than those found in the US, with a small freezer unit, and to have a separate larger freezer unit.

not asked to customers of distribution utilities in the South, so we do not have a counterfactual.

We present other attempts to shed light on the mechanisms of hysteresis in Appendices N to Q. The results there are also suggestive as we remain limited by the data available to study this question so many years after the crisis. First, we exploit the fact that the PPH surveys asked all households in both rounds and in both the South and the Southeast/Midwest about their appliance characteristics and their utilization habits. Unfortunately, difference-in-differences results based on these data always have large standard errors when clustered by distribution utility because the surveys only include two distribution utilities in the South. Nevertheless, this analysis reveals that the adoption of fluorescent lightbulbs was a nationwide phenomenon and was stronger in the two distribution utilities in the South than in those in the Southeast/Midwest. So the adoption of fluorescent lightbulbs may not explain any part of the long-run effect of the energy-saving program. In contrast, point estimates for specific changes in utilization habits are consistent with the households' answers in Panel B and are large enough to rationalize the long-run effect. Finally, we also document patterns consistent with our qualitative findings using time-series data on sales of appliances from manufacturers and microdata from the Brazilian Household Expenditure Surveys.

6 Implications for the welfare evaluation of corrective policies

We established that households reduced electricity use by about 12% in the long run following a short-run policy that induced a short-run reduction of about 23%. In this section, we combine these estimates with the conceptual framework and the formulas derived in Section 1 to illustrate the possible implications of such hysteresis for the welfare evaluation of corrective policies.

6.1 Benchmark framework

We begin by considering the welfare implications of hysteresis under the assumptions of the benchmark framework. In that case, one would underestimate the welfare effect of short-run and long-run corrective policies by using textbook formulas that assume away hysteresis (see Section 1).

The welfare effect of a short-run policy is larger in the presence of hysteresis because it includes an additional correction of the externality in the long run. We illustrate the magnitude of the bias from assuming away hysteresis in our context as follows. First, we assume a fixed marginal damage, such that the bias in equation (7) is equal to the (discounted) product of the marginal damage and the long-run reduction in electricity use. Second, we approximate the marginal damage by the social value of reducing electricity use based on the revealed preferences of the government (like in, e.g., Ashenfelter and Greenstone, 2004).[51] Every year, PROCEL reports the resources that

[51] Estimating the marginal damage in our context through other approaches is beyond the scope of this paper.

it invested and an estimate of the electricity saved as a result. These reports thus provide yearly values of the government's willingness-to-pay to reduce electricity use (see procelinfo.com.br). Specifically, from 2003 to 2014, PROCEL reports to have reduced electricity use by a total of 66.2 billion kWh, or much less than our estimate of the long-run effect of the energy-saving program, at an average willingness-to-pay of US$12.9 million per billion kWh. Using this figure, the government would have been willing to pay US$874 million for the long-run electricity savings due to the short-run policy between 2003 and 2014 (we discount values at a 5% yearly rate starting in 2001 in all calculations).[52] The bias would be even larger if we extended the horizon beyond 2014.

In the case of a long-run policy, the welfare effect is larger because one would overestimate the loss in private welfare (i.e., consumer surplus) from the policy by assuming away hysteresis. We illustrate the magnitude of this bias as follows. First, as discussed in Section 1, we focus on the first source of bias in equation (9). Second, we consider a hypothetical long-run policy aimed at reducing electricity use by $\frac{x_{10}-\bar{x}_1}{x_{10}} = \frac{x_{20}-\bar{x}_2}{x_{20}}$ =23% for the 163 months between June 2001 and December 2014. Period 1 in our framework (short-run) is assumed to correspond to the first 9 months of the policy and period 2 (long-run) to the following 154 months. In this case, the bias arises because the policy must only reduce consumption by $\frac{x_2(\bar{x}_1)-\bar{x}_2}{x_2(\bar{x}_1)} = \frac{(1-.12)\cdot x_{20}-(1-.23)\cdot x_{20}}{(1-.12)\cdot x_{20}} = 12.5\% < 23\%$ in period 2, using our estimate of the long-run effect of the short-run policy. This also implies a smaller change in the household's marginal utility from electricity use. For instance, we estimate a price elasticity of average residential electricity use of $-.31$ in the Southeast/Midwest after the crisis (see Appendix A).[53] Using this elasticity, the impact of the policy on the marginal utility from electricity use in period 2 amounts to a $\frac{-.23}{-.31} = 74.2\%$ increase in electricity tariffs if we assume away hysteresis, but only to a $\frac{-.125}{-.31} = 40.3\%$ increase given our estimate of hysteresis. Third, we use a linear approximation for the demand curve, such that the first source of bias amounts to:

$$|DWL_{NoH}| - |DWL| = \frac{1}{2}\sum_{t=10}^{163} \beta^t \cdot [(.23 \cdot x_{t0}) \cdot (.742 \cdot p_t) - (.125 \cdot x_t) \cdot (.403 \cdot p_t)] = \text{US\$7.75 billions},$$

where p_t and x_t are the average price and the total residential consumption in the Southeast/Midwest in each month after the crisis, respectively, and $x_{t0} = \frac{x_t}{(1-.12)}$ is the counterfactual consumption absent hysteresis (we found no impact on prices). A bias of US$7.75 billions corresponds to 7.18% of the discounted sum of residential electricity bills in the Southeast/Midwest over the period. Our estimates also imply that, assuming away hysteresis, one would overestimate the loss in consumer surplus in period 2 by $\frac{.23\cdot .742/(1-.12)-.125\cdot .403}{.125\cdot .403} = 285\%$. Moreover, this figure may constitute a lower bound because it only takes into account the hysteresis due to the first nine months of the

[52]This is calculated as $\sum_{t=2003}^{2014} \beta^{t-2001} \cdot \$12.9 \cdot (.12 \cdot \frac{X_t}{1-.12})$, where X_t is total yearly residential consumption in the Southeast/Midwest and $\frac{X_t}{(1-.12)}$ is the counterfactual consumption absent hysteresis.

[53]We use quasi-experimental variation in yearly electricity tariffs across distribution utilities.

policy and only until 2014, and because it abstracts from the third source of bias in equation (9).

6.2 Extensions to the benchmark framework

Heterogeneity. We have shown that the sufficient statistics formulas are essentially unchanged with heterogeneous households, when corrective policies are implemented efficiently. We recognize that the energy-saving program is unlikely to have been implemented efficiently, but we cannot quantify its allocative inefficiencies. These inefficiencies would increase the loss in private welfare whether or not one assumes away hysteresis, so it is unclear how they would affect our takeaways.

Uncertainty. We have shown that the sufficient statistics formulas are similar in a rational inattention model introducing uncertainty to our framework. Some qualitative evidence provides support for such a model. Panel C in Table 4 displays the share of households in the Southeast/Midwest that chose different answers to a question from the 2004-2005 PPH surveys about their experience during the crisis. Most households that complied with their quota report that it was either "not difficult at all" (43%) or "not very difficult" (48%) to do so. These answers appear at odds with the assumptions of the benchmark framework: if the costs of large short-run reductions in electricity use were small, households should have changed their behavior on their own given the long-run savings in electricity bills. Yet, these answers are consistent with a model in which the relevant costs are mostly attention costs, rather than the costs of electricity conservation per se.

