Online Appendix

The Costs of Misaligned Incentives: Energy Inefficiency and the Principal-Agent Problem

Joshua Blonz

A Refrigerator Policy Background

Refrigerators have historically represented a large portion of household electricity consumption. Over the last 30 years, the consumption of electricity by refrigerators in homes has dropped because of two federal minimum appliance standards that mandated certain levels of performance. For example, before the first standards were enacted, a typical refrigerator manufactured in 1985 consumed 127 kWh/month and represented one-fifth of total household electricity consumption (Energy Information Administration (EIA) 1987). The first national standards were implemented in 1990, causing the same refrigerator to consume only 96 kWh/month. Starting in 1993, the standard was tightened, and a similar model consumed only 65 kWh/month.[1] The standard was strengthened again in 2001, requiring the same type of refrigerator to consume no more than 52 kWh/month. By 2009, refrigerators made up 11 percent of a household's electricity consumption (EIA 2009). Finally, in 2014, the standard was set so that a similar refrigerator would consume only 38 kWh/month.[2] Each new standard was associated with a discontinuous jump in refrigerator efficiency from the previous regime, which suggests that the minimum efficiency standards were binding and affected what refrigerators were manufactured and sold.

In 1996, the US Environmental Protection Agency introduced a voluntary labeling program for refrigerators called Energy Star, which allowed manufacturers to signal to consumers that their products were more energy efficient than the minimum standards. For example, between 2008 and 2014, a refrigerator model would receive an Energy Star certification if it exceeded the standard by 20 percent or more. Energy Star refrigerators have significant penetration in the appliance market. In 2011, 56 percent of the new refrigerators were Energy Star certified (EPA 2011). The replacement refrigerators provided in the ESA program were Energy Star certified.

Figure A1 uses household billing data from participants in the California Energy Savings Assistance (ESA) program to show monthly electricity consumption in the cooler regions of the Southern California Edison (SCE) service territory broken down by the standards regime of their refrigerators. The figure shows how more efficient refrigerators are associated with lower household electricity consumption.[3] The differences in consumption among the various groups illustrate

1. In Southern California, the average household uses 450 kWh/month, which means a refrigerator manufactured before 1993 consumed up to 20 percent of a household's electricity.

2. I report the values for specific refrigerators because the standards are differentiated by size and features. The numbers given reflect the usage for a 20-cubic-foot side-by-side refrigerator all made by the same manufacturer.

3. For example, it is possible that households with newer refrigerators are more likely to have other updated and more efficient appliances.

how the increasingly stringent standards over the last 30 years may have contributed to reduced household electricity consumption.

Figure A1: Household Electricity Consumption by Refrigerator Standard

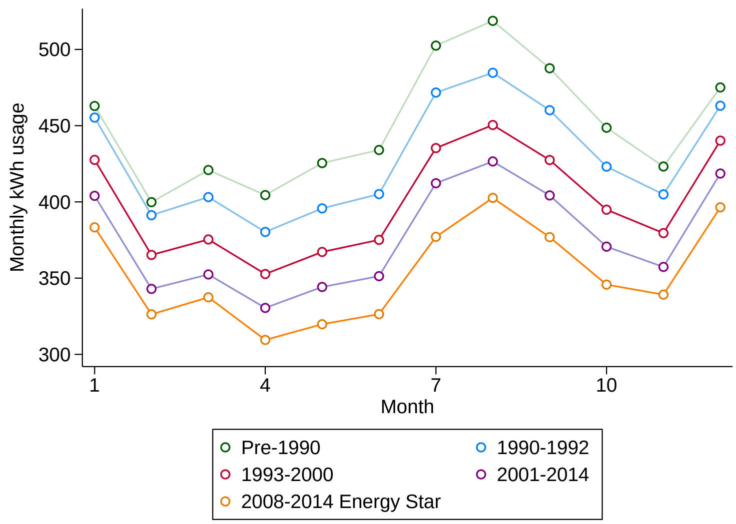

Note: This figure shows the monthly electricity consumption for households grouped by the year the refrigerator in each home was manufactured. Each group corresponds to a different federal minimum efficiency standard. The figure uses ESA program data and monthly billing data from 26,215 households in the cooler coastal climate zones. The figure shows how increasing minimum federal efficiency standards for refrigerators are associated with decreasing household electricity consumption.

The ESA program was purposely designed to replace the least efficient refrigerators—those manufactured before 1993—with new Energy Star models. The goal was to replace models that consumed over 100 kWh/month with Energy Star models that consumed 40 kWh/month or less, providing significant cost-effective savings. Figure A1 shows this potential savings as the difference between the top two consumption lines, both of which qualify for replacement in the ESA program, and the bottom Energy Star average consumption.

B Model

The derivations in Section II are conducted in this appendix. The subsections below set up and mathematically derive the surplus from the integrated and separated task contract. The two contracting approaches are compared in Section II.A.

B.1 Integrated-Task Contract

Under an integrated-task contract, the principal offers the agent a contract to both assess potential recipients and provide the service to those that qualify. When a potential recipient is deemed qualified during the assessment visit, the same agent schedules a second visit to provide the service. For all qualified potential recipients, even if the agent intentionally misreported them as eligible, the integrated-task agent provides the service and receives g. In the integrated-task contract, the agent must pay an "information acquisition cost" (F) to complete her task during the assessment step.

There are two types of potential recipients. A portion α of potential recipients qualify for the service under program rules.[4] The principal and the agent have a shared but noisy expectation about the value of α. The expected number of qualified recipients is $n \cdot \alpha$, and the expected number of unqualified participants is $n \cdot (1-\alpha)$. Subject to participation, the agent chooses a level of misreporting $Z_I \in [0,1]$ to maximize her expected payout, which is a static decision the agent makes at the beginning of the contract.

$$E(\max_{Z_I} \pi_I) = \overbrace{n\alpha(g+d) + n(1-\alpha)(g+d)Z_I}^{\substack{\text{Referral and service}\\ \text{provision payments}}} - \overbrace{n(1-\alpha)C(Z_I)}^{\substack{\text{Cost to agent}\\ \text{of misreporting}}} + \overbrace{n(A-F)}^{\substack{\text{Assessment payment}\\ \text{minus information}\\ \text{acquisition cost}}}. \tag{B1}$$

The I subscript indicates the integrated-task contract. The fixed fee for each assessment completed is A. The referral payment given to the contractor for each recipient she decides qualifies for the service is d. The cost to the agent of misreporting is $C(Z)$, which captures the disutility of misreporting and the effort spent doing that misreporting. $C(Z)$ is a function of the difference between the proportion of recipients that received a new refrigerator and the regulator's expected proportion of eligible households, α. The misreporting cost function is what deters agents from misreporting on every job they conduct, even though the principal is not able to monitor agent actions. Both $C'(Z_I)$ and $C''(Z_I)$ are positive, which reflects the idea that as misreporting rates rise, the chance of the principal catching the agent increases at an increasing rate. A higher chance of being caught by the principal increases the agent disutility of misreporting and the amount of effort the agent exerts to misreport. The principal engages in a simple form of monitoring where it compares the expected population eligibility rate, α, with the proportion of assessments that the agent deems eligible. There is no specific threshold at which the principal will detect misreporting and punish the agent, but the agent perceives the risk of the principal catching her as increasing with her misreporting rate.[5] Following Baker (2002), I use a functional form of $C(Z_I) = \frac{1}{2}\gamma_I Z_I^2$ for the misreporting cost function.[6] The parameter γ_I converts the misreporting rate into dollars for the integrated-task contract.

Plugging the misreporting cost function into equation (B1), I derive the optimal level of agent misreporting Z_I^* using the first-order condition as follows:

$$\frac{\partial \pi_I}{\partial Z_I} = n(1-\alpha)(g+d) - n(1-\alpha)\gamma_I Z_I^* \tag{B2}$$

$$Z_I^* = \frac{g+d}{\gamma_I}. \tag{B3}$$

4. In the ESA program, α reflects the proportion of ESA-qualified households with a refrigerator manufactured in 1992 or earlier.

5. For example, if the principal has an expectation that 13 percent of the population qualifies and observes 16 percent of the assessed population receiving the service, it likely would not suspect agent misreporting. If, however, 55 percent of the assessed population receives the service, then the principal is much more likely to terminate the contract because of suspected misreporting. The agent is aware of this, and therefore has increasing disutility and cost at higher levels of misreporting.

6. Baker (2002) uses the functional form to model the cost of effort for the agent. I adapt it to this setting to model the cost of misreporting.

Equation (B3) shows that the optimal misreporting rate is equal to the benefits of misreporting $(g+d)$, scaled by γ_I. This equation can be used to derive an estimate for γ_I when the values of Z_I^*, g, and d are known.

I next consider the expected value of this contract to the principal while taking into account the optimal level of agent misreporting, Z_I^*. The value to the principal of the integrated-task contract is as follows:

$$E(V_I^{\text{principal}}) = \overbrace{nb_1 n\alpha}^{\substack{\text{Net benefits}\\ \text{of qualified}\\ \text{service provision}}} + \overbrace{nb_2 n(1-\alpha)Z_I^*}^{\substack{\text{Costs}\\ \text{of unqualified}\\ \text{service provision}}} - \overbrace{n(1-\alpha)(g+d)Z_I^* - n\alpha(g+d) - nA.}^{\substack{\text{Payments to agent for assessment,}\\ \text{referral, and service provision}}} \tag{B4}$$

nb_1 and nb_2 are the net benefits of qualified and an unqualified service provision, respectively. Both nb_1 and nb_2 include the marginal cost to the agent of the referral and providing the service.[7] For simplicity, I assume $nb_1 \geq 0 > nb_2$, which means that a qualified service has positive net benefits and an unqualified one has negative net benefits.[8]

To calculate the total surplus of the contract, I add together the value to the principal (equation (B4)) and the profit of the agent (equation (B1)):

$$E(\text{Surplus}_I) = \overbrace{nb_1 n\alpha}^{\substack{(1)\\ \text{Net benefits}\\ \text{of qualified}\\ \text{service provision}}} + \overbrace{nb_2 n(1-\alpha)Z_I^*}^{\substack{(2)\\ \text{Costs}\\ \text{of unqualified}\\ \text{service provision}}} - \overbrace{\frac{1}{2}n(1-\alpha)\gamma_I Z_I^{*2}}^{\substack{(3)\\ \text{Cost to agent}\\ \text{of misreporting}}} - \overbrace{nF.}^{\substack{(4)\\ \text{Agent}\\ \text{information}\\ \text{acquisition cost}}} \tag{B5}$$

Equation (B5) is the main formulation of the surplus from the integrated-task contract divided into four components. Part (1) captures the benefits from qualified service provision. Part (2) is the cost of unqualified service provision. It is always negative, since $nb_2 < 0$. Part (3) is the cost of agent misreporting. The last term is the information acquisition cost paid by the agent at every assessment.

B.2 Separated-Task Contract

Under the separated-task contract, the "assessment agent" is responsible for the assessment, and a different "service agent" conducts the service step. The assessment agent's problem is similar to the integrated-task problem in equation (B1). The assessment agent can still intentionally misreport that a potential recipient household qualifies when it does not, but the assessment agent no longer receives g for providing services. The referral payment d continues to provide an incentive for the assessment agent to misreport.[9] The assessment agent's problem for the separated-task contract is specified as follows:

7. I choose to have nb_1 and nb_2 include the marginal cost of service provision because it yields a simpler expression of contract surplus. The results are unchanged if the marginal cost of providing the service is a separate term.

8. A more general formulation could allow $nb_1 > nb_2 > 0$ where nb_2 includes the opportunity cost of providing additional higher benefit qualified services.

9. I rule out any side deals between the assessment and service agents. I do not find any evidence of side deals between contractors in the ESA program.

$$E(\max_{Z_S} \pi_S^{\text{assessment}}) = \overbrace{n\alpha d + n(1-\alpha)(d)Z_S}^{\substack{\text{Assessment agent} \\ \text{referral payment}}} - \overbrace{n(1-\alpha)C(Z_S)}^{\substack{\text{Cost to agent} \\ \text{of misreporting}}} + \overbrace{n(A-F)}^{\substack{\text{Assessment payment} \\ \text{minus information} \\ \text{acquisition cost}}}. \tag{B6}$$

The S subscript indicates the integrated-task contract. The assessment agent chooses a misreporting rate Z_S. As with the integrated contract, the agent chooses Z_S at the beginning of the contract based on $C(Z_S)$ and the other parameters of the contract. The cost of misreporting function is the same as in the integrated contract, $C(Z_S) = \frac{1}{2}\gamma_S Z_S^2$.

