Title
ICE Calculator 2.0: Final Report for Phase 1 of the National Initiative to Update the Interruption Cost Estimate (ICE) Calculator

Authors
Larsen, Peter H
Carney, Kyle
Eto, Joseph H
et al.

Publication Date
2025-06-15

ICE Calculator 2.0

Final Report for Phase 1 of National Initiative to Update the Interruption Cost Estimate (ICE) Calculator

Principal Investigator: Peter H. Larsen [1]

Contributing Authors: Kyle Carney [2], Joseph H. Eto [1], George Jiang [2], Dhawal Joshi [2], Kristina H. LaCommare [1], Ridge Peterson [2], Chris Ramee [2], Anna-Elise Smith [2]

[1] LAWRENCE BERKELEY NATIONAL LABORATORY
[2] RESOURCE INNOVATIONS, INC.

MAY 2025

Disclaimer

Copyright Notice

ICE Calculator 2.0:
Final Report for Phase 1 of the National Initiative to Update the Interruption Cost Estimate (ICE) Calculator

Prepared for American Electric Power, Commonwealth Edison, Dominion Energy, Duke Energy, DTE Electric, Exelon, National Grid, and Puget Sound Energy

Principal Authors

Principal Investigator: Peter H. Larsen[1]

Contributing Authors: Kyle Carney[2], Joseph H. Eto[1], George Jiang[2], Dhawal Joshi[2], Kristina H. LaCommare[1], Ridge Peterson[2], Chris Ramee[2], Anna-Elise Smith[2]

[1]Lawrence Berkeley National Laboratory
[2]Resource Innovations, Inc.

Lawrence Berkeley National Laboratory
1 Cyclotron Road, MS 90R4000
Berkeley CA 94720-8136

May 2025

The work described in this study was funded by American Electric Power (Award No. AWD00005921), Commonwealth Edison (Award No. AWD00005868), Dominion Energy (Award No. AWD00005819), Duke Energy (Award No. AWD00005788), DTE Electric (Award No. AWD00005795), Exelon (Award No. AWD00006036), National Grid (Award No. AWD00006334), and Puget Sound Energy (Award No. AWD00006534) under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Acknowledgements

The authors would like to recognize and thank numerous organizations and individuals that have contributed to the completion of Phase 1 of the ICE 2.0 Initiative.

First and foremost, we thank the sponsoring utilities for their financial support of the Initiative, including representatives from the sponsoring utilities who served on the ICE 2.0 Project Executive Committee (PEC) or served as our day-to-day liaisons. PEC members reviewed the survey instruments and this final report, and they lead in the coordination of in-kind support for survey activities from across their organizations, especially support from account representatives in enlisting large customers to participate in the survey.

Jaime Ortega and Susanna Aquilar, Commonwealth Edison
Bridget Powers Beggs and Thomas Vaccaro, National Grid
Dave Isaacson, Jennifer Ellis and Daniel Coleman, American Electric Power
Robert Stewart and Harsha Jasti, Exelon
Bill Chiu and Ed Karpiel, DTE Electric
Mark Keener, Duke Energy
Jeff Kensok and Jeff Hahn, Puget Sound Energy
Gus Johnson and Rick Siepka, Dominion Energy

We also thank the members of the ICE 2.0 Project Advisory Committee (PAC) who supported the initiative by reviewing both the survey instruments and this final report and by providing input into the functional enhancements we have made to the updated online ICE Calculator.

Venkat Banunarayanan, National Rural Electrification Cooperative Association
Campbell Delahoyde, National Association of State Energy Offices
Jimmy Glotfelty, National Association of Regulatory Utility Commissions & Public Utility Commission of Texas
Fred Hoover, National Association of State Utility Consumer Advocates & Maryland Office of People's Counsel
Ji Yoon Lee, American Public Power Association
Yohannes K. G. Mariam, National Association of State Utility Consumer Advocates & Office of the People's Counsel – District of Columbia
William McCurry, National Association of Regulatory Utility Commissions
Phillip Moeller, Edison Electric Institute
Sara Mullen-Trento, Electric Power Research Institute
Jacob Nicodemus, National Association of Regulatory Utility Commissions & Public Utilities Commission of Ohio
Kirsten Verclas, National Association of State Energy Offices

We also thank the following independent experts who provided review comments on the draft of this report:

Granger Morgan, Carnegie Mellon University
Bernard Neenan, Energy Resource Economics LLC
Sunhee Baik, Lawrence Berkeley National Laboratory

Additionally, we're grateful to the following technical experts who guided project staff on selected aspects of the Initiative:

Michael Sullivan, Resource Innovations, provided guidance on the development of the survey instruments and their administration.
Michael Hanemann, U.C. Berkeley and Arizona State University, provided guidance on the development and analysis of results from the contingent valuation methods used to survey residential customers.
Josh Schellenberg, H&S Insights, provided valuable feedback on the development of the customer damage functions.