Panel C in Table 4 also displays the answers to a second question, in which only 28% of households report that the energy-saving program negatively affected their quality of life. Most households report no change in life quality (48%) or that they "learned to live comfortably while saving money" (24%). These answers are less consistent with a model in which attention costs are real. They are consistent with models in which households benefited from the policy, e.g., because they were ex-ante overestimating the loss in their utility from reducing electricity use, as with present bias or myopic beliefs, or because they derived utility from possible social incentives.[54]

Present bias. In the case of present bias, it is not obvious whether the welfare criterion should be based on the household "experienced" or "decision" utility. Under the first criterion, the key difference from the benchmark framework is the presence of an internality, as the household misestimates the loss in its utility from reducing electricity use in period 1. This term is difficult to estimate, but it likely increases the welfare effect of corrective policies in our setting and thus the bias from assuming away hysteresis. Indeed, if anything, Panel C in Table 4 suggests that households were ex-ante overestimating their loss from reducing electricity use. Under the second

[54] These answers are also consistent with a model in which households understood the private costs of changing their behavior, but in which these costs drop when everybody is subject to a policy because of, e.g., social learning (Dupas, 2014). The welfare effect in both the short run and the long run would be of course even larger in this case. However, the results in Section 4.2, using the idiosyncratic variation in quotas, show that the amount of information among one's peers (or the information due to media campaigns) cannot rationalize all our findings.

criterion, the key difference is that the long run should be further discounted in welfare calculations. For instance, if the household puts twice as much weight on the short run than on the long run ($\alpha = 1/2$), the bias from assuming away hysteresis would be cut in half for both short-run and long-run policies. Even in this extreme example, however, the bias would remain sizable.

Myopic beliefs. In the case of a household that is myopic about the impact of its choices on its future propensity to consume, the relevant welfare criterion is based on the household's "experienced" utility in our setting. Indeed, Panel C in Table 4 indicates that households were unlikely to be making a mistake by consuming less electricity after the crisis. Under this criterion, there are two additional welfare effects for both short-run and long-run policies. First, as in the case of a present bias, the household misperceives its utility loss from reducing electricity use in period 1, holding fixed its plans regarding period-2 choices. This again likely increases the welfare effect of corrective policies in our setting and thus the bias from assuming away hysteresis. Second, the household experiences a gain in private welfare from reducing electricity use in period 2 below the level it previously expected. Setting expectations at counterfactual levels ($\frac{x_t}{(1-.12)}$) and using a linear approximation for the demand curve and the same elasticity as above, this welfare gain amounts to US\$2.848 billions.[55] Although merely illustrative, this figure shows that the bias from assuming away hysteresis can considerably increase in a model with myopic households.

Social incentives. We have shown that the sufficient statistics formulas include an additional term when policies involve social incentives, namely the direct effect of the social incentives on the household utility. This term appears in the formulas whether or not one assumes away hysteresis, however. So the fact that social incentives may have played a role in our setting does not affect the welfare implications of hysteresis. Yet, it may have affected the welfare effect of the policy; for instance, Panel C in Table 4 suggests that it may have mitigated the loss in consumer surplus. We have also shown that the welfare implications of hysteresis differ if the persistent effect of a short-run policy is due to persistent changes in social incentives. However, this could not rationalize the broadly comparable degree of hysteresis estimated using the variation in quotas among movers.

Producer surplus. Finally, we have shown that hysteresis implies an additional welfare loss for the short-run policy if electricity is not priced at a constant marginal cost. The model with an upward-sloping supply curve unlikely applies to our setting. The marginal cost of electricity supply is likely flat, and we found no evidence for an impact of the short-run policy on electricity prices.[56] Yet, the consumer price is likely above the marginal cost because the regulatory framework factors fixed costs into marginal electricity prices. We do not have an estimate of this mark-up

[55]This is calculated as: $\frac{1}{2}\sum_{t=10}^{163}\beta^t \cdot \left[(.12 \cdot \frac{x_t}{(1-.12)}) \cdot (\frac{-.12}{-.31} \cdot p_t)\right]$, where $.12 \cdot \frac{x_t}{(1-.12)}$ is the reduction in electricity use below expectations, and $\frac{-.12}{-.31} \cdot p_t$ captures the impact on the marginal utility from electricity use in period 2.

[56]Electricity generation comes mostly from hydropower. The assumption of a fixed marginal cost is also common for welfare calibrations in the literature (Allcott and Kessler, forthcoming; Ito, Ida and Tanaka, 2017).

that would allow us to quantify the long-run loss in producer surplus, but we can use our results to derive bounds for the mark-up such that the bias from assuming away hysteresis remains positive. For instance, if the only additional source of welfare gain is a long-run correction of the externality valued at US$874 million, the bias would become negative with a mark-up of only 6.42%.[57] Therefore, the welfare implications of hysteresis for the short-run policy may in fact be negative in our setting unless welfare gains from the long-run correction of possible internalities were also sizable. The welfare implications are unchanged, however, for the long-run policy.

7 Conclusion

This paper combined theory and evidence to study the implications of hysteresis for the welfare evaluation of corrective policies. We derived sufficient statistics formulas to show how estimates of the persistent effect of a short-run policy could inform the welfare evaluation of both short-run and long-run policies. We then estimated such an effect and used it to illustrate the welfare implications of hysteresis for corrective policies in the context of energy use in a developing country setting.

The sufficient statistics formulas that we derived are not specific to our application, and a number of studies have provided evidence of hysteresis in a wide range of settings and behaviors. As such, these formulas could be used to evaluate the implications of their estimates for the welfare evaluation of corrective policies in their setting. These formulas could also inform the design of future studies that aim to understand the welfare implications of hysteresis in other contexts.

Our setting is also particularly useful. There is a lot of interest in policies aimed at large and persistent changes in behaviors in the context of energy use. However, such policies are rarely implemented, partly because of the severe loss in private surplus that they would supposedly create. The policy that we study, which was only implemented because of a dramatic energy crisis, thus gives us a rather unique opportunity to study the consequences of policies aimed at changes in behaviors of this magnitude. In particular, it allows us to show that failing to take hysteresis into account may greatly overestimate the loss in private surplus for policies aimed at large long-run changes in energy use. Moreover, that we find persistent effects for all distribution utilities, which differ in the characteristics of their local demand, and for millions of customers of one distribution utility brings some reassurance regarding the external validity of our evidence and conclusions. An exciting question that naturally arises from these findings, but that is beyond the scope of this paper, is which policies would best induce persistent changes in households' choices at an earlier stage of development to limit the growing energy demand in the developing world.

[57] This mark-up μ is such that the long-run loss in producer surplus is equal to the value of the long-run correction of the externality: $\sum_{t=2003}^{2014} \beta^{t-2001} \cdot (.12 \cdot \frac{x_t}{(1-.12)}) \cdot (\mu \cdot p_t) = \$874,000,000$.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association*, 105(490): 493–505.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2012. "The Environment and Directed Technical Change." *American Economic Review*, 102(1): 131–166.