The optimal level of misreporting for the assessment agent in the separated contract is

$$Z_S^* = \frac{d}{\gamma_S}. \tag{B7}$$

The service agent provides the service to all of the potential recipients that are deemed eligible by the assessment agent and receive the same net benefit g. The service agent does not make decisions about which potential recipients receive the service, and they do not report unqualified assessments to the principal, because the piece-rate compensation gives them the incentive to provide as many services as possible. The service agent pays an information acquisition cost for every service she provides, which captures that in the separated-task contract there are costs to having two different agents working with the same recipient. The service agent takes the assessment agent's decisions as given and receives the following expected payout based on Z_S^*:

$$E(\pi_S^{\text{service}}) = \overbrace{n\alpha g + n(1-\alpha)Z_S^* g}^{\substack{\text{Service provision} \\ \text{payment}}} - \overbrace{F[n(1-\alpha)Z_S^* + n\alpha]}^{\substack{\text{Service agent} \\ \text{information} \\ \text{acquisition cost}}}. \tag{B8}$$

The service agent pays the information acquisition cost (F) for each service she provides. The information acquisition cost is the same for both assessment and service agents, but it is paid by the service agent only when the service is provided.[10] Based on the misreporting decisions of the assessment agent and the work of the service agent, the expected value to the principal from the separated-task contract is

$$E(V_S^{\text{principal}}) = \overbrace{nb_1 n\alpha}^{\substack{\text{Net benefits} \\ \text{of qualified} \\ \text{service provision}}} + \overbrace{nb_2 n(1-\alpha)Z_S^*}^{\substack{\text{Costs} \\ \text{of unqualified} \\ \text{service provision}}} - \overbrace{n(1-\alpha)(g+d)Z_S^* - n\alpha(g+d) - nA}^{\substack{\text{Payments to agent for assessment,} \\ \text{referral, and service provision}}} \tag{B9}$$

The net benefits of the services to the principal (nb_1 and nb_2) are the same for the integrated- and separated-task contracts. I add the value to the principal (Equation B9) to the value to the

10. I assume the information acquisition costs are the same for agents in both steps and contract types for algebraic simplicity. Allowing information costs to differ between agents and contract types does not change the results of the model.

assessment and service contractors (Equations B6 and B8) to calculate the total surplus from the contract.

The surplus from the separated-task contract is

$$
\begin{aligned}
E(\text{Surplus}_S) = \; & \overbrace{nb_1 n\alpha}^{\substack{(1)\\ \text{Net benefits}\\ \text{of qualified}\\ \text{service provision}}} + \overbrace{nb_2 n(1-\alpha)Z_S^*}^{\substack{(2)\\ \text{Costs}\\ \text{of unqualified}\\ \text{service provision}}} - \overbrace{\frac{1}{2}n(1-\alpha)\gamma_S Z_S^{*2}}^{\substack{(3)\\ \text{Cost to agent}\\ \text{of misreporting}}} \\
& - \underbrace{nF}_{\substack{(4)\\ \text{Assessment agent}\\ \text{information}\\ \text{acquisition cost}}} - \underbrace{F[n(1-\alpha)Z_S^* + n\alpha]}_{\substack{(5)\\ \text{Service agent}\\ \text{information}\\ \text{acquisition cost}}}.
\end{aligned}
\tag{B10}
$$

The main difference between the separated-task surplus (Surplus_S) and the integrated-task surplus (Surplus_I) is term (5), which captures that in the separated-task contract, the information acquisition cost must be paid a second time during the service provision. In Section II.A I calculate the relative benefit of an integrated-task contract by comparing the surplus from each contract type.

B.3 Discussion of Model Assumptions

There are a number of assumptions in the model that are used to simplify the exposition and remove details that add complication but do not change the result. In this section, I discuss some of these assumptions and their implications.

- No costless information transmission across agents (page 7). I assume that information cannot be passed from the assessment contractor to the installation contractor in the separated-task contract. An alternate formulation is that the information can be transferred at a price that is higher than the cost of reacquiring the information.
- Risk neutrality of agents (page 7). I assume that agents are risk neutral, which means their goal is to maximize the expected value of their payout from the contract. The principal-agent literature examines the role of risk-averse agents in many contexts and how it might affect the optimal contract.[11]
- Static optimization for agents (Appendix page 3). I assume that agents choose the optimal level of misreporting Z^* when they agree to a contract. Agents do not dynamically adjust their misreporting rate throughout the contract. Relaxing this assumption would create a more realistic optimization problem, but it would not add further insight to the model.

11. See Sappington (1991) for a more detailed discussion.

C SCE ESA program

C.1 Supplementary Information on ESA Program

Households are eligible for the ESA program if their income is below 200 percent of the federal poverty line.[12] About one-third of Californians qualify. The ESA program is large, upgrading 400,000 households annually at a cost of $285 million per year. Homeowners and renters are eligible, and the ESA program has a goal of upgrading the entire low-income housing stock by 2020 (Research into Action 2011).[13]

The CPUC sets the guidelines for the program and determines enrollment eligibility and which upgrades the program provides. SCE is tasked with implementing the ESA program in its service territory based on the CPUC's rules.[14] About one-third of the 4.4 million households that SCE services are eligible for the ESA program (SCE 2012). Contractors implementing the ESA program are prohibited from offering any other services not provided by the ESA program.

The SCE ESA program is mostly unchanged as of revising this paper in early 2022. Small adjustments have been made to the income eligibility thresholds to account for changes in the federal poverty line. Most of the upgrades are the same, but with a few minor additions. For example, SCE now provides smart power strips to most households that participate in the program. Households are eligible for free refrigerator replacements if their existing units were manufactured in 1998 or earlier, which is the same cutoff that was set in September 2012.

The analysis in this paper focuses on the ESA program administered by SCE. As an electricity utility, SCE is only permitted to run an ESA program that installs energy efficiency measures intended to reduce electricity consumption. The majority of SCE's service territory overlaps with the natural gas utility Southern California Gas (SoCalGas), which runs its own ESA program focused on reducing natural gas consumption for its customers. SoCalGas has a referral system to pass households that received upgrades through its program to SCE. One-third of the households that are enrolled in the SCE ESA program were referred by SoCalGas. In the analysis, I omit the SoCalGas referrals, which are called joint-utility upgrades, because contractors face different incentives and the service territories of the two utilities do not fully overlap. Referrals were not made from the SCE program to the SoCalGas program during the period I study.

Table C1 shows the upgrades that are provided at no cost to the households between 2009 and 2012. Compact fluorescent light (CFL) bulbs are the most common upgrade, with 290,645 provided to 64,581 households. ESA program rules state that homes "must receive all feasible measures offered under the ESA program" and also require that a household must save a minimum

12. Between 2009 and 2013, a household with four occupants would qualify if their combined income was below $44,100 per year. The ESA program shares the same income thresholds with the California Alternate Rates for Energy (CARE) program, which provides a 30-35 percent discount on electricity bills for low-income consumers.

13. Renters make up 68 percent of the enrolled population. Contractors are required to get a signed consent form from landlords before any work is conducted. I find that this requirement does not change the rate of upgrades in the program.

14. This is a common structure for many energy efficiency programs in the United States. In some cases, the regulator will give utilities additional compensation to offset any revenue loss from decreased demand caused by the energy efficiency program. The SCE is not eligible for any financial compensation from savings from the ESA program.

amount of energy—based on ex ante projections—to qualify for an upgrade (CPUC 2013). CFL installations are considered "minor" upgrades in the ESA program and on their own are not enough to count as a full upgrade. To receive any upgrades, including lighting, a household must be eligible for a "major" upgrade, which consists of refrigerators, evaporative coolers, HVAC upgrades, pool pumps, and room AC units. Of the major upgrades conducted in the ESA program, 55 percent were refrigerator replacements.

Table C1: Count of Upgrades in ESA Program

Upgrade	Count
Refrigerator	41,529
CFL bulb (total installed)	290,645
CFL bulb (households received)	64,581
Evaporative cooler	20,153
HVAC package	1,575
HVAC condenser	3,744
Floor lamp	5,157
Pool pump	3,470
Room AC unit	1,498
Total households enrolled	180,105

Note: This table shows the type and number of upgrades that were provided in the SCE ESA program between January 2009 and August 2012 for the sample studied.

Table C2 shows the most frequent combinations of upgrades that households receive in the ESA program. Refrigerators and CFLs were the most common upgrades installed together, with 27,730 households receiving this combination. Despite the major upgrade requirement, 8,767 households received CFLs but no major upgrade. In most of these cases, the households were found to be eligible for a major upgrade, but one was never installed due to the household refusing or other logistical issues. It is worth noting that with the exception of the refrigerator-only category, CFLs were provided to almost all households that received at least one major upgrade.

Table C2: Combinations of Upgrades in ESA Program

Upgrades conducted together	Household count
CFLs, evaporative cooler, and central air	1,170
CFLs and pool pump	1,347
CFLs and floor lamp	2,934
Refrigerator, CFLs, and evaporative cooler	3,350
Refrigerator only	6,852
CFLs only	8,767
CFLs and evaporative cooler	10,683
Other combinations	12,292
Refrigerator and CFLs	27,730
No upgrade	104,980
Total households enrolled	180,105

Note: This table shows the most frequent combination of upgrades that households received in the SCE ESA program between January 2009 and August 2012 for the sample studied. Other combinations refers to the 67 other combinations of upgrades, many of which include less common ESA upgrades.

C.2 Ride-Alongs to ESA Home Upgrade Visits

I went on a number of ride-alongs with the ESA program in the PG&E service territory as part of my research. I observed both assessments and upgrades, and I also was able to speak to contractors about their role in both steps. The ride-alongs were helpful in crafting the model and writing the paper.

I observed that it was challenging to work directly with households, which is frequently observed in-home energy efficiency retrofit programs (Fowlie, Greenstone, and Wolfram 2018). The ESA program required a number of visits from contractors, each of which needed an adult to be present in the dwelling. The adult had to be home during a pre-specified window, which usually fell during the workweek. The contractors were aware that frequently people were not home during their appointment, which made scheduling and completing upgrades much harder.

Contractors frequently encountered households that were skeptical that the contractor was offering utility services instead of an unwanted solicitation (Research into Action 2011). Contractors were instructed to hold up their badges with the firm name on them in order to demonstrate their legitimacy to the households. It is possible that an integrated-task contract, where the same contractor conducts the assessments and upgrades, could reduce these types of barriers that are common to energy efficiency programs. These benefits are captured in the information acquisition cost.

D Contractor Misreporting Behavior

D.1 Assessment Data

I use confidential assessment data to conduct my analysis in Section III. The raw data contains assessments for 482,439 households in the ESA program between 2009 and 2014. For the main analysis, I drop ESA households that were referred to SCE from SoCalGas. This leaves 271,409 households that enrolled in the ESA program directly. I limit the sample to contractors that conducted at least 250 assessments between 2009 and 2014. This drops 283 assessments conducted by 6 contractors. These contractors conducted a small number of assessments in my sample because they primarily conducted joint-utility enrollments. The final sample consists of 271,126 assessments conducted by 22 contractors.

D.2 ESA Program Advertising

The ESA program did not advertise the eligibility criteria for refrigerator upgrades, and its solicitations to apply focused on income eligibility and what appliances one could get upgraded for free. Unfortunately, I could not access a copy of the brochures SCE used between 2009 and 2012. However, I was able to find a 2020 example of SCE's ESA advertisements. Appendix Figure D2 shows a 2020 email sent to potential enrollees that lists the appliances a household could have upgraded, but it does specify the eligibility criteria.

Figure D2: ESA Program Email from 2020

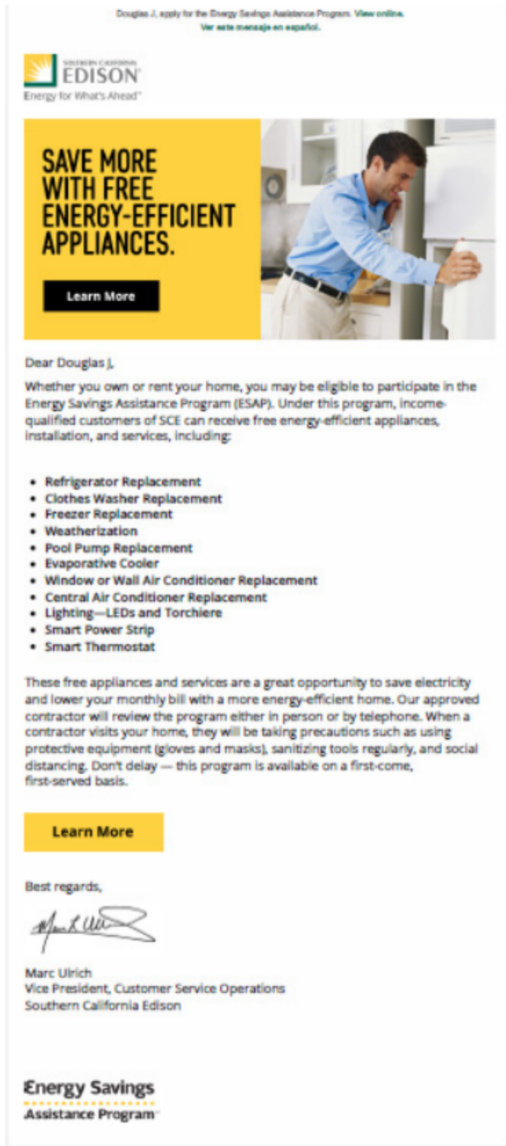

Note: This is a 2020 email for the ESA program. It lists which appliances a household may have upgraded, but it does not discuss the eligibility criteria. Source: https://liob.cpuc.ca.gov/wp-content/uploads/sites/14/2020/12/SCE-AUGUST2020-Low-Income-Monthly-Report.pdf

D.3 Contractor Misreporting Behavior at Eligibility Change

Figure 2 shows the ratio of refrigerators that were reported to be manufactured in 1992 to those in 1993 over time. Immediately after the eligibility cutoff shifted from 1992 to 1998, contractors began reporting an equal number of 1992 and 1993 refrigerators. The horizontal lines show the results from a regression discontinuity, which is estimated using the following equation:

$$ratio_t = \beta_1 \mathbf{1}[\text{1998 eligibility cutoff}= 1]_t + \beta_2 \mathbf{1}[\text{1998 eligibility cutoff}= 1]_t Date_t + \beta_3 Date_t + \epsilon_t. \quad \text{(D11)}$$

The variable $ratio_t$ is the ratio of reported 1992 divided by reported 1993 assessments in month of sample t; $\mathbf{1}[\text{1998 eligibility cutoff}= 1]_t$ is an indicator for the policy change, which turns to 1 in September 2012; $Date_t$ is the number of months the observation is from the September 2012 eligibility cutoff; and ϵ_t is the error term. Errors are clustered at the month-of-sample level.