We thank Berkeley Lab administrative staff who worked behind the scenes to ensure the success of the Initiative:

Joanna Santoro, Strategic Partnerships Office
Michelle Wong, Office of Laboratory Counsel
Virginia de la Puente, Intellectual Property Office
Erica Ibrahim, Procurement & Property Department
Kelsey Miller, Human and Animal Regulatory Committees (HARC) Office
Kristan Johnson, Electricity Markets and Policy Department

Finally, we thank the Department of Energy's Office of Electricity for their support for the creation of the ICE Calculator and for their commitment to maintain the ICE Calculator in the public domain:

Gil Bindewald, Principal Deputy Assistant Secretary
Bill Parks, Deputy Assistant Secretary
Sandra Jenkins, Program Manager
Joe Paladino, Senior Program Manager (retired)

The views and opinions of authors expressed throughout this study do not necessarily state or reflect those of our employers, the U.S. Department of Energy, the sponsoring utilities, or organizations represented by the Project Executive or Advisory Committees.

Table of Contents

List of Figures

List of Tables

Acronyms and Abbreviations

AIC	Akaike information criterion
APPA	American Public Power Association
BIC	Bayesian Information Criterion
CDF	Customer damage function
CMI	Customer minutes interrupted
DOE	Department of Energy
EEI	Edison Electric Institute
EPRI	Electric Power Research Institute
GLM	Generalized linear model
ICE	Interruption Cost Estimate
LASSO	Least absolute shrinkage and selection operator
LNR	Large non-residential
MAPE	Mean absolute percentage error
NAICS	North American Industry Classification System
NARUC	National Association of Regulatory Utility Commissioners
NASEO	National Association of State Energy Offices
NASUCA	National Association of State Utility Consumer Advocates
NRECA	National Rural Electric Cooperative Association
OHDC	One-half-bound dichotomous choice
QMLE	Quasi-maximum likelihood estimator
RMSE	Root mean squared error
SMNR	Small non-residential
WFH	Working from home
WTP	Willingness-to-pay
ZCTA	Zip Code Tabulation Area

Executive Summary

About the Interruption Cost Estimate (ICE) Calculator

The ICE Calculator is a publicly available online tool that estimates the economic costs electricity customers experience due to power interruptions.[1] It was first developed over 15 years ago for the U.S. Department of Energy by Berkeley Lab and Freeman, Sullivan & Co. The tool is used routinely by utility planners and decision makers to estimate the economic benefits of grid reliability and resilience improvements.

National Initiative to Update the ICE Calculator

In 2021, Berkeley Lab and Resource Innovations, Inc. launched the “ICE 2.0 Initiative” – a national study to refresh the underlying data and enhance the functionality of the ICE Calculator. The Initiative involves Berkeley Lab contracting with sponsoring utilities to administer identical, updated and comprehensive interruption cost surveys to statistically representative samples of each utility’s customers. Berkeley Lab and Resource Innovations then pool the survey results across the utilities and use them to update the analytical engines that drive the ICE Calculator.

The ICE 2.0 Initiative is being conducted in phases. Each phase involves the administration of interruption cost surveys to the customers of sponsoring utilities, followed by an update to the ICE Calculator based on analysis of the pooled survey results.

This report describes the activities and findings from Phase 1 of the ICE 2.0 Initiative. Phase 1 was sponsored by eight utilities, as shown in ES Figure 1: American Electric Power, Commonwealth Edison, Dominion Energy, Duke Energy, DTE Electric, Exelon, National Grid, and Puget Sound Energy. Phase 1 involved 11 customer interruption cost survey activities representing a total of 24 electricity distribution service territories, 23 of them located in the Eastern and Midwestern regions of the U.S. and one located in the Pacific Northwest.[2]

[1] https://icecalculator.com/home.

[2] Phase 2 of the Initiative involves surveys of Pacific Gas and Electric, San Diego Gas and Electric, Southern California Edison, and Ameren Missouri customers. Phase 3 utilities are now being recruited.

ES Figure 1. ICE 2.0 Phase 1 Sponsoring Utilities and Distribution Service Territories