Acland, Dan, and Matthew R Levy. 2015. "Naiveté, projection bias, and habit formation in gym attendance." *Management Science*, 61(1): 146–160.

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally–Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review*, 105(8).

Allcott, Hunt, and Judd Kessler. forthcoming. "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons." *American Economic Journal: Applied Economics*.

Allcott, Hunt, and Sendhil Mullainathan. 2010. "Behavior and energy policy." *Science*, 327(5970): 1204–1205.

Allcott, Hunt, and Todd Rogers. 2014. "The Short–Run and Long–Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review*, 104(10): 3003–3037.

ANEEL. 2005. "Tarifas de fornecimento de energia elétrica." *Agência Nacional de Energia Elétrica, Cadernos Temáticos ANEEL*, 4.

Ashenfelter, Orley, and Michael Greenstone. 2004. "Using Mandated Speed Limits to Measure the Value of a Statistical Life." *Journal of Political Economy*, 112: 226–267.

Bardelin, Cesar. 2004. "Os efeitos do racionamento de energia elétrica ocorrido no Brasil em 2001 e 2002 com ênfase no consumo de energia elétrica." *Dissertação de Mestrado em Sistemas de Potência, Universidade de São Paulo: Escola Politécnica*.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. 2016. "Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century." *Journal of Political Economy*, 124(1): 105–159.

Becker, Gary, and Kevin Murphy. 1988. "A Theory of Rational Addiction." *Journal of Political Economy*, 96(4): 675–700.

Borenstein, Severin. 2009. "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing." *Mimeo, Energy Institute, University of California at Berkeley*.

Brandon, Alec, Paul Ferraro, John List, Robert Metcalfe, Michael Price, and Florian Rundhammer. 2017. "Do The Effects of Social Nudges Persist? Theory and Evidence from 38 Natural Field Experiments." *NBER Working Paper*, 23277.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Cameron, C., J. Gelbach, and D. Miller. 2008. "Bootstrap-Based Improvements for Inference With Clustererd Errors." *Review of Economics and Statistics*, 90(3): 414–427.

Charness, Gary, and Uri Gneezy. 2009. "Incentives to Exercise." *Econometrica*, 77(3): 909–931.

Davis, Lucas W, Alan Fuchs, and Paul Gertler. 2014. "Cash for coolers: evaluating a large–scale appliance replacement program in Mexico." *American Economic Journal: Economic Policy*, 6(4): 207–238.

de Oliveira, Luiz Antônio Pinto, and Antônio Tadeu Ribeiro de Oliveira. 2011. *Reflexões sobre os deslocamentos populacionais no Brasil.* IBGE.

Dixit, Avinash. 1989. "Entry and Exit Decisions under Uncertainty." *Journal of Political Economy*, 97: 620–638.

Dubin, Jeffrey, and Daniel McFadden. 1984. "An econometric analysis of residential electric appliance holdings and consumption." *Econometrica*, 345–362.

Dupas, Pascaline. 2014. "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence From a Field Experiment." *Econometrica*, 82(1): 197–228.

Farhi, Emmanuel, and Xavier Gabaix. 2017. "Optimal taxation with behavioral agents." *NBER Working Paper*, 21524.

Ferraro, Paul J, and Michael K Price. 2013. "Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment." *Review of Economics and Statistics*, 95(1): 64–73.

Fujiwara, Thomas, Kyle Meng, and Tom Vogl. 2016. "Estimating Habit Formation in Voting." *American Economic Journal: Applied Economics*, 8(4): 160–188.

Giné, Xavier, Dean Karlan, and Jonathan Zinman. 2010. "Put your money where your butt is: a commitment contract for smoking cessation." *American Economic Journal: Applied Economics*, 213–235.

Greenstone, Michael, and Rema Hanna. 2014. "Environmental regulations, air and water pollution, and infant mortality in India." *The American Economic Review*, 104(10): 3038–3072.

Hintermann, Beat, and Andreas Lange. 2013. "Learning abatement costs: On the dynamics of the optimal regulation of experience goods." *Journal of Environmental Economics and Management*, 66(3): 625–638.

Imbens, Guido, and Karthik Kalyanaraman. 2012. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." *The Review of Economic Studies*, 1–27.

IPCC. 2014. "Climate Change 2014: IPCC Fifth Assessment Report."

Ito, Koichiro. 2015. "Asymmetric incentives in subsidies: Evidence from a large-scale electricity rebate program." *American Economic Journal: Economic Policy*, 7(3): 209–37.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2017. "The Persistence of Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand." *American Economic Journal: Economic Policy*.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417–1438.

Kelman, Jerson. 2001. "Desequilíbrio entre Oferta e Demanda de Energia Elétrica." *Relatório da Comissão de Análise do Sistema Hidrotérmico de Energia Elétrica*.

Kling, J., J. Liebman, and L. Katz. 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica*, 75(1): 83–119.

Larcom, Shaun, Ferdinand Rauch, and Tim Willems. 2017. "The benefits of forced experimentation: striking evidence from the London underground network." *The Quarterly Journal of Economics*, 132(4): 2019–2055.

Levinson, Arik. 2014. "California energy efficiency: Lessons for the rest of the world, or not?" *Journal of Economic Behavior and Organization*, 107: 269–289.

Mation, Lucas, and Claudio Ferraz. 2011. "How Do Firms React to Infrastructure Constraints? Evidence from Brazil's Energy Shortage." *Working paper, Department of Economics, Pontifícia Universidade Católica do Rio de Janeiro*.

Matsuura, Kenji, and Cort Willmott. 2012. "Terrestrial precipitation: 1900-2010 gridded monthly time series (1900-2010)(v 3.01 added 6/14/12). University of Delaware." *http://climate.geog.udel.edu/ climate/*.

Maurer, Luiz, Mario Pereira, and José Rosenblatt. 2005. "Implementing Power Rationing in a Sensible Way: Lessons Learned and International Best Practices." *Energy Sector Management Assistance Program Report, World Bank*.

Meier, Alan. 2005. "Saving Electricity in a Hurry: Dealing with Temporary Shortfalls in Electricity Supplies." *OECD/IEA*.

Miller, Conrad. 2016. "The Persistent Effect of Temporary Affirmative Action." *American Economic Journal: Applied Economics*, 9(3): 152–190.

ONS. 2011. "Plano Decenal de Energia." *Tech. report, Operador Nacional do Sistema Eletrico.*

Pigou, Arthur. 1920. *The Economics of Welfare.* London: McMillan & Co.

Pimenta, Amanda, Hilton Notini, and Luiz Felipe Maciel. 2009. "Brazilian Electricity Demand Estimation: What Has Changed after the Rationing in 2001? An Application of Time Varying Parameter Error Correction Model." *Mimeo, Getulio Vargas Foundation*.

Reiss, Peter, and Matthew White. 2008. "What Changes Energy Consumption? Prices and Public Pressures." *RAND Journal of Economics*, 39(3): 636–663.

Sallee, James. 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics*, 57(3): 781–820.