β_1, the coefficient of interest, represents how the ratio of 1992 to 1993 reported refrigerator ages changed at the September 2012 policy shift. I estimate the regression using data from the 22 months before and after the policy shift. I weight the regression by the number of assessments that reported 1992 or 1993 in each month of the sample. The results are shown in Table D3. The coefficient on the 1998 eligibility cutoff variable shows that the ratio of 1992 to 1993 reported refrigerators dropped by 1.16 when the policy shifted starting in September 2012. The ratio drops from 2.15 to around .99, meaning that a similar number of 1992 and 1993 refrigerators are being reported after the change in the eligibility cutoff. The results are statistically significant and confirm the change in reporting behavior that is observed in Figure 2.

Table D3: Regression Discontinuity Results on Contractor Reporting

	(1) Ratio of reported 1992 to 1993
1998 eligibility cutoff	-1.163 (0.193)
Date	0.002 (0.013)
1998 eligibility cutoff x date	0.001 (0.015)
Observations	42

Note: This table reports the results from estimating equation (D11). It shows that the ratio of 1992 to 1993 reported refrigerator ages discontinuously changed in September 2012 when the refrigerator eligibility cutoff shifted from 1992 to 1998.

D.4 SDG&E Reported Refrigerator Manufacturer Year Data

Figure 3 shows the data from the 106,179 assessments from the SDG&E ESA program between 2009 and 2012. SDG&E structured its ESA program differently, which reduced the incentive for contractors to misreport model numbers. SDG&E hired one firm to manage the program and conduct all the assessments instead of hiring individual contractors to do the assessments.[15] Using this approach removed the incentive for the assessment contractor to misreport manufacture years because they were not the residual claimants of misreporting. The lack of misreporting was a benefit of the SDG&E contract approach, but I am not able to measure if whether this contract structure

15. This approach was feasible partially because SDG&E is a much smaller utility geographically and enrolls only 20,000 households a year, compared with 90,000 in the SCE program.

created additional costs. For example, the SDG&E approach may have removed some of the incentives to work quickly, since contractors were not paid per assessment and upgrade they conducted.

D.5 Predicting Qualified Eligibility Calculations

I use data collected during 2013 to predict what proportion of the ESA population had refrigerators that were manufactured before 1992 during the sample period of 2009 to 2012. I observe that 11.6 percent of refrigerators in the 2013 data were manufactured in 1992 or earlier. I adjust for the natural replacement rate of refrigerators over time to account for the fact that more 1992 and earlier refrigerators would otherwise have been replaced by 2013 than in 2009 to 2012. I approximate the proportion of refrigerators that were manufactured in 1992 or earlier during my sample window with the proportion of refrigerators that were manufactured in 1993 or earlier in 2013. Using this approach, I estimate that 13.8 percent of refrigerators were manufactured before 1993 for my sample window of 2009 to 2012.

Figure D3 shows the differences between the observed distribution and the counterfactual based on 1993 data shown in Figure 4. The vertical bars are the 95 percent confidence intervals from pairwise t-tests between the two distributions at each reported manufacture year. The differences are statistically significant at the 5 percent level for each manufacturer year from 1975 and newer. The significant misreporting for the older refrigerators is due to the large sample size, but the implied misreporting for the older manufacture years is not economically meaningful. Furthermore, there are few assessments reporting early manufacture years, with only 4.9 percent of the assessments listing 1985 or earlier.

The two distributions are statistically different from each other as well. A joint test of significance for all of the coefficients in Figure D3 has an F-value of 157. The estimate of the overall misreporting rate is also precise. The 95 percent confidence interval of the misreporting estimate is 13.0 to 14.0 percent misreporting.

I also consider how within-contractor behavior might affect misreporting. It is possible that errors could be correlated within a contractor firm. The relatively small number of clusters (22) means the standard cluster-robust errors could be too small. To address this concern, I use cluster wild bootstrapped errors in a simple regression framework to conduct the statistical test. I find the misreporting rate remains significant at the 1 percent level, and the 95 percent confidence interval is between 8.1 and 22.0 percent.

I choose the approach in Section III.C because it is simple and transparent. It assumes that the stock of 1993 refrigerators observed in 2013 is similar to the stock of 1992 refrigerators during my sample window. It is possible that refrigerators manufactured in 1993 may age differently than those manufactured in 1992. Without a long-run panel of observed refrigerator manufacture years, it is hard to make this determination. Despite the potential shortcomings, this straightforward approach uses few assumptions and gives a reasonable approximation of the proportion of the sample that was eligible for refrigerator replacements between 2009 and 2012.

Figure D3: Difference Between Observed Distribution and the Counterfactual

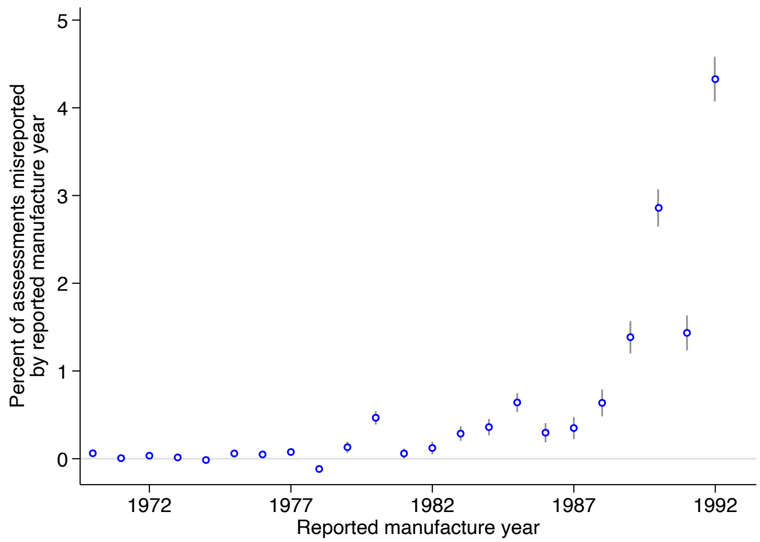

Note: This figure shows the differences between the observed distribution and the counterfactual based on 1993 data shown in Figure 4. The vertical grey lines show the 95 percent confidence intervals from a t-test between the two distributions.

D.6 Contractor Size and Implied Misreporting

The contractors in the ESA program misreport at different rates. Table D4 reports the implied misreporting rates for each of the 22 contractors in the ESA program along with the number of assessments they conducted between 2009 and 2012.[16] The table shows that there is heterogeneity in implied misreporting rates across contractors and contract types. For example, the implied misreporting rate varies from as little as 0.7 percent to as high as 21.8 percent for separated-task contractors. The spread is much larger for integrated-task contractors, with an implied misreporting rate ranging from 8 percent all the way to 53.1 percent.

Contractors in the ESA program are paid the same amount for the services they provide. The large differences in implied misreporting rates may reflect different misreporting costs and benefits across firms.[17] This is analogous to firms having different levels of γ in Section II. I observe all 22 contractors in each year between 2009 and 2012. I do not have data on how long the contractors have worked in the ESA program, which prevents me from studying how misreporting rates vary based on tenure.

16. The table only shows ranges for the number of assessments conducted to avoid revealing the identity of individual firms.

17. The contractor-level implied misreporting rate suffers from measurement error. I use the same predicted eligibility rate of 13.8 percent for all contractors to calculate the implied misreporting rate. On average, the implied misreporting rate is accurate, but it may not be accurate for every contractor.

Table D4: Implied Misreporting Rate by Contractor

Anonymous contractor ID	Contractor type	Assessments conducted	Reported refrigerator eligiblity rate	Implied misreporting rate
1	Separated task	>15,000	14.5%	0.7%
2	Separated task	2,000-15,000	20.8%	7%
3	Separated task	>15,000	22%	8.2%
4	Separated task	>15,000	24.5%	10.7%
5	Separated task	2,000-15,000	25.8%	12%
6	Separated task	2,000-15,000	29.6%	15.8%
7	Separated task	2,000-15,000	32%	18.1%
8	Separated task	2,000-15,000	32.9%	19.1%
9	Separated task	2,000-15,000	33.1%	19.3%
10	Separated task	<2,000	33.5%	19.6%
11	Separated task	<2,000	35.6%	21.8%
12	Integrated task	>15,000	21.8%	8%
13	Integrated task	>15,000	23.9%	10.1%
14	Integrated task	>15,000	25.5%	11.7%
15	Integrated task	<2,000	31.1%	17.2%
16	Integrated task	>15,000	34.2%	20.3%
17	Integrated task	<2,000	37.6%	23.8%
18	Integrated task	2,000-15,000	39.7%	25.8%
19	Integrated task	>15,000	47%	33.2%
20	Integrated task	<2,000	50.5%	36.6%
21	Integrated task	<2,000	54.9%	41.1%
22	Integrated task	2,000-15,000	67%	53.1%

Note: This table shows the implied misreporting rate by ESA contractor. Contractor size is reported in three bins to preserve confidentiality.

D.7 Assignment of Households to Contractors

The assignment of contractors to households is closely related to how the SCE ESA encouraged potential recipients to apply. Despite offering free retrofits, the ESA program spent considerable resources to encourage customers to apply (Research into Action 2011).[18] The need for direct outreach to increase participation in a free energy efficiency program is not unique to the ESA program. Fowlie, Greenstone, and Wolfram (2015) found lower enrollment rates in the federal Weatherization Assistance Program and estimated that there are large non-monetary costs to households to participate in no-cost home energy efficiency retrofit programs.

The ESA program used many different marketing approaches to encourage applicants to apply for the program. SCE used both mailers and automatic call centers to contact potential applicants. A survey of ESA participants reported that 41 percent of households learned about the program from bill inserts and 22 percent through word of mouth (Research into Action 2011). SCE estimates that it cost the company $48 per enrolled lead through mailers (SCE 2010). Other strategies, such as a telethon on Univision, were also reported to have effectively generated leads (Research into Action 2011). The eligibility requirements for the ESA program are the same as those for the CARE program, which has very high enrollment rates, giving SCE a database of potentially eligible

18. While in graduate school in Northern California, I received two calls from my local utility encouraging me to apply for its ESA program.

customers to target.[19] Contractors also knocked directly on the doors of households and left door hangers, but this was reported in only 4 percent of households (Research into Action 2011).[20]

A potential concern is that contractors might try to select homes that either are more likely to have pre-1992 refrigerators or may be better in some ways for misreporting. If this were the case, the analysis could reflect some level of differential selection conducted by integrated- or separated-task contractors. In practice, assessments and upgrades in the ESA program were split up geographically by SCE and assigned to contractors in a somewhat quasi-random manner. ESA program documents state the following:

> Contractors are nominally assigned to geographic areas by ZIP code, but are explicitly not guaranteed sole right to work in those areas. In the relatively few ZIP codes in which more than one contractor is assigned, the system (EMAPS) uses a "round robin" algorithm to assign batches of neighborhoods to alternating contractors authorized for that ZIP (Research into Action 2011).

The result of this round-robin assignment process is that, conditional on the geographic region a contractor was assigned to, contractor assignment to individual houses by SCE is quasi-random. Contractors do have some ability to perform assessments in areas convenient to their staffs (e.g., close to where they live), meaning the assignment is not fully random.

D.8 Refrigerator Replacement Completion Percentage

Not every household that is qualified for a refrigerator replacement will receive one, because of logistical constraints or households not wanting the upgrade. Refrigerator replacements happen on average 62 days after the assessment visit. Sometimes it is difficult for the contractor to schedule a follow-up appointment. Other issues, such as an ungrounded electricity outlet, can prevent an installation from being completed.[21]

The refrigerator replacement completion percentage—defined as the percentage of eligible households (either qualified or intentionally misreported) that received a replacement refrigerator—is 84.2 percent. Appendix Table D5 shows the refrigerator replacement completion percentage broken down by contract type and by whether the assessment contractor correctly recorded the model number. The refrigerator replacement completion percentage for separated- and integrated-task contractors is 83.0 percent and 84.9 percent, respectively. The small difference between the two suggests that the refrigerator replacement completion percentage is not related to the type of contractor that conducts the assessment.

The refrigerator replacement completion percentage can also be used to check for differences in program outcomes across subgroups of households. Landlords are required to provide permission

19. The CARE program requires only self-reporting of income, resulting in many households being enrolled that likely do not qualify. As a result, not all CARE participants are eligible for ESA.

20. Reported percentages come from household survey data on a small sample of participants. Program data did not collect detailed information on how customers learned about the program.

21. During one of my ride-alongs, I observed a household with a refrigerator that was manufactured in 1992 or earlier, but the contractor could not complete the installation because the outlet was not grounded.

Table D5: Refrigerator Replacement Completion Percentage by Sample and Contract Type

Sample	All contractors	Separated-task contractors	Integrated-task contractors
All assessments	84.2%	83.0%	84.9%
Omitted model numbers	83.1%	81.0%	83.9%
Matched model numbers	83.4%	82.4%	84.2%

Note: This table shows the percentage of households that were determined to be eligible for refrigerator replacements during the assessment step that actually received the replacement. The refrigerator replacement completion percentage is similar across all groups. This suggests that contract structure and information about the model number are not correlated with how likely it is for an eligible household to have a refrigerator installed.

for any upgrades conducted for renters in the ESA program.[22] I find that renters and owners have completion rates of 84.4 percent and 83.6 percent, respectively, which suggests that requiring the landlord's permission does not reduce the rate of installations.