Sandmo, A. 1975. "Optimal Taxation in the Presence of Externalities." *Swedish Journal of Economics*, 77(1): 86–98.

Wolfram, Catherine, Orie Shelef, and Paul Gertler. 2012. "How Will Energy Demand Develop in the Developing World?" *Journal of Economic Perspectives*, 26: 119–138.

Figure 1: The impact of the temporary energy-saving program on residential electricity use (raw data)

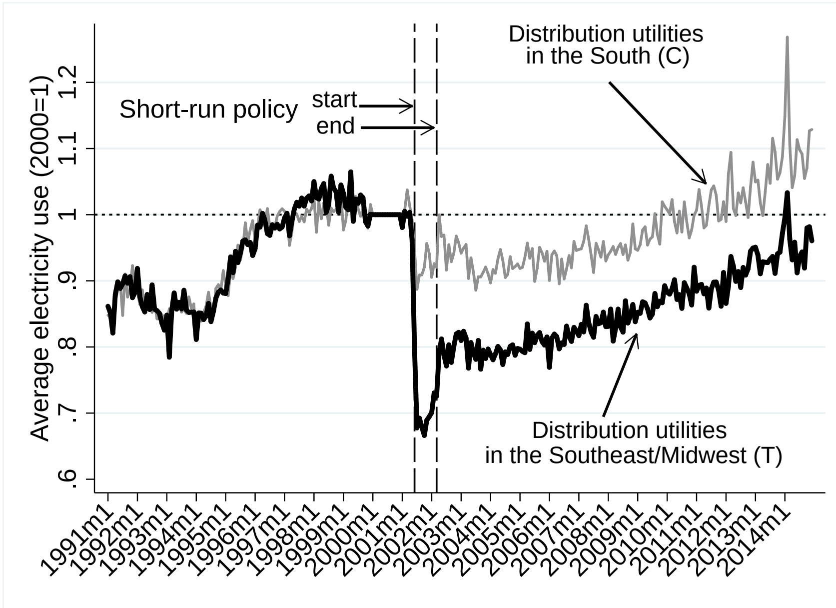

The figure displays the overall impact of the energy-saving program on monthly average residential electricity use per customer for distribution utilities in the Southeast/Midwest where the policy was implemented (Treatment group), and in the South where the policy was not implemented (Control group). We use the raw administrative data for each distribution utility and present unweighted averages in each month, normalized with respect to the same month in 2000 to net out the seasonality. Trends were similar prior to the start of the policy in June 2001. Average consumption then dropped, particularly for distribution utilities in the Southeast/Midwest (there were some spillovers of the policy for those in the South). It rebounded after February 2002, once the policy ended, but only partially. Comparing patterns in the Southeast/Midwest and in the South in later years suggests that the policy had a stable and persistent impact on average residential electricity use, until at least the end of our sample (December 2014).

Figure 2: The loss in private welfare from corrective policies with and without hysteresis

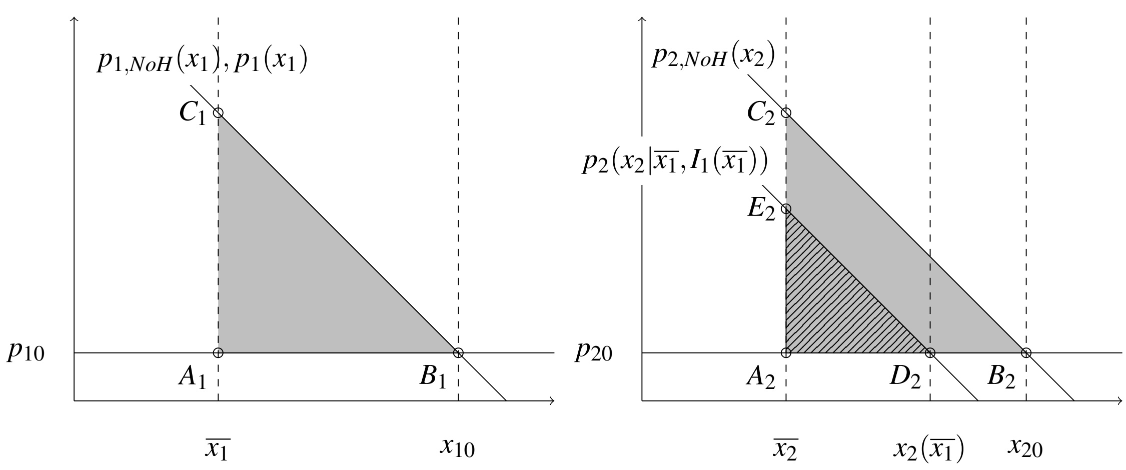

(a) Period 1 (b) Period 2

The figure illustrates the loss in private welfare in the benchmark framework (the private welfare corresponds to the consumer surplus) for the short-run and the long-run corrective polices. For simplicity, we assume linear demand curves and no change in the slope of the demand curve following a short-run policy in the presence of hysteresis. The triangles $A_1B_1C_1$ illustrates the loss in private welfare for the short-run policy with and without hysteresis. The demand curve $p_2(x_2|\overline{x_1}, I_1(\overline{x_1}))$ illustrates that the demand curve shifts inward in period 2 following a short-run policy in the presence of hysteresis. The triangles $A_1B_1C_1$ and $A_2B_2C_2$ illustrate the loss in private welfare for the long-run policy when assuming away hysteresis. The triangles $A_1B_1C_1$ and $A_2D_2E_2$ illustrate the comparable loss in private welfare when allowing for hysteresis. The trapezoid $D_2B_2C_2E_2$ illustrates the first source of bias from assuming away hysteresis for the long-run policy. See text for more details.

Figure 3: Causes of the electricity supply crisis and its temporary energy-saving program

(a) Stocked energy of the hydroelectric reservoirs

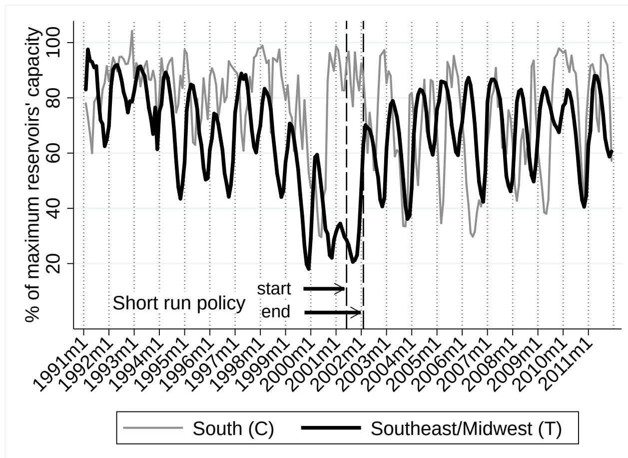

(b) Flow into the reservoirs (Southeast/Midwest)

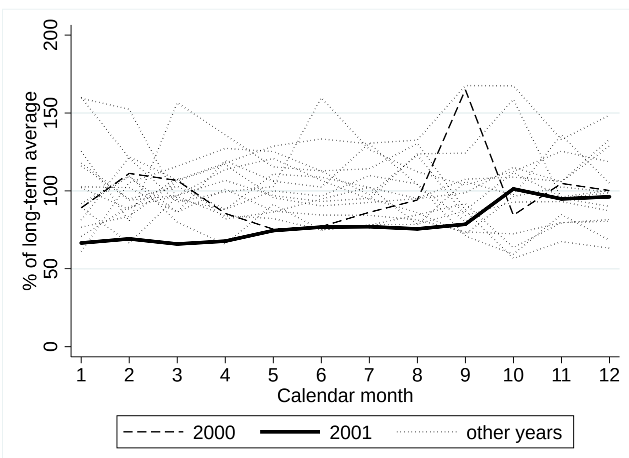

(c) Flow into the reservoirs (South)

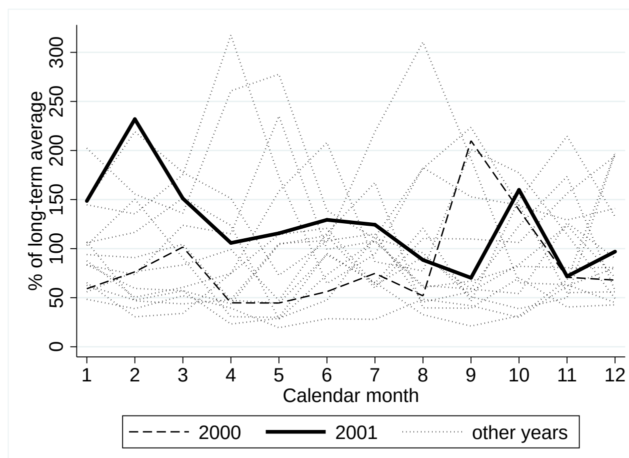

Panel (a) displays the evolution of the water level in the reservoirs of hydroelectric power plants in the Southeast/Midwest and in the South in percentage of their maximum capacity. Official data from ONS (the National System Operator); dotted lines indicate January in each year. *Panels (b) and (c)* display the streamflow level of the rivers serving the reservoirs in these two subsystems in percentage of the long-term average streamflow levels in each month in each subsystem. Panel (a) shows that reservoirs' water levels followed a clear seasonal pattern in the Southeast/Midwest before the crisis, with heavy rain replenishing the reservoirs downstream during the austral summer. Water levels were low in the two subsystems by 2000. Panels (b) and (c) then show that the crisis and the differential treatment between subsystems were due to a weather shock (and not to a demand shock). Streamflow levels were higher than average in both regions around September 2000. This is a period of low streamflow in the Southeast/Midwest but of high streamflow in the South. As a result, reservoirs were rapidly replenished in the South but not in the Southeast/Midwest. The beginning of every year is a period of high streamflow in the Southeast/Midwest. However, at the beginning of 2001, streamflow levels were much lower than the average for those months. As a result, reservoirs' water levels did not increase in the Southeast/Midwest, as they usually do at the beginning of every year. In March 2001, at the end of the rainy season, reservoirs' levels reached their lowest point in 40 years (for the season). This caused the stress of the system. In contrast, streamflow levels were higher than average over the same period in the South, and reservoirs' water levels remained high. After the crisis, the government heavily invested in transmission capacity to reduce the risk of electricity shortages. As a result, the South increased its supply to the rest of the country and panel (a) shows that water levels in the South started to follow the seasonal pattern observed in the Southeast/Midwest.

Figure 4: Treatment effect of the temporary energy-saving program on average residential electricity use

(a) Difference-in-differences (overall effect) (b) Synthetic control methods (utility-specific effect)

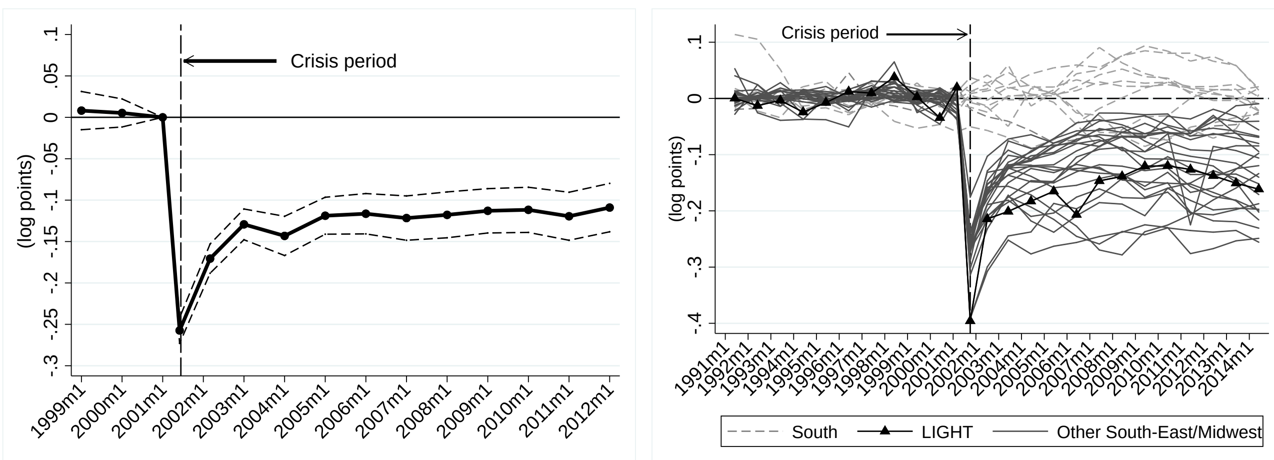

Panel (a) displays the estimated $\widehat{\delta_p}$ and their 95% confidence intervals using the specification in equation (10). The estimates during and after the crisis capture the average treatment effect of the policy on the treated under a common-trend assumption; the estimates in preceding periods test for the presence of a common-trend prior to the crisis. Monthly observations are divided into yearly periods before 2001 and after 2002; we divide 2001 and 2002 into three periods: pre-crisis (early 2001), crisis (June 2001-February 2002), and post-crisis (rest of 2002). The pre-crisis period (early 2001) is the reference period. The regression includes fixed effects for each distribution utility, for each period, and for each month per subsystem (e.g., June in the Southeast/Midwest), as well as the following controls (in log) at the distribution utility level in each year: main residential electricity tariff, population size, GDP, formal employment levels, and average temperature. Standard errors are clustered by distribution utility (43 clusters). *Panel (b)* displays the estimated effect for each distribution utility using synthetic control methods, as detailed in section 3.4. Darker lines correspond to utilities in the Southeast/Midwest. Lighter lines correspond to placebo estimates in which we compare a given control distribution utility to a weighted average of the other ones. The synthetic controls can closely match the trends pre-crisis. The estimated short-run effect is large for all the treated distribution utilities at between -.18 and -.40 log points. The long-run effect is also negative for all these utilities. Both the median and the average of these effects are equal to -.14 log points in 2014. The median and the average effect of the placebo estimates are close to 0 in all years.