D.9 Observable Characteristics by Contractor Type

Table D6 shows the available observable characteristics by contract type. The table shows that separated- and integrated-task firms conduct a similar number of assessments per month and work in similar numbers of zip codes. A t test between the two groups shows that one cannot reject that these values are the same statistically for the two categories.

Table D6: Observable Characteristics of Contractors by Contract Type

	(1) Separated-task contractors	(2) Integrated-task contractors	(3) P-value of difference
Assessments conducted per contractor per month	262	229	0.193
Total zip codes serviced per contractor	265	239	0.711

Note: This table reports the observable characteristics for the 11 separated- and 11 integrated-task contractors.

D.10 Classification of Contractors

Most of the contractors I study do not work exclusively for the ESA program. To better understand the types of contractors and their misreporting behavior, I group them into three contractor classifications.[23] Table D7 shows these classifications by contract type along with the number of assessments conducted and the implied misreporting rates by group. The first classification is "general low-income services" contractors, which provide a wide range of services to low-income communities. For example, a low-income service contractor might also be a food bank, provide worker training, and offer GED classes. Working for the ESA program may be a small portion of the services it provides. This class of contractors misreports on 17.5 percent and 36.2 percent of assessments for the separated and integrated contracts, respectively. Both of these misreporting rates are higher than for the full population of contractors, which suggests that this type

22. Landlords that also pay the utility bill are required to pay $200 toward the refrigerator replacement. This is not common in the ESA program because California has few buildings where landlords pay the electricity bill.

23. I do this by searching the internet based on the firm name and cataloging the non-ESA services they provide.

of contractor might have higher preferences for misreporting.[24] The integrated-task contractors misreport at a higher rate, which highlights that even within a specific classification of firms, the integrated task contract increases the incentive to misreport.

Table D7: Misreporting Rate by Contractor Classification

Contract structure	Contractor classification	Assessments	Number of firms	Reported refrigerator eligibility rate	Implied misreporting rate
Separated	General low-income services	10,985	2	29.1%	17.5%
Integrated	General low-income services	24,265	4	47.8%	36.2%
Separated	Utility subsidy contractors	4,210	2	19.5%	7.9%
Integrated	Utility subsidy contractors	41,107	3	26.4%	14.8%
Separated	General contractors	32,046	4	24.1%	12.5%
Integrated	General contractors	24,553	4	25%	13.4%
Separated	Other/unclassified	42,689	3	16.6%	5%

Note: This table reports the implied misreporting rate by contract type and contractor classification. The classifications are based on the types of non-ESA services each contractor provides and are described in Appendix D.10. The table shows that the implied misreporting rate varies across the contractor classifications but that within each classification, integrated-task contractors have a higher implied misreporting rate.

The second classification of contractors is "utility subsidy contractors," which focus on providing households and businesses with upgrades that are subsidized by the utility or the state. They provide services to homes or businesses such as free or subsidized weatherization through the Low-Income Home Energy Assistance Program (LIHEAP), hot water heater insulation, or air conditioner tune-ups. Some of these firms primarily conduct ESA upgrades for SCE and Southern California Gas (SoCalGas), which provides natural gas service in the same area. The misreporting rate among the utility subsidy contractors is 7.9 percent for the separated- and 14.8 percent for the integrated-task contractors. Both of these are lower than the implied misreporting rate for the overall sample, which shows that these firms misreport less often regardless of the contract structure. Despite the lower levels of misreporting, the integrated contractors appear to misreport at a relatively higher rate within the utility subsidy contractor classification.

The third classification is "general contractors," which provide a wide range of contracting services to homes and businesses. These might include attic insulation, electrician services, furnace repair, commercial energy management systems, and lighting maintenance services. Some of the general contractors advertise utility subsidies for their services, while others are more focused on the wide range of services they provide. It appears that general contractors misreport at a lower rate than the overall population. Furthermore, integrated contractors misreport on just a small amount more than separated contractors, which suggests that general contractors might not be as influenced by the contract structure to misreport.

The last classification is "other/unclassified" contractors, those for which I was not able to find sufficient information to put them into one of the three main classifications. These contractors

24. Many of these organizations are funded by grants, which introduces uncertainty into their funding and ability to hire or retain workers (Research into Action 2011). This need for a more certain revenue stream might explain why the misreporting rates are higher for these firms. It is also possible that these contractors were motivated to some degree by altruism and the desire to provide benefits to enrolling households.

are only on separated contracts and misreport at a lower rate than any of the other types of contractors. Overall, these findings suggest that there is not selection of specific contractor classification into the two contract structures. Furthermore, within each classification of contractor, the rate of misreporting varies in the predictable manner by contract type. Taken together, this suggests that selection of types of contractors into integrated versus separated contracts is not responsible for the differences in Z_S and Z_I calculated in Section III.E.

E Data

E.1 Matching Model Numbers to Manufacture Years

This section describes the process used to verify the year a refrigerator was manufactured using model number data collected by the ESA program. The goal of model number matching is to verify whether a refrigerator was manufactured before or after the 1992 eligibility cutoff. Therefore, the matching process does not need to precisely identify which year a refrigerator was manufactured. During assessments, contractors use serial numbers, not model numbers, to determine the precise year the refrigerator was manufactured, because there is a specific digit within each serial number that identifies the manufacture year. A single model number, by contrast, can be manufactured in multiple years. Fortunately, a given model number rarely spans a standard regime, meaning that I can use the model number to identify whether a refrigerator was manufactured before or after the 1992 eligibility cutoff.

I attempted a number of approaches to determine the year a refrigerator was manufactured. I find that the best approach is to use data from the SDG&E ESA program. The SDG&E and SCE ESA programs recorded the same model number and manufacture year data during the assessment process. The SDG&E program was administered differently, with one firm conducting all the assessments. This same firm was not responsible for any refrigerator replacements, which did not give an incentive to intentionally misreport. Figure 3 shows the reported distribution of refrigerator ages from SDG&E assessments, which shows no bunching to the left of the 1992 eligibility threshold. The absence of bunching suggests that there was no intentional misreporting, which allows me to use it as a reference for the year a refrigerator was manufactured.

One advantage to using the SDG&E data is that the full model numbers were recorded. I attempted to use a number of other databases, including those from the California Energy Commission and the US Environmental Protection Agency and one collected by Midwest Energy Performance Analytics.[25] These databases frequently lacked enough model number digits to accurately match to the SCE data. The SDG&E data, having been recorded using the same process as the SCE data, were a much closer match.

I use SDG&E data from 165,116 assessments collected between 2007 and 2014 as my database against which the reported SCE model number data are matched. There are 64,052 unique model numbers in the SDG&E data. I drop model numbers that appear in the data only once, leav-

25. The last database was graciously made available to me by Jim Cavallo and can be found at http://www.kouba-cavallo.com/refmods.htm.

ing 11,331 unique model numbers. I use the median value of the reported refrigerator age in the SDG&E data as the manufacture year for that model number. Because model numbers rarely span a change in standards, the median value from the SDG&E data is sufficient to identify whether a replacement was qualified or unqualified. I merge these SDG&E model numbers with the 180,105 SCE assessments and find exact matches for 57,847 assessments.

To validate this approach, I manually verify 459 model numbers, those of the more common refrigerators in my data set. These refrigerator model numbers were reported in 43,229 assessments. I find that the SDG&E database performed well compared with the hand-verified model numbers: only one of the 459 reported that a refrigerator was on the other side of the 1992 eligibility cutoff. This high degree of match suggests that using the SDG&E data is an effective approach to identify whether a refrigerator was manufactured before or after 1993.

There are around 104,108 assessments that did not omit the model numbers but also do not contain enough specificity to be matched to a refrigerator manufacture year with a high level of confidence. I calculate the implied misreporting rate for this sample using my information on the misreporting rates for all assessments, the omitted model number sample, and the matched model number sample. I estimate an implied misreporting rate of 12.7 percent for the unmatched sample of 104,108 assessments, which is similar to the implied misreporting rate of 13.5 percent for all assessments. The similarity of the implied misreporting rates suggests that model number transcription errors are responsible for the lack of matching for these 104,108 assessments, rather than strategic contractor misreporting.

Section III.D shows that around 10% of assessments did not include a model number. Without a model number for verification, I am forced to drop these observations. If these households systematically use their refrigerators in a different way than households where the assessment had an accurately recorded model number, than the analysis in Section IV may not reflect the average effect of a refrigerator replacement in the ESA program. Fortunately, households have little discretion over how much energy their basic replacement refrigerator uses (unlike, for example, air conditioning), suggesting that this is likely not a large concern.

E.2 Electricity Consumption Data Cleaning

I use electricity consumption data in Section IV to estimate the effects of refrigerator replacements on household electricity use. I start with the 57,847 assessments for which I am able to verify the year of manufacture based on their model numbers.[26] I merge these households with monthly billing data, leaving me with 3,347,654 monthly observations from 55,340 households. I drop 1,829 households that consumed over 2,000 kWh/month during the sample. This drops the highest electricity-consuming households and data errors that mistakenly report very high levels of consumption. I next drop the 14,542 households that do not have at least 6 months of data before and after the upgrade was conducted. An additional 46 households are dropped due to missing fields that are required for the analysis. I restrict the analysis to electricity consumption within

26. Refer to Appendix E.1 for a description of the model matching process.

two years of a refrigerator replacement, which leaves me with 1,689,239 observations from 38,923 households that I use to conduct the main analysis. The average household has between 20 and 21 months of data before and after it enrolled in the program.

E.3 Weather Data

I use weather data obtained from MesoWest to calculate the monthly heating and cooling degree days. The data come from a balanced panel of 117 weather stations across Southern California. Households are matched to the closest weather station, which is on average 8.6 kilometers from the home.

E.4 Comparing Regression Sample to Full ESA Population

The regression sample in Section IV is a subset of the establishments that participated in the ESA program between 2009 and 2012. The subset is based on the households where I can use the model number to correctly identify the age of a refrigerator to estimate the electricity savings from a qualified and unqualified replacement. The regression sample is described in Appendix E.1 and Appendix E.2.

Appendix Table E8 compares the observable characteristics of households in the Full ESA sample to the regression subsample.[27] It shows that households in the full ESA sample consume more electricity and experience more cooling and heating degree days. The reason for these differences can be seen in the proportion of households in coastal, near-coastal, and inland climate zones.[28] The regression subsample has more households in coastal and near-coastal climates and fewer in inland climate zones. In SCE's service territory, living closer to the coast is associated with milder weather and lower electricity consumption. The regional patterns of households in the regression subsample result in different observable characteristics. There are also statistically significant differences in characteristics such as proportion renter and proportion living in a single-family home. However, the differences are only one percentage point, which is not economically meaningful, suggesting that the large sample size caused the statistically significant differences between the two categories.

The observable differences likely come from the model matching process I used to find the regression subsample. I could only verify model numbers for contractors that reported a full and accurate model number. It is possible that contractors in inland regions were less likely to report accurate model numbers, meaning they are underrepresented in the regression subsample.

However, the differences between the two groups do not bias the empirical results. The goal of Section IV is to determine the energy savings from qualified and unqualified replacements. The energy savings from a refrigerator replacement is not significantly affected by the climate zone or

27. There are 828 fewer households in the full sample in Appendix Table E8 because it was not possible to match these households to their billing data.

28. Climate zones are administratively determined by the California Energy Commission to designate areas with similar climates within California.

Table E8: Summary Statistics for ESA Participating Households Compared to Regression Subsample

Variable	Full ESA sample	Regression subsample	P-value of difference
Pre-period use (kWh/month)	491	441	<.01
	(304)	(233)	
Average monthly CDD	99	91	<.01
	(75)	(53)	
Average monthly HDD	156	147	<.01
	(87)	(63)	
Proportion renter	.66	.65	<.01
	(.47)	(.48)	
Proportion single-family home	.67	.66	<.01
	(.47)	(.47)	
Proportion disabled resident	.07	.08	.23
	(.26)	(.27)	
Proportion English speaker	.55	.49	<.01
	(.50)	(.50)	
Coastal climate zone	.11	.13	<.01
	(.32)	(.34)	
Near-coastal climate zone	.66	.71	<.01
	(.47)	(.45)	
Inland climate zone	.22	.15	<.01
	(.42)	(.36)	
Households	179277	38007	
Observations	6702335	1650101	

Note: This table shows pre-period summary statistics for the sample of households in Section III to the regression subsample of households used to estimate the energy savings from qualified and refrigerator replacements in Section IV. The far right column reports p-values of the difference between the two groups. CDD and HDD signify cooling and heating degree days, respectively.

outdoor temperature.[29] Furthermore, the household-month-of-year and climate zone-by-month-of-sample fixed effects will absorb the differences between the samples. As a result, the energy savings estimates in Section IV are likely reflective of the savings from a qualified or unqualified replacement in the full ESA sample.

F Measuring Electricity Consumption Effects of Refrigerator Replacements

F.1 Refrigerator Replacements Pre-trends

Figure F4 compares the monthly electricity consumption between the eligible and ineligible households in the two years before a household enrolled for the ESA program. The figure reports the residual electricity consumption after account and time fixed effects are removed to better reflect the regression in equation (2). Panel A shows the consumption in the warmer inland regions, which have higher cooling demands in the summer. Panel B shows the pre-upgrade consumption for households

29. A refrigerator may require a small amount more energy to cool as indoor temperatures rise. However, most homes in this region that experience high temperatures also have air conditioning, which will minimize any climate zone differences.

that live in the cooler coastal areas, where summer demand for cooling is lower. The juxtaposition shows how the climate in each region greatly affects how households consume electricity.

In most cases, households with refrigerators manufactured in 1993 or later consume less electricity than households in the same region with refrigerators manufactured in 1992 or earlier, which can be seen in Figure A1. Figure F4 removes the differences in levels to show that the consumption follows similar trends before a household enrolls in the ESA program. It appears that ineligible households (group 3) in the control group consume more electricity in the summer in the inland climate zones. This is probably due to higher air-conditioning usage, which usually spikes during the summer months. The comparison in Figure F4 is not a definitive test of parallel trends, however, because equation (2) includes more fixed effects. Furthermore, the average monthly pre-period is not necessarily balanced across years and regions, which means any differences (or similarities) in Figure F4 may reflect variation that would be absorbed in the regression. Despite these differences, Figure F4 shows that ESA-enrolled households that received qualified refrigerator replacements were similar to ineligible households before they were upgraded.

Figure F5 compares the pre-ESA participation electricity consumption for households that received unqualified refrigerator replacements with those that did not. The figure presents the same type of comparison as Figure F4 for ineligible households that received upgrades due to contractor misreporting. The same caveat applies that this representation cannot directly test the parallel trends assumption, but it can provide some insight on how the two groups of households used electricity before enrolling in the ESA program. Both Panels A and B show similar overall consumption patterns across months, which suggests that contractors are not systematically targeting specific types of electricity consumers when they misreport.

Figure F4: Comparison of Pre-period Electricity Consumption for Qualified Replacements

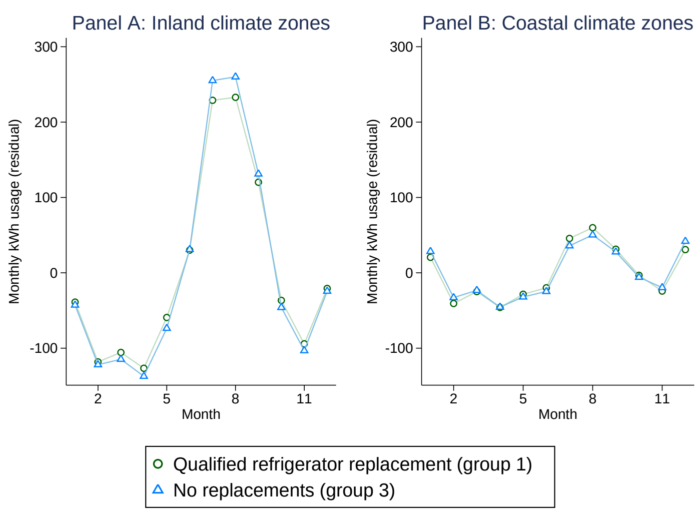

Note: This figure shows the monthly pre-ESA program participation electricity consumption for households that received qualified upgrades (group 1) and control households (group 3). Electricity consumption is shown conditional on household fixed effects.

Figure F5: Comparison of Pre-period Electricity Consumption for Unqualified Replacements

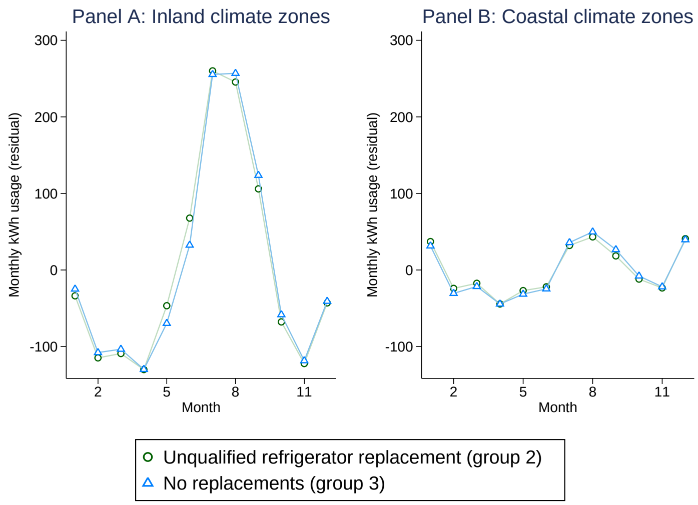

Note: This figure shows the monthly pre-ESA program participation electricity consumption for households that received unqualified upgrades due to contractor misreporting (group 2) and control households (group 3). Electricity consumption is shown conditional on household fixed effects.

F.2 Placebo Event Study

Figure F6 assigns a placebo event study indicator to households in the control group (group 3) that went through the ESA program but did not receive any upgrades. The placebo event study tests whether there were any trends or changes in consumption by the control group due to enrollment in the ESA program. It shows a flat response and no change in consumption at enrollment, which suggests that these households are good controls for both of the event studies.

Figure F6: Event Study: Placebo Treatment on Control Group

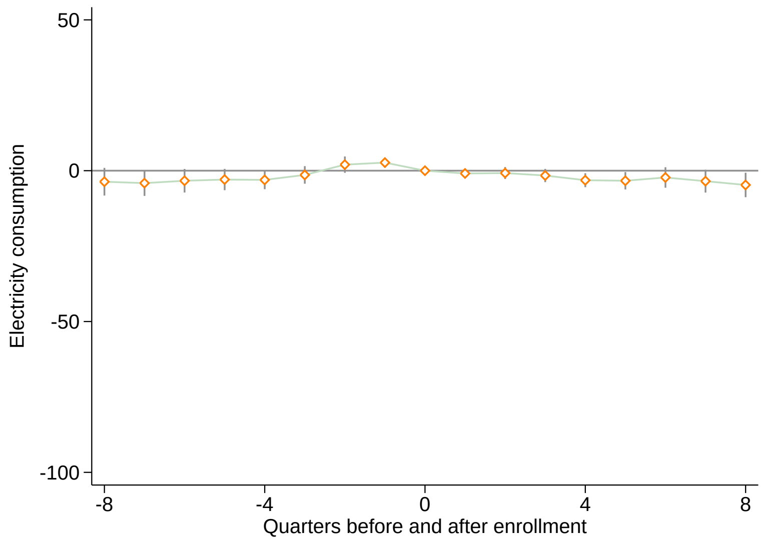

Note: This figure shows an event study for households that participated in the ESA program but did not receive any upgrades. It assigns a placebo event study indicator to households for the date they were enrolled in the program and estimates equation (3) in the same way as Figure 7. The results shows no change in monthly electricity consumption due to enrollment in the ESA program, indicating that the reduction in consumption in Figure 7 was caused by refrigerator replacement.

F.3 Full ESA Sample

In this section, I estimate the average effect of a refrigerator replacement in the ESA program. I use 28,800 refrigerator replacements where I cannot necessarily verify their manufacture year. I estimate equation (2), but I do not differentiate between qualified and unqualified replacements. This approach contrasts with Section IV, where the analysis focused on a subsample of around 5,000 refrigerators where I could confirm the true manufacture year using the reported model number. The results of this analysis are shown in Table F9. I find that the average refrigerator replacement in the estimation sample saves 61 kWh/month. The savings reflect a mix of qualified and unqualified replacements in the ESA program.

The analysis in Table F9 has advantages and disadvantages. The main advantage is that the analysis does not depend on the subsample of refrigerators where I verified the refrigerator manufacture year. The estimates based on the larger sample also better reflect the average savings of a refrigerator replacement in the ESA program. However, two disadvantages prevent me from using it in the main analysis. First, it does not allow me to calculate the net benefits separately for qualified and unqualified replacements. Second, while the sample is larger, I can only use 28,833 of the 41,500 refrigerator replacements conducted by the ESA program between 2009 and 2012. I did not include households in the estimation sample if they did not have sufficient or reliable billing data. Appendix E.2 details the screening process, which is responsible for cutting many of the refrigerator replacements. For example, I dropped households if they did not have more than six months of data before and after the upgrade.

Table F9: The Average Effect of Refrigerator Replacements in the Full ESA Program

	(1) All ESA households
Refrigerator replacement	-60.91***
	(1.50)
Cooling degree days	0.11***
	(0.02)
Heating degree days	0.03***
	(0.01)
Controls for other upgrades	Yes
Pre-period consumption (kWh/month)	474
Refrigerator replacements	28,833
Households	124,984
Observations	5,401,982

Note: - This table reports regression coefficients from one regressions using equation (2) that estimates the average effect of a refrigerator replacement in the ESA program. The dependent variable is monthly household electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. The refrigerator replacement coefficient that in the estimation sample, a refrigerator replacement saved 61 kWh/month. The regression includes household by month-of-sample and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

F.4 Non-Refrigerator ESA Upgrades

The panel fixed effects results in Section IV.E control for all the non-refrigerator upgrades conducted in the ESA program. Table F10 shows the same results as Table 2, including the individual coefficients of the non-refrigerator upgrades. Most of the measures are associated with small reductions or insignificant changes in electricity consumption. The coefficients should not be interpreted as causal impacts of non-refrigerator upgrades.

The savings for non-refrigerator upgrades are similar for the qualified replacements regression in Column 1 and the unqualified replacements regression in Column 2. Enrolling in the ESA program is associated with a small increase in electricity consumption, which could be due to households enrolling in the ESA program after they face positive idiosyncratic shocks to their consumption. Evaporative coolers provide a significant reduction of 31 kWh/month. In dry climates like SCE's service territory, evaporative coolers are substitutes for air-conditioning that can use one-quarter the electricity to cool a home (DOE 2018). Each compact fluorescent light (CFL) bulb replacement is associated with a small reduction in consumption between 1.2 and 1.5 kWh/month.

Many of the upgrades in the ESA program were uncommon, which results in large standard errors when estimating their impacts. Table F11 shows the number of upgrades by group. The estimates in Column 1 of Table F10 include households in group 1 and group 3, while the estimates in Column 2 include include households in group 2 and group 3. Table F11 shows that there are a relatively small number of HVAC packages, floor lamps, pool pumps, and room air conditioner units.

Table F10: The Effect of Refrigerator Replacements on Monthly Electricity Consumption

	(1) Qualified replacements	(2) Unqualified replacements
Refrigerator replacement	-73.45	-38.02
	(2.68)	(3.20)
ESA enrollment	2.65	2.90
	(0.81)	(0.82)
HVAC condenser	-12.35	-10.22
	(5.69)	(5.88)
HVAC package	-14.27	-12.69
	(8.86)	(9.09)
Evaporative cooler	-31.15	-31.24
	(6.76)	(6.79)
Torchiere CFL bulb	-1.52	0.53
	(3.02)	(3.16)
Pool pump	6.47	9.87
	(7.19)	(7.37)
Room AC	11.02	5.94
	(9.42)	(9.02)
CFLs installed (number)	-1.22	-1.53
	(0.44)	(0.48)
Cooling degree days	0.12	0.11
	(0.02)	(0.02)
Heating degree days	0.03	0.03
	(0.01)	(0.01)
Pre-period consumption (kWh/month)	448	445
Refrigerator replacements	3,715	1,261
Households	36,747	34,293
Observations	1,581,024	1,474,841

Note: This table reports regression coefficients from two regressions using equation (2) with the coefficients for the non-refrigerator upgrades displayed. The dependent variable in both regressions is monthly household electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. Table F11 shows the count of upgrades for each regression by group. Both regressions include household by month-of-sample and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

Table F11: Number of Upgrades Conducted in ESA Program by Group

Upgrade	Group 1 Qualified replacements	Group 2 Unqualified replacements	Group 3 No replacements (control)
Refrigerator	3,709	1,260	0
CFL bulb (total installed)	12,952	4,436	29,019
CFL bulb (households received)	3,029	1,003	6,313
Evaporative cooler	290	185	3,258
HVAC condenser	39	6	622
HVAC package	23	4	256
Floor lamp	136	32	1,119
Pool pump	52	13	485
Room AC unit	39	15	206
Households	3,713	1,261	33,023

Note: This table shows the number of upgrades conducted for the panel fixed effects regression in Table 2. The upgrade counts are broken down into the three groups as defined in Section IV.C. The results in Column 1 of Table 2 include the households in group 1 and group 3, while the results in Column 2 include the households in group 2 and group 3.

F.5 Alternate Model Number Matching Approach

In this section, I use a less stringent model matching approach to determine the year each refrigerator was manufactured. I do this by using the same matching technique as described in Appendix E.1, but I drop the requirement that the refrigerator appears in the SDG&E data at least once. This approach expands the list of model numbers I can match with the SCE ESA data, but it reduces the confidence I have in the match. It makes it more likely that I will misclassify a refrigerator replacement as either qualified or unqualified, which could affect the regression results.

Table F12 runs the regressions in Table 2 using the alternate model number matching approach. Column 1 shows that qualified replacements save 70 kWh/month, which is 3 kWh/month less than the primary estimates using the more stringent model number matching approach. The sample includes about 14,000 more households, 3,000 of which received refrigerator replacements. The small difference between the two estimates shows that the results are robust to using an alternate model number matching approach.