Figure 5: Anatomy of the reduction in average residential electricity use using household-level billing data

(a) Comparing time-series in average residential electricity use in the utility-level data and the microdata

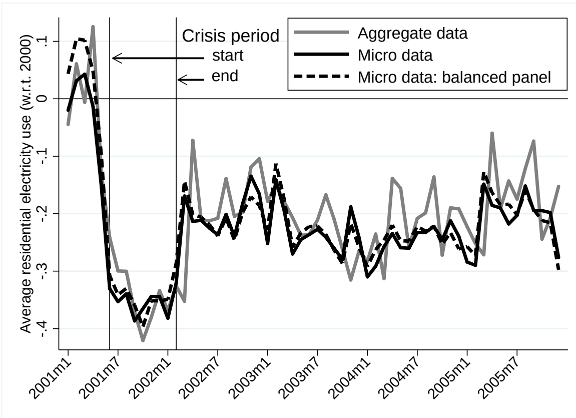

(b) Distribution of average monthly electricity use over time for a balanced panel of customers

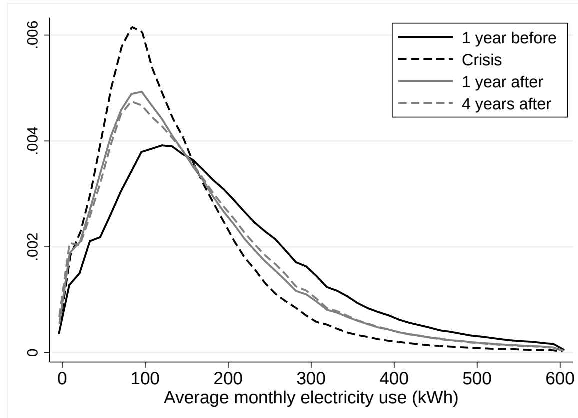

(c) Distribution of changes in average monthly electricity use for a balanced panel of customers with the same quota level (300 kWh/month)

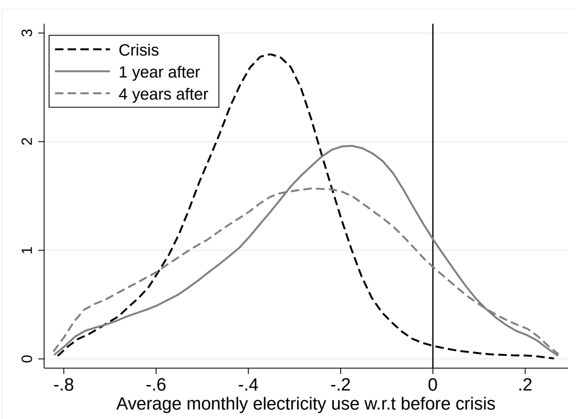

(d) Correlation between changes in electricity use during and after the crisis for a balanced panel of customers with same quota level (300 kWh/month)

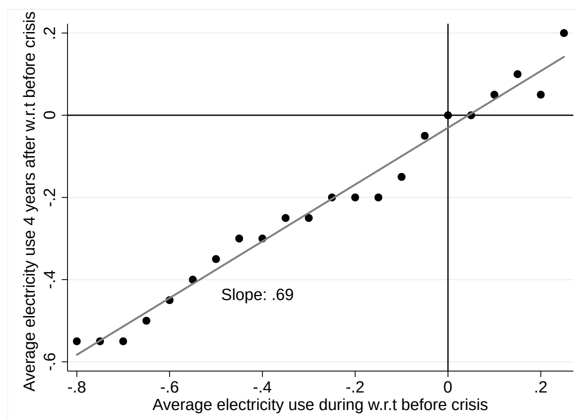

The figure uses individual monthly billing data for the universe of residential customers of LIGHT, a distribution utility subject to the energy-saving program during the crisis, from January 2000 to December 2005. *Panel (a)* displays the average electricity use for LIGHT customers in each month compared to the same month in 2000. It shows that the time-series is almost identical when (i) we use the aggregate data at the distribution utility level used in previous figures, (ii) the microdata at the household level in each month, and (iii) a subsample of the microdata with those customers metered regularly from 2000 to 2005 (balanced panel; 1,185,146 customers). This provides additional evidence that composition effects, absent from the balanced panel by construction, are unlikely to drive our results. *Panel (b)* shows that the average changes in electricity use came from sizable reductions at every level of consumption. It uses the same balanced panel to investigate changes in the distribution of electricity use over time. It displays kernel densities for average monthly electricity use before, during, and after the crisis. Kernel densities are based on data from July to December, such that we can compare consumption levels up to four years after the crisis. The density during the crisis is first-order stochastically dominated by the other densities. Densities one year and four years after the crisis are very similar and fall exactly between the crisis and pre-crisis densities. *Panel (c)* shows that the average changes in electricity use came from large reductions from most customers at a given quota level. It displays the distribution of changes in electricity use during and after the crisis compared to the same months before the crisis, for a subset of the sample in panel (b), in which customers had about the same quota (10% above and below 300 kWh/month; 83,966 customers), and thus about the same incentive schedule during the crisis. Kernel densities are based on electricity use during the first five months of the crisis (and in the same months in other years), before any change in quotas. During the crisis, 98% reduced electricity use and the median customer reduced usage by -34%. Four years after the crisis, 78% were still using less electricity than before the crisis; the median customer was using 22% less electricity. *Panel (d)* displays the correlation between individual changes in electricity use during and after the crisis compared to the same months before the crisis, for the sample used in panel (c). Customers are averaged by bins of 5% changes in electricity use during the crisis. The strong correlation (.69) provides household-level evidence that changes in electricity use that took place during the crisis were persistent. Kernel densities use Epanechnikov kernels and optimal bandwidths. We show similar patterns as in panels (c) and (d) for other baseline consumption levels in Appendix L.

Figure 6: Seasonality in residential electricity use and variation in the quota of "movers"

(a) Seasonality in electricity use (balanced panel) and the average quota of movers by moving month

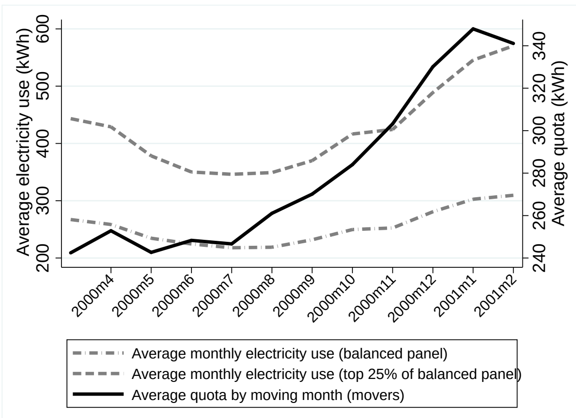

(b) Distribution of quotas for winter (May 2000-Jul 2000) vs. summer movers (Dec 2000-Feb 2001)