Table F12: The Effect of Refrigerator Replacements on Monthly Electricity Consumption: Alternate Model Number Matching

	(1) Qualified replacements	(2) Unqualified replacements
Refrigerator replacement	-69.81	-36.75
	(2.23)	(2.87)
Cooling degree days	0.12	0.11
	(0.02)	(0.02)
Heating degree days	0.04	0.04
	(0.01)	(0.01)
Controls for other upgrades	Yes	Yes
Pre-period consumption (kWh/month)	453	449
Refrigerator replacements	6,801	1,677
Households	50,816	45,692
Observations	2,191,318	1,969,357

Note: This table reports regression coefficients from two regressions using equation (2) using a different model number matching approach than the main results in Table 2. The dependent variable in both regressions is monthly household electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. The refrigerator replacement coefficient in Column 1 shows that qualified replacements cause a 70 kWh/month reduction in consumption. Column 2 shows that unqualified replacements cause a 37 kWh/month reduction in consumption. Both regressions include household by month-of-sample and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

Column 2 estimates the effect of unqualified replacements on electricity consumption using the alternate model number matching approach. It shows a savings of about 37 kWh/month, which is 1 kWh/month less than the primary estimates in Table 2. The sample in Column 2 has around 11,000 more households and 400 more unqualified refrigerator replacements. The similarity between the two estimates shows that the unqualified results are also robust to an alternate model number matching approach.

F.6 Household-Specific Time Trends

In this section, I estimate the effect of qualified and unqualified refrigerator replacements using equation (2) with household-specific time trends instead of household-month-of-year fixed effects. This regression specification controls for trends in household electricity consumption over the sample. The treatment effect is estimated from a change in trends. Household-specific time trends can absorb any linear trends in electricity consumption that might be correlated with a household receiving a refrigerator.

The results from the household-specific time trends specification are shown in Table F13. Qualified replacements save 4 kWh/month more and unqualified replacements save 2 kWh/month less than in the main specification in Table 2. I cannot statistically reject that the estimates are the same. The similarity across these two specifications shows that the results are robust to household-specific time trends and that trends in consumption are not driving the main results.

Table F13: The Effect of Refrigerator Replacements on Monthly Electricity Consumption: Household-Specific Time Trends

	(1) Qualified replacements	(2) Unqualified replacements
Refrigerator replacement	-77.37	-36.12
	(3.44)	(4.38)
Cooling degree days	0.11	0.10
	(0.02)	(0.02)
Heating degree days	0.04	0.04
	(0.01)	(0.01)
Controls for other upgrades	Yes	Yes
Pre-period consumption (kWh/month)	448	445
Refrigerator replacements	3,715	1,261
Households	36,747	34,293
Observations	1,581,024	1,474,841

Note: This table reports regression coefficients from two regressions using equation (2) with household-specific time trends instead of household-month-of-year fixed effects. The dependent variable in both regressions is household monthly electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. The refrigerator replacement coefficient in Column 1 shows that qualified replacements cause a 77 kWh/month reduction in consumption. Column 2 shows that unqualified replacements cause a 36 kWh/month reduction in consumption. These values are similar to the primary saving estimates in Table 2, showing that the results are robust to household-specific time trends. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

F.7 Main Results with only Households that Received Refrigerators and CFLs

In this section, I restrict the main sample to only include households that received either a Refrigerator or a CFL. The motivation for this exercise is that contractors may be strategically installing other upgrades that could introduce bias into the refrigerator energy savings estimation. If, for example, contractors that frequently misreport on refrigerators also do so on evaporative coolers, then the refrigerator savings estimate for those firms in the main specification may be off.

Table F14 shows the results from two regressions using equation (2) where I exclude households that received any upgrade besides a CFL or Refrigerator. The results are similar to those in Table 2, which suggests that other upgrades are not playing a meaningful role in the refrigerator savings estimates.

Table F14: The Effect of Refrigerator Replacements on Monthly Electricity Consumption: Refrigerators and CFLs only

	(1) Qualified replacements	(2) Unqualified replacements
Refrigerator replacement	-72.58	-34.01
	(2.80)	(3.58)
Cooling degree days	0.12	0.11
	(0.02)	(0.02)
Heating degree days	0.05	0.04
	(0.01)	(0.01)
Controls for CFLs	Yes	Yes
Pre-period consumption (kWh/month)	422	417
Refrigerator replacements	3,181	1,017
Households	30,798	28,634
Observations	1,319,248	1,225,710

Note: This table reports regression coefficients from two regressions using equation (2) where I exclude households that received any upgrade besides a CFL or Refrigerator. This is to see if other upgrades are affecting the refrigerator replacement results. I cannot reject that the results are the same as in the primary specification in Table 2, which suggests that other non-refrigerator or CFL upgrades are not playing a large role in the refrigerator savings estimates. Both Column 1 and Column 2 use the same control group of customers that went through the ESA program but did not qualify for refrigerator replacements. Both regressions include weather controls, household-month-of-year fixed effects, and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

F.8 Main Results by Minimum Efficiency Standard of Existing Refrigerators

In this section, I separately estimate the savings from each column of Table 2 based on the minimum efficiency standard regime of the existing refrigerators. The results for qualified refrigerator replacements are shown in Table F15. Column 1 shows that replacing refrigerators manufactured before 1990 with new units saves 85 kWh/month. This estimate is higher than the coefficient for all qualified replacements, because refrigerators manufactured before 1990 did not have to comply with a federal minimum efficiency standard. As a result, this category of refrigerators consumed more electricity than any other group of refrigerators in the program.

Column 2 estimates the effects of qualified refrigerator replacements when the existing unit was manufactured between 1990 and 1992. These replacements save 59 kWh/month. About half of the qualified refrigerator replacements in this sample were manufactured between 1990 and 1992. The qualified replacements savings estimate shown in Table 2 of 73 kWh/month is a weighted average of these two groups.

Table F16 reports the effects of unqualified refrigerator replacements broken down by the standard of the existing refrigerators. Column 1 shows that unqualified refrigerator replacements where the existing units were manufactured between 1993 and 2000 saves 40 kWh/month. In this regression, I limit the control group to households with existing refrigerators manufactured during the same period that did not receive refrigerators.

Beginning in 2001, a new federal minimum efficiency standard went into effect. Column 2 estimates the effects of replacing refrigerators that were manufactured between 2001 and 2012 that satisfied this new standard. The results show a much lower savings of 20 kWh/month, highlighting

Table F15: The Effect of Qualified Refrigerator Replacements on Monthly Electricity Consumption by Standards of Existing Refrigerators

	(1) Pre-1990 replacements	(2) 1990-1992 replacements
Refrigerator replacement	-84.88	-58.54
	(3.64)	(3.47)
Cooling degree days	0.12	0.13
	(0.02)	(0.02)
Heating degree days	0.04	0.04
	(0.02)	(0.02)
Controls for other upgrades	Yes	Yes
Pre-period consumption (kWh/month)	453	450
Refrigerator replacements	1,815	1,900
Households	14,115	14,200
Observations	612,645	616,647

Note: This table reports regression coefficients from two regressions using equation (2). The dependent variable in both regressions is household monthly electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. Column 1 shows the energy savings of 85 kWh/month from replacing a refrigerator manufactured in 1989 or earlier with a new energy-efficient model. These refrigerators consumed more electricity on average because there was no federal minimum efficiency standard when they were manufactured. Column 2 shows the electricity savings of 59 kWh/month caused by replacing a refrigerator manufactured between 1990 and 1992 with a new energy-efficient model. These refrigerators were subject to the first federal minimum efficiency standards and consumed less electricity than the units manufactured before 1990. Both regressions include household by month-of-sample and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

that unqualified replacements of the newest refrigerators resulted in much lower savings than all other types of replacements.

Table F16: The Effect of Unqualified Refrigerator Replacements on Monthly Electricity Consumption by Standards of Existing Refrigerators

	(1) 1993-2000 refrigerators	(2) 2001-2012 refrigerators
Refrigerator replacement	-40.00	-19.86
	(3.88)	(5.75)
Cooling degree days	0.12	0.11
	(0.02)	(0.02)
Heating degree days	0.04	0.03
	(0.02)	(0.01)
Controls for other upgrades	Yes	Yes
Pre-period consumption (kWh/month)	448	443
Refrigerator replacements	968	293
Households	13,268	21,025
Observations	576,374	898,456

Note: This table reports regression coefficients from two regressions using equation (2). The dependent variable in both regressions is household monthly electricity consumption. Refrigerator replacement is an indicator for a household's refrigerator being replaced. Column 1 shows the energy savings of 40 kWh/month from replacing a refrigerator manufactured between 1993 and 2000 with a new energy-efficient model. Column 2 shows the electricity savings of 20 kWh/month caused by replacing a refrigerator manufactured between 2001 and 2012 with a new energy-efficient model. These refrigerators were subject to a more stringent minimum efficiency standard than those in Column 1, resulting in only half the savings from a replacement. Both regressions include household by month-of-sample and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

F.9 Additional Regression Specifications

Table F17: The Effect of Refrigerator Replacements on Monthly Electricity Consumption: Jointly Estimated

	(1) All replacements
Refrigerator replacement	-73.37
	(2.65)
Refrigerator replacement * unqualified	34.58
	(3.48)
Cooling degree days	0.12
	(0.02)
Heating degree days	0.03
	(0.01)
Controls for other upgrades	Yes
Pre-period consumption (kWh/month)	449
Unqualified refrigerator replacements	1,261
Qualified refrigerator replacements	3,715
Households	38,008
Observations	1,635,198

Note: This table reports regression coefficients from a modified version of equation (2) where I pool the households in groups 1, 2, and 3. I estimate the effect of a refrigerator replacement on electricity usage with an interaction term for the refrigerator replacement being unqualified. The results show that a qualified replacement reduces electricity consumption by 73 kWh/month. The savings from unqualified replacements are 35 kWh/month lower. The regression controls for non-refrigerator upgrades conducted in the ESA program and includes weather controls, household by month-of-sample fixed effects, and month-of-sample by climate-zone fixed effects. Standard errors are in parentheses and are clustered at the household and month-of-sample levels.

G Net Benefit Analysis

G.1 Determining Counterfactual Refrigerator Replacement Rate

I use 2013 ESA assessment data to model the replacement rate of old refrigerators. These data were collected in the ESA program after the eligibility threshold was moved to 1998, which removed a contractor's incentives to misreport that refrigerators were manufactured in 1992 or earlier. To approximate the retirement of old refrigerators, I observe the distribution of existing refrigerators and assume that the current stock will decay in a similar manner. I fit a quadratic function through the observed data to produce a smoothed distribution to use for the analysis.

Using real-world data to model the replacement rate in the counterfactual scenario provides a more accurate characterization of the impacts of the ESA refrigerator replacement program. For example, using this approach, I am able to differentiate between how long a refrigerator manufactured in 1994 would remain in the absence of the ESA program compared with one manufactured in 1992. This approach also allows for a proportion of the refrigerators to be retired each year, instead of the standard assumption used in other papers that the program moves up replacements by a fixed number of years (Davis, Fuchs, and Gertler 2014). Furthermore, the assessment data reflects the behavior of the ESA eligible population including the potential purchase of used refrigerators.

If, for example, households in the counterfactual did not dispose of their old units but resold them to other low-income households, that would be reflected in the data and my replacement estimates.

Figures G7 and G8 show the estimated replacement curves for qualified and unqualified replacements, respectively. The solid lines show the fitted curves, which illustrate what proportion of the original stock of refrigerators remains in each year. I adjust the fitted curves to start at one in year zero, which creates a small kink. The qualified refrigerator stock is lower at all points than the unqualified refrigerator stock because qualified refrigerators are older and thus would not last as long in the absence of the ESA program before being replaced.

I rely on only one year of data to model the replacement rate. I assume that the current age distribution reflects how future refrigerators will be replaced. If, for example, newer refrigerators are poorly made and more likely to fail, then this approach will understate the rate at which refrigerators are replaced. I also assume that the replaced refrigerators are scrapped in the counterfactual and not resold or used as second units outside of the sample I am studying. If this were not the case, the ESA program's energy savings relative to the counterfactual would be higher, which would increase the benefits of the ESA program.

Figure G7: Qualified Refrigerators Replacement Rate

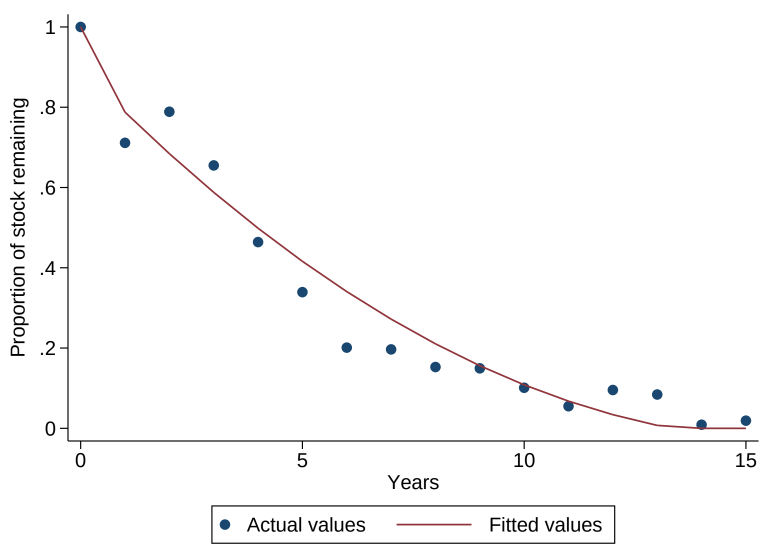

Note: This figure shows how qualified refrigerators would have been replaced in the absence of the ESA program. I fit a quadratic function through the observed points. The fitted function is used to estimate how refrigerators would have been replaced in the absence of the ESA program.