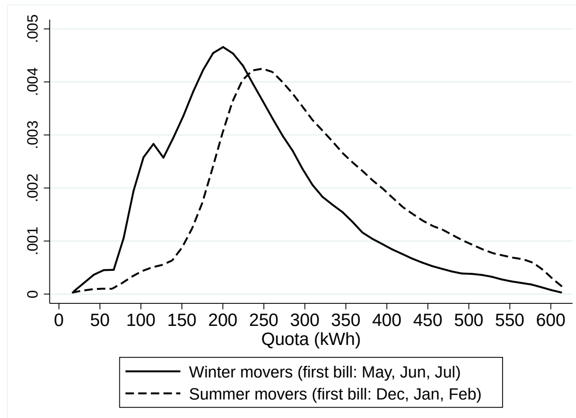

The figure shows how seasonality in residential electricity use created quasi-experimental variation in quotas among LIGHT customers whose first bill was sent between March 2000 and February 2001 ("movers"). The quota was set at 80% of their average consumption between May and July 2000 for those households that moved into their housing unit before May 2000, but at 80% of the average consumption in their first 3 billing months for those households that moved after May 2000. Because of strong seasonality in electricity use, these rules created wide variation in quotas among movers entirely because of differences in moving dates. Panel (a) shows the seasonality in electricity use for LIGHT customers. It displays the average consumption in each month, between March 2000 and February 2001, for the overall balanced panel of customers in Figure 5a and for the top quartile of the distribution in this sample. Consumption levels are higher in the austral summer and lower in the austral winter. The difference reaches more than 60% in the top quartile of the distribution, showing that the seasonality is particularly strong among large consumers. Panel (a) also shows the impact of the seasonality in electricity use on the movers' quotas. It displays the average quota of customers who moved into their housing unit in the same month, as measured by the date of their first electricity bill. The sample is restricted to large consumers given the stronger seasonality in their electricity use, that is movers whose average consumption prior to the crisis (March-May 2001) was in the top quartile of the distribution (66,037 customers; we also restrict attention to those observed at least until December 2002). Customers who moved into their housing unit before May 2000, whose quota was based on their consumption in the typical baseline period, had similar average quotas. In contrast, the quota of customers who moved in after May 2000, which was based on their consumption in their first 3 billing months, closely follows the seasonality. *Panel (b)* shows the distribution of quotas for movers who moved in the winter (May 2000-July 2000; 13,557 customers) and in the summer (December 2000-February 2001; 11,725 customers). The quota distribution for summer movers first-order stochastically dominates the distribution for winter movers; at the median, the difference reaches 32%, which is equivalent to offering a non-binding quota to summer movers. Kernel densities use Epanechnikov kernels and optimal bandwidths.

Table 1: Descriptive statistics for distribution utilities in the four subsystems

	Descriptive statistics in 2000 Mean [min-max]					**Differential trends 2010 vs. 2000** Coefficient, log points (s.e.)		
	South (1)	LIGHT (2)	Southeast/Midwest (3)	Northeast (4)	North (5)	SE/MW vs. S (6)	NE vs. S (7)	N vs. S (8)
Average residential electricity	166	226	190	107	180	-.113***	.001	-.004
use (kWh/month)	[133–191]		[122–261]	[72.7–129]	[128–267]	(.022)	(.037)	(.051)
Main residential electricity	.179	.216	.193	.179	.177	-.093	-.001	-.083
tariff (US$/kWh)	[.158–.201]		[.166–.222]	[.165–.202]	[.152–.194]	(.062)	(.065)	(.098)
Number of customers	360	2864	861	834	238	.012	.146***	.187***
(1000's)	[1.63–2200]		[17.7–4160]	[63–2405]	[12.2–857]	(.041)	(.045)	(.042)
Population size	1457	9025	3235	4323	1605	.031	.027	.164***
(1000's)	[12.1–9208]		[64.8–16661]	[193–13014]	[124–6122]	(.03)	(.028)	(.047)
Share of households	.982	.999	.984	.894	.83	-.002	.081***	.109***
with electricity	[.949–1]		[.896–1]	[.759–.989]	[.645–.989]	(.007)	(.028)	(.037)
Share of households	.795	.991	.861	.72	.734	-.02	-.01	-.003
in urban areas	[.587–.916]		[.654–.991]	[.613–.851]	[.414–.979]	(.017)	(.017)	(.025)
Median household income	737	944	761	342	496	-.104**	.081*	-.057
(US$/month)	[507–944]		[450–1180]	[282–413]	[354–802]	(.045)	(.043)	(.082)
Share of households	.937	.972	.921	.646	.708	.009	.258***	.148***
with refrigerator	[.828–.994]		[.806–.985]	[.538–.758]	[.522–.916]	(.018)	(.037)	(.056)
Share of households	.436	.554	.366	.089	.199	-.008	.369***	.06
with washing machine	[.145–.661]		[.1–.618]	[.039–.137]	[.075–.311]	(.089)	(.092)	(.118)
Share of households	.584	.479	.58	.61	.543	.002	-.09*	-.041
with TV	[.477–.682]		[.45–.701]	[.511–.666]	[.436–.64]	(.053)	(.048)	(.047)
Share of households	.071	.311	.06	.044	.125	n.a.	n.a.	n.a.
with air conditioner	[.005–.183]		[.008–.311]	[.014–.08]	[.047–.228]			
Observations	17	1	26	11	8	86	56	50

The table uses utility-level administrative data for distribution utilities in the four subsystems in 2000 and 2010, and census data matched to the concession area of these distribution utilities in the same years (the 2010 census does not have information on air conditioners). Columns (1)-(5) display descriptive statistics in 2000 (prior to the crisis) for the variables listed in the left-hand side column for distribution utilities in the South (column 1), in the Southeast/Midwest (LIGHT utility in column 2, all distribution utilities in column 3), in the Northeast (column 4), and in the North (column 5). Columns (6)-(8) display estimates of a long-term difference-in-differences estimator comparing the logarithm of these variables in 2010 vs. 2000 for distribution utilities in the Southeast/Midwest (column 6), in the Northeast (column 7), and in the North (column 8) compared to distribution utilities in the South. Significance levels: *10%, **5%, ***1% (s.e. clustered by distribution utility). Regressions include fixed effects for each distribution utility and each year. All monetary values are expressed in US$ from 2012. A similar table with additional variables is provided in Appendix A. We argue in the text that the information in this table shows that distribution utilities in the South constitute a credible control group for distribution utilities in the Southeast/Midwest, but not for those in the other two subsystems.