Figure G8: Unqualified Refrigerators Replacement Rate

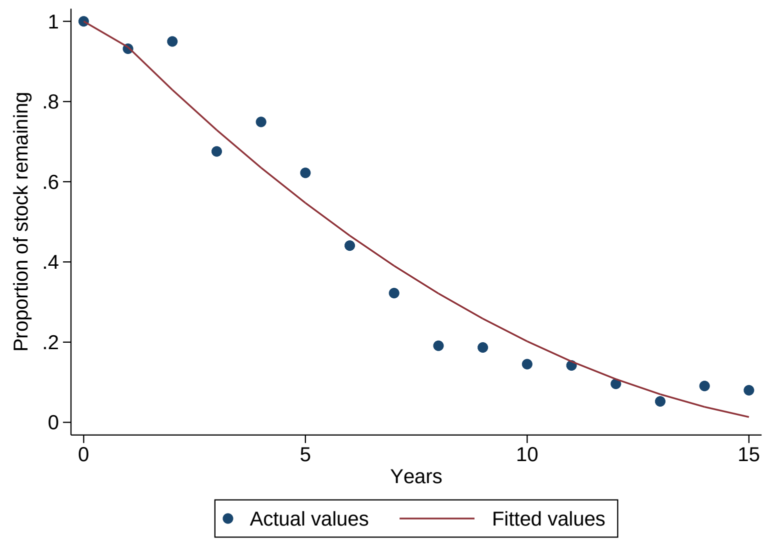

Note: This figure shows how unqualified refrigerators would have been replaced in the absence of the ESA program. I fit a quadratic function through the observed points. The fitted function is used to estimate how refrigerators would have been replaced in the absence of the ESA program.

G.2 Extended Net Benefit Calculation Discussion

This section provides additional details on the assumptions behind the calculations in Table 3 and who directly benefits from each component. The electricity bill savings are calculated using the electricity tariff paid by low-income customers in SCE, which is the California Alternate Rates for Energy (CARE) rate. The income qualifications for CARE are the same as for the ESA program, and all ESA participants are automatically enrolled in CARE if they were not before applying to the ESA program. I assume all customers pay $.127/kWh, which is the second tier of the CARE tariff during this period. The benefits of the electricity bill savings go directly to the occupant of the dwelling that received the refrigerator upgrade.

The second row of Table 3 calculates the reduction in externalities using a $50 social cost of carbon and local air pollution costs in California from Muller, Mendelsohn, and Nordhaus (2011). The externality net benefit is a function of carbon's social cost and the local air pollution costs. If, for example, the social cost of carbon was $100, the total externality benefits from qualified and unqualified replacements would increase to $143 and $99. The externality benefits are society-wide. The social cost of carbon considers the planet-wide cost of emissions, making these benefits global. The local air pollution reductions primarily benefit people living in SCE's service territory because they would lead to small improvements in local air quality.

The third row of Table 3 is the electricity grid fixed costs not recovered. The details behind this calculation are affected by California electricity pricing, where customers do not pay a fixed monthly charge. As a result, all the fixed costs associated with providing electricity must be recovered volumetrically, leading to average prices above the private and social marginal costs of generation (Borenstein 2012). When a household receives a new refrigerator, their reduction in electricity consumption reduces revenue for SCE. Around half of the lost revenue would have paid for the electricity that is no longer sold due to the more efficient replacement refrigerator. There are also externality benefits associated with that reduced electricity generation. The generation and externality costs together are the social marginal cost of generation, which is $.10/kWh in California (Borenstein 2012).[30] The electricity grid fixed costs not recovered is the difference between the $.10/kWh social marginal cost of generation and the $.127/kWh that CARE customers pay during this period. This revenue is needed to pay for the transmission and distribution infrastructure, and those costs do not meaningfully decrease when electricity demand is reduced. The transmission and distribution costs will eventually be spread across all SCE customers in the form of higher electricity rates.

The fourth row of Table 3 is the capital replacement and labor costs. It is the cost of scrapping an old but working refrigerator earlier than would have happened in the absence of the ESA program. I make this calculation using the non-ESA counterfactual refrigerator replacement path described in Section V.A. SCE pays contractors $850 for a refrigerator replacement, which includes the labor and disposal of the existing refrigerator. I value each year that the refrigerator purchase would be moved up in time at its rental rate. I calculate the rental rate by dividing the full $850 cost by 15, which is the assumed lifespan of the replacement refrigerator. I discount future years' payments with a 3 percent discount rate. For example, if the ESA program resulted in a replacement being accelerated by one year, the cost would be $57. However, if the ESA upgrade moves the replacement date 10 years forward, the cost is $498. I assume the household would eventually pay this price for a refrigerator replacement and that the ESA program moved the replacement forward in time. The capital replacement and labor costs are borne by society because they come from removing a working refrigerator before it would have been replaced in the absence of the ESA program.

The fifth row of Table 3 is the distortion associated with raising the revenue to fund this program. In my sample, the ESA program replaced 41,529 refrigerators at a total cost of $35 million. The ESA program raises revenue with a surcharge on all electricity consumed in California. I assume the $35 million is raised over four years using a small rate increase on all the 4.3 million residential customers served by SCE. Using EIA 861 data on SCE, I estimate that this would increase prices by $0.0003 per kWh. To estimate the change in consumption, I use a demand elasticity of .088. This value is taken from Ito (2014), which is estimated on the border of the SCE and SDG&E service territories, and directly applies to the customers I study. I find that this small change in price results in a reduction in consumption of 1.25 kWh per household per year. The distortion associated with this reduction in consumption is $0.03 per refrigerator that is installed by the ESA program. The distortion from raising $8.8 million per year is small because electricity demand is inelastic and the cost is spread across 4.3 million SCE customers.

30. I use Borenstein (2012), which includes capacity construction costs, instead of the more recent estimates in Borenstein and Bushnell (2018) because Borenstein and Bushnell (2018) reflects lower marginal costs due to excess generation capacity in California, which could be transitory.

G.3 Net Benefit Calculation Assumptions

The calculations in Section V.A capture the key components of refrigerator replacement net benefits, but there are some that I am not able to capture. I assume these are small relative to the other costs and benefits I measure. I do not include the costs associated with administering the ESA program in the calculation of refrigerator impacts. Refrigerator replacements, while a significant expenditure in the program, are a relatively small component of running the larger ESA program. There is also no clear estimate of the costs of administering only the refrigerator replacement component of the policy.

For simplicity, I assume that each household would have purchased a similar Energy Star model when its existing unit broke or was replaced. An alternate formulation would have households purchasing refrigerators that met the minimum efficiency standard, which is 9 percent lower than the efficient refrigerators given out by the ESA program until late 2014. In September 2014, the minimum efficiency standard for that same refrigerator increased to be 7 percent higher than the Energy Star model given out by the ESA program between 2009 and 2012.[31] If households purchased refrigerators that just met the minimum standard, then under the replacement patterns used for my calculations, roughly half of the replacements would meet the old standard, and the other half would meet the new standard.[32] These countervailing effects would be close to canceling, leading to a similar energy savings under either set of assumptions. Most refrigerators either meet the minimum efficiency standard or are Energy Star Certified, which suggests the approach I take does a reasonable job at capturing the counterfactual energy usage.

The estimates do not capture the net benefit to the home occupant from being present during the ESA refrigerator delivery. These costs are relatively small, because the ESA program is moving the replacement of a refrigerator forward in time, and a similar time cost would have been experienced in the absence of the ESA program when the existing refrigerator was eventually replaced. Appliance deliveries from major appliance retailers function in a similar way to the ESA program and would have similar costs.[33]

I am not able to capture two benefits in my analysis. The first is the non-energy benefits a household might receive from a new refrigerator. These are likely to be small because the new refrigerators were basic and did not have any special features (e.g., an ice maker or water filtration). Some households may have large improvements in the quality of the refrigeration from a new unit, but, on average, this effect should be small since the ESA program required that the replaced unit was functioning. Any non-energy benefits should be higher for qualified replacements than unqualified re-

31. The 2014 standard increased the efficiency of a typical refrigerator by 27 percent above the 2001 standard. Energy star models sold between 2008 and 2014 were required to exceed the 2001 standard by 20 percent. As a result, the difference between a refrigerator that met the 2014 standard and the 2008-2014 Energy Star refrigerators distributed by the ESA program is relatively small. The new 2014 Energy Star standard required refrigerators to be at least 10 percent more efficient than the 2014 minimum standard. As of late 2020 there are no plans to update the 2014 refrigerator standards.

32. Around half of the replacements in the counterfactual happen within 5 years, which makes late 2014 approximately the halfway point for the counterfactual refrigerator replacements during my 2009 to 2012 study window.

33. Contractors in the ESA program allowed households to schedule replacement appointments outside of standard business hours (Research into Action 2011).

placements because qualified replacements are older refrigerators. As a result, including non-energy benefits would increase the wedge between the effects of qualified and unqualified replacements.

Second, I do not include any refrigerator producer surplus benefits from selling additional refrigerators. I assume the market for refrigerators is large and competitive. Therefore, the additional 40,000 refrigerators sold to the ESA program over four years will not affect the market or generate any measurable producer surplus. It is possible that the ESA contractors have market power in their negotiations with SCE and the PUC over compensation. However, ESA budgets are not binding, which means the market power might increase compensation but not reduce the number of upgrades conducted. As a result, contractor market power will not affect the net benefit calculation because the potential elevated payment to contractors would be a transfer.

The net benefit calculations focus on the energy usage associated with the operation of the refrigerator but do not include other energy use associated with the refrigerator's total life cycle. For example, it takes energy to build a new refrigerator and dispose of an old one. Kim, Keoleian, and Horie (2006) show that electricity consumed operating a refrigerator accounts for 90 percent of the electricity usage in a refrigerator's life cycle (Davis, Fuchs, and Gertler 2014). I am not able to directly account for the energy used to manufacture new refrigerators or dispose of existing refrigerators. However, these life-cycle costs are small because the ESA program does not cause new refrigerators to be manufactured but moves their purchases forward in time. It is possible that the production and disposal technologies could improve over time, but these changes likely are small and would not significantly affect the results.

The net benefit calculation assumes that contractors disposed of the refrigerators they removed from households during replacements. It is a risky for contractors to leave the existing refrigerator when providing a replacement because the ESA program conducts audits after upgrades are complete to verify that the stated work was conducted. The empirical evidence in Section IV shows that not removing existing refrigerators is uncommon, or I would find an increase in consumption after an upgrade. It is still possible, however, that the old refrigerators are not being removed for a small proportion of upgrades. There is no anecdotal evidence in any of the available documents or in my conversations with the regulators that this is happening. Unfortunately, there is no way to conclusively identify if old refrigerators are not removed directly with the available data. Part of the challenge is month-to-month variation in electricity usage can be high and is based on the weather and characteristics of the dwelling. For example, 54 percent of the 26,601 households that went through the ESA program but did not receive an upgrade had an increase in electricity consumption in the 6 months following enrollment compared to the 6 previous months.

I take two approaches to investigate if contractors removed old refrigerators. First, I examine households that received ineligible refrigerator replacements, but no other upgrades. I find that 21 percent of them increased their consumption by 35 kWh/month in the 6 months after a replacement compared to before. I use 35 kWh/month because that is approximately what the refrigerators provided in the ESA program consume, and that increase in consumption would be consistent with the old refrigerator not being taken away. However, 33 percent of households that received no upgrades had an increase of 35 kWh/month. This suggests that the households receiving refrigerators that had an increase in consumption are likely caused by month-to-month variation in consumption, and not a failure to remove the old refrigerator.

Second, I visually examined the consumption of the 75 households that had their usage increase after a refrigerator replacement.[34] Most of the increases in consumption were not concurrent with the upgrade, and appeared to be correlated with changes in weather. These two approaches are not conclusive evidence that contractors are removing old refrigerators when an upgrade is provided, but suggest that it is not a frequent occurrence.

If a non-trivial proportion of replacements did not include disposal, then the net benefit calculations would be different. There would no longer be a cost of an early replacement of a working refrigerator, because the old refrigerator would remain in use. The effect on overall electricity consumption would depend on the purpose of the used refrigerator. If the used refrigerator is added as a second unit (e.g., in someone's garage), then the replacement could increase aggregate consumption leading to a reduction in net benefits.[35] If the used refrigerator replaced a less efficient unit, it could cause a small reduction in electricity consumption and an increase in net benefits. Ultimately, it is hard to know the net benefits of a refrigerator replacement if the used unit is not destroyed.

G.4 Net Benefit Calculations Using Alternate Discount Rates

The net benefit calculations in Table 3 use a 3 percent annual real discount rate to calculate both the costs and benefits of a refrigerator replacement. In this section, I consider how this assumption affects the net benefit calculations. Table G18 shows the net benefit calculations from the base case using three discount rates. Column 1 uses a 1 percent discount rate, finding that unqualified replacements reduce net benefits by $110. Column 3 uses a 10 percent discount rate, finding that unqualified replacements reduce net benefits by $95. The small change in net benefits shows that the calculation is not sensitive to the discount rate. Similarly, the net benefits from qualified replacements change from $69 with a 1 percent discount rate to $35 with a 10 percent discount rate.

Table G18: Net benefits of Qualified and Unqualified Refrigerator Replacements: Discount Rate Robustness

	(1) 1 percent discount rate	(2) 5 percent discount rate	(3) 10 percent discount rate
Lifetime kWh savings from qualified replacement	3,677	3,677	3,677
Lifetime kWh savings from unqualified replacement	2,598	2,598	2,598
Net benefit of qualified replacement (nb_1)	$69	$51	$35
Net benefit unqualified replacement (nb_2)	-$110	-$103	-$95
Pre-1993 refrigerators in year 5 without ESA program	42%	42%	42%
Post-1993 refrigerators in year 5 without ESA program	55%	55%	55%

Note: This table shows the net benefits of qualified and unqualified replacements for three scenarios using the electricity savings estimates from Table 2. Each scenario uses a different discount rate, showing similar results for 1, 5, and 10 percent discount rates.