Table 2: Difference-in-differences results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treat × Crisis	-.2641***	-42.78***	-.2615***	-.2785***	-.2482***	-.2448***	-.2714***	-.2612***
	(.0085)	(2.392)	(.0085)	(.015)	(.0099)	(.0075)	(.0116)	(.0102)
Treat × Postcrisis until 2005	-.1442***	-24.58***	-.1406***	-.16***	-.1474***	-.1636***	-.1545***	-.1784***
	(.0092)	(2.109)	(.0102)	(.0145)	(.0094)	(.0078)	(.0119)	(.0125)
Treat × Postcrisis after 2005	-.1227***	-20.64***	-.1233***	-.1511***	-.1318***	-.1204***	-.1297***	-.134***
	(.0137)	(2.366)	(.0152)	(.0274)	(.0141)	(.0132)	(.0168)	(.0186)
Treat × Time trend × Before crisis						-.0005		.0001
						(.0004)		(.0002)
Treat × Time trend × Crisis						-.0025**		-.0036***
						(.0012)		(.0013)
Treat × Time trend × Postcrisis until 2005						.0012***		.0009***
						(.0002)		(.0003)
Treat × Time trend × Postcrisis after 2005						.0002		0
						(.0002)		(.0001)
Estimated 10-year effect	-.1227***	-20.64***	-.1233***	-.1511***	-.1318***	-.1078***	-.1297***	-.1334***
	(.0137)	(2.366)	(.0152)	(.0274)	(.0141)	(.0151)	(.0168)	(.018)
Number of distribution utilities	43	43	43	43	43	43	35	35
Specification in levels	No	Yes	No	No	No	No	No	No
Exclude outliers	No	No	Yes	No	No	No	No	No
Weights by customer base	No	No	No	Yes	No	No	No	No
Only winter months	No	No	No	No	Yes	No	No	No
Include time trends	No	No	No	No	No	Yes	No	Yes
Years 1991 to 2014	No	No	No	No	No	No	Yes	Yes

The table displays the results from estimating variants of the difference-in-differences specifications in equations (11) and (12). It summarizes and evaluates the robustness of our results in Figure 4a to several specification and sample checks, as detailed in Section 3.3 (the number of distribution utilities is slightly smaller in columns (7) and (8) because a few distribution utilities were split or created after 1991). The specifications in columns (1)-(6) use the same sample and the same set of controls used in Figure 4a. The specifications in columns (7)-(8) use the full sample from 2011 to 2014, and thus do not include any controls. In all specifications, we find that the short-run policy led to large reductions in average residential electricity use in the short-run, and that about half of the short-run effect persisted in the long run. Significance levels: *10%, **5%, ***1% (s.e. clustered by distribution utility).

Table 3: The impact of quasi-experimental variation in quotas among movers on their electricity use

	Log Quota	Log Consumption	Log Consumption July-December						
		Mar-May 2001	2001		2002	2003	2004	2005	2003–2005
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. *Average effect*									
Log Moving-Week Quota	.785***	.000	.137***	.138***	.087***	.051*	.054**	.041	.049**
	(.223)	(.023)	(.034)	(.039)	(.024)	(.029)	(.027)	(.034)	(.024)
Panel B. *Median effect*									
Log Moving-Week Quota	.662***	.018	.115***	.153***	.087***	.093***	.073***	.088***	.086***
	(.028)	(.023)	(.02)	(.013)	(.023)	(.023)	(.023)	(.03)	(.013)
Observations	38,207	38,207	38,207	38,207	38,207	38,207	38,207	38,207	114,621
Log Cons. Mar-May 2001	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	No	No	No	No	No	No	No	No	Yes

The table shows the effect of quasi-experimental variation in quotas among LIGHT customers whose first bill was sent between March 2000 and February 2001 ("movers") on their electricity use in different periods. The sample is restricted to those movers whose average consumption in the 3 months prior the crisis was in the top quartile of the distribution, and who are observed regularly until December 2005. The quota assignment rules generated variation in quotas entirely because of differences in moving dates in this sample (seasonality, see text). The table displays the coefficients from regressing different outcomes (listed above each column) on the variable that captures the variation in quotas from differences in moving dates; it is the logarithm of the average of the quota of all movers (excluding i) who received their first bill in the same week as household i ("Log Moving-Week Quota"). Panel A use OLS regressions including neighborhood fixed effects. Panel B uses quantile regressions (median effect). Significance levels: *10%, **5%, ***1% (standard errors clustered by moving week in parentheses – 50 clusters – and obtained using 1000 bootstrap replications for the quantile regressions). The average quota of same-week movers strongly predicts variations in the movers' own quotas (column 1) but not on their average consumption levels between March and May 2001, prior to the crisis (column 2). We find, however, that the average quota of same-week movers is positively correlated with consumption levels in the first six months of the crisis (column 3). The effect is larger when we control for pre-crisis consumption levels (in log) to absorb part of the household heterogeneity (column 4; we do so in the next columns as well). The effect on consumption levels in the same months decrease in 2002 (column 5). The point estimates become more stable in the following years (columns 6 to 8), but they are noisier. In column 9, we then consider a specification that pools the years between 2003 and 2005 (adding year fixed effects).

Table 4: Qualitative evidence from household surveys conducted in the Southeast/Midwest

	Main domestic appliances					Other domestic appliances		
	Electric shower	Refrigerator	Freezer	Light	TV	Air Conditioner	Washing machine	Microwave
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. *Main sources of electricity use prior to the crisis (1999 survey; N=6,482)*								
Average quantity before the crisis	.97	.99	.20	8.45	1.39	.10	.53	.21
Average kWh/month before the crisis	58.14	41.71	7.88	42.54	15.63	2.78	3.37	2.97
Panel B. *Reported mechanisms of long-run changes in electricity use (2005 survey)*								
Conditional on owning a given type of appliance before the crisis, share of respondents who report that they:								
(1) Use appliance as much as before the crisis	.60	.90	.61	.51	n.a.	.32	.44	.56
(2) Use appliance less than before the crisis	.39	.07	.21	.41	n.a.	.60	.55	.41
(3) Disconnected or disposed of the appliance	.01	0	.17	0	n.a.	.05	.01	.03
(4) Substituted a more energy-efficient model	0	.03	.01	.08	n.a.	.03	0	0
Number of respondents	*4,225*	*4,432*	*784*	*4,500*		*257*	*3,004*	*1,399*

	Share of respondents
Panel C. *Households' reported life quality during the crisis (2005 survey)*	
If your consumption reduction was sufficient to meet your quota, how difficult was it? (N=3375)	
It was very difficult (*muito*)	.09
It was not so difficult (*pouco*)	.48
It was not difficult at all (*nenhuma*)	.43
How do you evaluate your change in life quality caused by the energy-saving program? (N=4376)	
I did not experience any change in life quality (*não houve variação*)	.48
I experienced some discomfort (*causou desconforto*)	.20
I experienced some severe discomfort (*causou muito desconforto*)	.08
I learned to live with the same comfort while saving money (*aprendi a viver com o mesmo conforto economizando dinheiro*)	.24

The table uses household-level data for 8 distribution utilities in the Southeast/Midwest from the two most recent rounds of the PPH surveys (1998/1999 and 2004/2005). *Panel A* displays the average number per household for the appliances listed on top of each column and the inputted average monthly electricity use in 1999 (before the crisis). The latter is calculated by multiplying the average quantity (from the survey) by the average utilization in 1999 (share of appliances owned frequently in use; also reported in the survey) and by the kWh consumption of the average model of each appliance from PROCEL estimates (see Table A.5 in Appendix A). *Panel B* tabulates the share of households that chose each statement as the best answer to a question about their usage of the appliance after the crisis (conditional on owning the appliance prior to the crisis). *Panel C* tabulates the share of households that chose each statement as the best answer to two questions on their experience during the crisis. In Portuguese, the questions were "*As medidas adotadas para atingir as metas durante o período de racionamento foram suficientes ou mais que suficientes. Dificuldade?*" and "*Como o(a) sr.(a) avalia a variação de qualidade de vida causada pelo racionamento?*." We display the answers in Portuguese in parenthesis.