34. I limited my analysis to households that had an unqualified replacement and only received a refrigerator.

35. The access to refrigeration could also increase net benefits, depending on the value the additional refrigeration provides.

The estimates are not sensitive to the discount rate because refrigerator replacements generate an annual stream of costs and benefits. In this calculation, the costs and benefits are similar. As a result, changing the discount rate does not significantly affect the net benefit estimates.

G.5 Alternate Marginal Cost of Service Provision

There are costs to the contractor to install a refrigerator. The costs include fuel costs to drive to the household receiving a replacement and the labor costs of removing the old refrigerator and installing a new one. The marginal cost reduces the contractor's incentive to misreport. In the base case in Table G20, I assume the marginal cost of providing a refrigerator replacement is $30.

Table G19 shows the same calculations as Table G20 with different marginal costs. Column 1 assumes there is no marginal cost to replace a refrigerator, which results in a cost of integration of $13.71 per assessment, which is $1.39 higher than the base case estimate. The value is higher because the larger benefit to the contractor of misreporting results in a higher estimate of the integrated misreporting cost parameter (γ_I), which increases the relative value of a separated-task contract. The opposite effect can be seen in Column 2, which uses a marginal cost of $60 per refrigerator replacement and estimates a cost of integration of $10.84 per assessment. The lower value of a refrigerator replacement to the contractor results in a lower value of γ_I than in the base case, which makes the integrated-task contract not as costly as it is in the primary specification. Both scenarios show that the main results are robust to a range of estimates of the marginal cost of replacing a refrigerator.

Table G19: Benefits of an Integrated-Task Contract Estimates: Alternate Marginal Costs of Refrigerator Replacement

	(1) No marginal cost of refrigerator replacement	(2) Double ($60) marginal cost of refrigerator replacement
Benefits of integration per assessment	-$13.71	-$10.84
Information acquisition cost (F)	$35	$35
Net benefit of unqualified replacement (nb_2)	-$106	-$106
Integrated misreporting cost parameter (γ_I)	735	551
Separated misreporting cost parameter (γ_S)	319	319

Note: This table shows the benefits of an integrated-task contract relative to a separated-task contract using alternate marginal costs of refrigerator replacement. Column 1 has no marginal cost to replace a refrigerator, which increases the value of a separated-task contract compared with the base case in Table G20. Column 2 has a marginal cost of refrigerator replacement of $60, which lowers the relative value of a separated-task contract. Both scenarios show that the separated-task contract is preferred under a range of marginal costs of refrigerator replacement scenarios.

G.6 Own-Firm Upgrade Rate

In the ESA program, an integrated-task contractor who conducts the assessment is not guaranteed to provide the refrigerator replacement. SCE has targets for how long a contractor has to install a replacement, which can result a different firm providing the final replacement (Research into Action 2011).[36] Of the 25,000 refrigerators that were replaced by integrated-task contractors, 69 percent had the same firm conduct the assessment and provide the final upgrade. The own-firm

36. The maximum time to complete a refrigerator replacement is not public to the best of my knowledge.

upgrade rate is lower than 100 percent because of scheduling issues or other logistical constraints that may have prevented the original contractor from providing the upgrade.

Contractors do not know at the time of assessment whether they will provide the refrigerator replacement, which gives them an incentive to misreport on all jobs. It takes on average 62 days from household enrollment to refrigerator replacement. There is significant variation across firms, and there is a negative correlation between the number of days it takes a firm to provide a replacement and the own-firm upgrade rate. This suggests that firms are more likely to provide replacements to households they assessed if they can do so quickly. Overall, having a same-firm upgrade rate below 100 percent reduces the monetary incentive to misreport during the assessment step, but it does not eliminate it.

G.7 Benefits of Integration Calculation Details

There are a number of parameters used in the the calculation of the benefits of an integrated task contract that are not directly observable. I calculate the values for γ_I and γ_S, which convert agent misreporting rates into a misreporting cost, using first-order conditions (equations (B3) and (B7)) and estimates of d, g, Z_I, and Z_S.[37] I leverage the functional form assumption that the cost of misreporting is $C(Z) = \frac{1}{2}\gamma Z^2$. Using equations (B3) and (B7), I calculate a value of 646 and 319 for γ_I and γ_S, respectively.

I also do not observe the information acquisition cost F, and I use a range of information acquisition cost estimates. I use the \$70 SCE pays contractors to conduct household assessment visits as the basis for my estimates. The assessment visit includes a variety of costs to the contractor including scheduling an appointment, transportation to a household, looking at a household's W2 income statement, walking through every room in the dwelling, and recording data. Because the full \$70 payment includes agent activities that do not involve acquiring information related to the refrigerator replacement, I approximate the information acquisition cost for refrigerator replacements in the primary specification as \$35.

Table G20 shows the calculation of the benefits of integration per assessment using the base case refrigerator replacement scenario. The first column shows the results using an information acquisition cost of \$35 per refrigerator replacement, which results in a benefit of integration of -\$12.32 per assessment. In this case, it is beneficial to separate the assessment and refrigerator replacement between two different contractors. In total, integrated contract assessments reduce net benefits of the ESA program by \$1,109,000 across 90,084 assessments. The assessments resulted in 25,000 refrigerator replacements (about 58 percent unqualified), which created \$6,774,000 in benefits from reducing electricity consumption. The distortion introduced by using an integrated-task contract rather than a separated-task contract is equal to 16 percent of the total benefits of the refrigerator replacement program.

37. d and g are \$25 and \$98 and are calculated in Section V.C.

Table G20: Benefits of an Integrated-Task Contract Estimates

	(1) Base case	(2) High information acquisition cost	(3) Zero agent misreporting cost
Benefits of integration per assessment	-$12.32	-$5.11	-$3.07
Information acquisition cost (F)	$35	$70	$35
Net benefits of unqualified replacement (nb_2)	-$106	-$106	-$106
Integrated misreporting cost parameter (γ_I)	646	646	0
Separated misreporting cost parameter (γ_S)	319	319	0

Note: This table shows the benefits of an integrated-task contract relative to a separated-task contract, calculated using equation (1). Each column uses a different set of assumptions. Column 1 shows that the costs of integration are high when using the base case parameters from Table 4. Columns 2 and 3 make progressively stronger assumptions to benefit the integrated-task contract, but the separated-task contract is still preferred.

In Column 2, I present a scenario where the information acquisition cost is $70. Under this conservative assumption, the benefits of integration are -$5.11 per assessment. Even assuming this high information acquisition cost, the benefits of integration are negative.

To tilt the balance further in favor of the integrated-task contract, I zero out the cost associated with contractor misreporting while assuming contractors still misreport at the same rate. The results of the zero agent misreporting cost scenario are shown in Column 3 of Table G20. This scenario does not reflect a plausible scenario, but it is a useful lower bound because it enables me to estimate the importance of the misreporting cost for the benefits of integration. I find that even with no agent misreporting cost, the benefits of integration are -$3.07 per assessment, and thus a separated-task contract is still preferable.

I next calculate the information acquisition cost that would justify an integrated-task contract in the absence of a contractor misreporting cost. I find that the value of integration would have to be at least $50 per refrigerator replacement. This estimate serves as a lower bound on how large the information acquisition cost would have to be to justify task integration. It does not require a functional form assumption on the misreporting cost function. Even in this scenario, designed to favor task integration, I find that a high information acquisition cost is necessary to justify the integrated-task contract.

G.8 Information Acquisition Cost and the Benefits of Integration

According to the contractors in the PG&E ESA program, it was frequently difficult to receive permission to enter the home and conduct the assessment or provide services. Another major concern was that someone would be present at the home during the scheduled window so the contractor could gain access. Contractors presented their badges, which gave their firm names and indicated that they were licensed ESA contractors working with the utility. Having the same contracting firm conduct both the assessment and refrigerator replacement may facilitate this process.

Installation contractors in the ESA program reported that they did not always get enough information from the assessment. Some said that "installation crews typically determined what a home needed then returned to their warehouse for supplies before completing the installation

(Research into Action 2011)." Integrated-task contractors might be more motivated to collect the necessary information on the assessment visit, which could reduce the overall installation costs.

On one ride-along I participated in, the house was designated to receive a refrigerator replacement, but the electrical outlet by the refrigerator was not grounded. The ESA program rules state that a grounded outlet is required, which meant that the contractor could not provide the refrigerator replacement. An integrated-task contractor would have had a greater incentive to check all household qualifications during the assessment visit, which could have avoided the oversight.

G.9 Generalizability

It is important to consider the generalizability to other settings of the results and policy fixes in this paper. List (2020) lays out four conditions – selection, attrition, naturalness, and scaling – to discuss the external validity of the results and policy implications. First, in terms of selection, this paper studies an existing real-world program typical of low-income energy efficiency home retrofits. Large federal programs such as the Weatherization Assistance Program have a similar structure with third-party contractors providing free energy efficiency retrofits to low-income households. Many other energy efficiency programs pay third-party contractors with piece-rate wages to decide which subsidized energy efficiency upgrades they are paid to provide to customers. Outside of energy efficiency, physicians, contractors, and other specialists paid via piece-rate wages frequently face decisions where they could conduct unqualified services to increase their compensation. The insights about contract structure and agent incentives in the ESA program can be scaled to help think about other similar principal-agent problems.

Second is attrition out of the refrigerator replacement program. I find that 84.2 percent of households that were said to be eligible for a replacement ultimately received one. This attrition level is similar between renters, owners, households with separated-task contractors, and households with integrated task contractors. Using these numbers, one could compare the level of attrition in the ESA program to a target program of interest. Third is naturalness, which also depends on the policy or situation compared to the ESA program. There are many other settings, such as ordering extra unnecessary non-invasive medical testing or an auto mechanic doing unnecessary work comparable to contractors' decisions in the ESA program. In different settings that have higher financial or safety stakes, the external validity is likely lower.

The last condition is scaling. This paper's insights are straightforward to scale horizontally to other low-income refrigerator replacement programs implemented outside California. Vertically scaling the findings on the ESA program to other energy efficiency programs or similar principal-agent problems depends on the context. Overall, this paper's insights about contract structure and misreporting are helpful in different settings with comparable incentives. The role that contract structure can play in exacerbating the principal-agent problem applies to other programs and policies with piece-rate wages where monitoring is ineffective to avoid similar pitfalls experienced in the ESA program. When comparing my findings to different situations, the context and details may differ, but insights learned about the principal-agent problem can be applied to other piece-rate wage settings.

References

Baker, George. 2002. "Distortion and Risk in Optimal Incentive Contracts." *The Journal of Human Resources* 37 (4): 728–751.

Borenstein, Severin. 2012. "The Private and Public Economics of Renewable Electricity Generation." *Journal of Economic Perspectives* 26 (1): 67–92.

Borenstein, Severin, and James B Bushnell. 2018. "Do Two Electricity Pricing Wrongs Make a Right? Cost Recovery, Externalities, and Efficiency." NBER Working Paper No. 24756, NBER, Cambridge, MA.

CPUC, California Public Utilities Commission. 2013. *Statewide Energy Savings Assistance Program Policy and Procedures Manual.* San Francisco: CPUC.

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): 207–38.

DOE. 2018. *Evaporative Coolers, Department of Energy.* https://www.energy.gov/ energysaver/home-cooling-systems/evaporative-coolers.

EIA. 2009. *Residential Energy Consumption Survey.* Table CE4.1. Washington, DC: EIA.

Energy Information Administration (EIA). 1987. *Residential Energy Consumption Survey.* DOE/EIA-0321/1(87). Washington, DC: EIA.

EPA. 2011. *ENERGY STAR Unit Shipment and Market Penetration Report Calendar Year 2011 Summary.* Washington, DC: EPA.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2015. "Are the Non-monetary Costs of Energy Efficiency Investments Large? Understanding Low Take-Up of a Free Energy Efficiency Program." *American Economic Review* 105 (5): 201–204.

———. 2018. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program." *The Quarterly Journal of Economics* 133 (3): 1597–1644.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* 104 (2): 537–63.

Kim, Hyung Chul, Gregory A. Keoleian, and Yuhta A. Horie. 2006. "Optimal Household Refrigerator Replacement Policy for Life Cycle Energy, Greenhouse Gas Emissions, and Cost." *Energy Policy* 34 (15): 2310–2323.

List, John A. 2020. *Non Est Disputandum de Generalizability? A Glimpse into The External Validity Trial* NBER Working Paper No 27535.

Muller, Nicholas Z., Robert Mendelsohn, and William Nordhaus. 2011. "Environmental Accounting for Pollution in the United States Economy." *The American Economic Review* 101 (5): 1649–1675.

Research into Action. 2011. *Low Income Energy Efficiency Program 2009-2010 Process Evaluation.* Portland, OR: Research into Action.

Sappington, David E. M. 1991. "Incentives in Principal-Agent Relationships." *The Journal of Economic Perspectives* 5 (2): 45–66.

SCE, Southern California Edison. 2010. *Southern California Edison Company's (U 338-E) Annual Progress Reports for Low Income Programs for the Period January-December 2009.* Rosemead, CA: SCE.

———. 2012. *The Joint Utilities' 2011 Amended Report on CARE Eligibility and Penetration Rates.* Rosemead, CA: SCE.