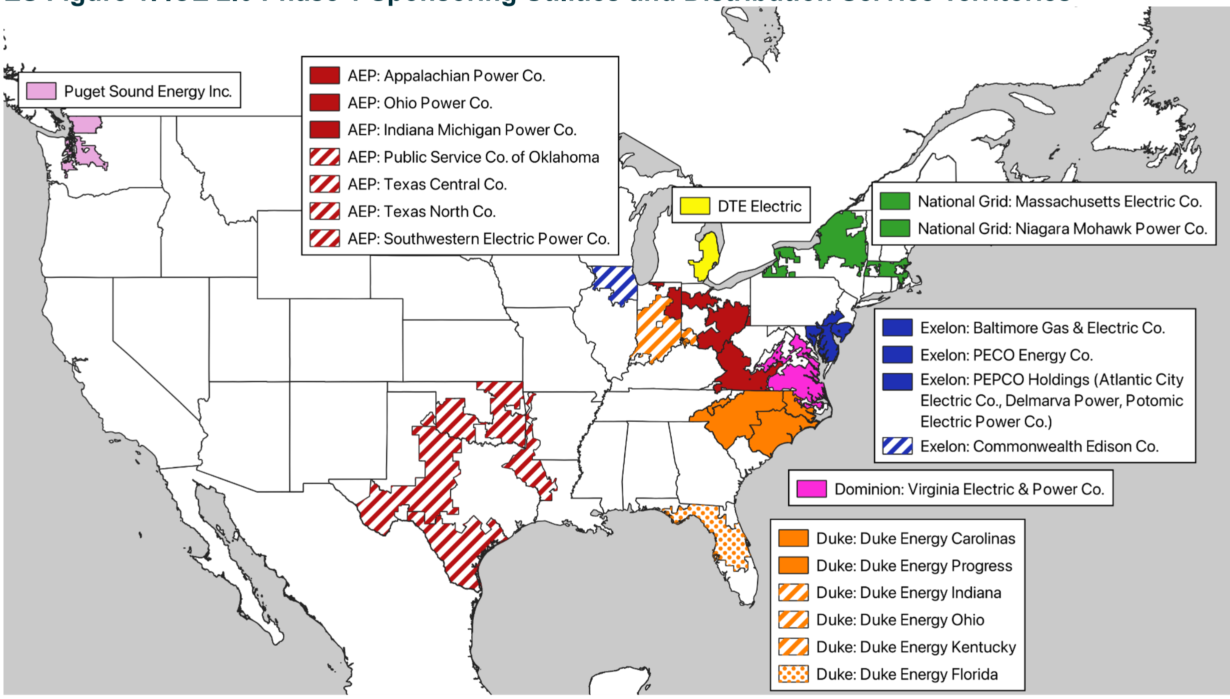

In addition to the support of sponsoring utilities, the ICE 2.0 Initiative received input from a Project Advisory Committee consisting of representatives from national organizations interested in the economic value of electricity reliability, including the U.S. Department of Energy (U.S. DOE), National Association of Regulatory Utility Commissioners (NARUC), National Association of State Utility Consumer Advocates (NASUCA), National Association of State Energy Offices (NASEO), Edison Electric Institute (EEI), American Public Power Association (APPA), National Rural Electric Cooperative Association (NRECA), and the Electric Power Research Institute (EPRI).

Rationale for Updating the ICE Calculator

The ICE 2.0 initiative was developed to address two concerns that have been raised regarding the continued usefulness of the original ICE Calculator.

First, in view of the magnitude of the investments being considered to improve reliability and address resilience, information about the prospective economic benefits of these investments must be current, robust, and easily accessible. The original ICE Calculator was developed based on surveys conducted from 1989 to 2012. As utility customer electricity usage has changed, in some instances significantly (e.g., work from home, digital controls for manufacturing processes, etc.), it is reasonable to expect that the value customers place on reliable electric service may have also changed.

Second, it is critical that the ICE Calculator is truly national in scope, given that reliability-enhancing investments are being considered throughout the country. The original ICE Calculator was developed based on surveys conducted independently by utilities located primarily on the West Coast and in the Southeast. Surveying customers in all U.S. regions ensures the ICE Calculator captures any regional differences in the value of reliable electric service.

Updated and Expanded Interruption Cost Surveys

Phase 1 survey activities involved administering updated and expanded interruption cost surveys to statistically-representative samples of the sponsoring utilities' residential and non-residential customers. In total, more than 3,000 residential and nearly 4,000 non-residential validated surveys were used to support the Phase 1 update of the ICE Calculator.

The ICE 2.0 Initiative features a number of methodological and analytical advances compared to prior customer interruption cost surveys:

- Interruption costs borne by residential customers were estimated using a state-of-the-art willingness-to-pay valuation method;

- Interruption costs for both residential and non-residential customers were estimated for scenarios ranging from momentary power interruptions (lasting up to five minutes) to sustained interruptions lasting up to 24 hours, considering different times of the day, days of the week, and seasons of the year; and

- The customers surveyed were selected by applying formal statistical sampling procedures to ensure that results are representative of all customers served. For residential customers, the procedure considered average annual electricity consumption, age, and household income. For non-residential customers, the procedure considered average annual electricity consumption and type of firm. The very largest non-residential customers were over-sampled given the extreme diversity of this customer group which, in aggregate, often represents the largest fraction of a utility's electricity sales.

ES Table 1 shows the customer populations, customers sampled, response target, responses, and response rate for the two customer segments. This table also includes the number of valid responses and valid response rate for each segment, which represents the total survey responses used in the final modeling. Appendix A details the criteria by which valid and invalid responses were identified.

ES Table 1. Survey Response Results

Segment	Customer Sampling Population	Customers Sampled	Response Target	Total Responses	Overall Response Rate	Validated Responses [3]	Validated Response Rate
Residential	22,276,695	35,743	2,750	3,316	9.3%	3,026	8.5%
Non-residential	2,141,558	90,464	3,487	4,579	5.1%	3,874	4.3%

Updated Customer Damage Functions

Analysis of pooled results from the Phase 1 surveys yields a series of econometric equations, known as customer damage functions (CDFs), that quantitatively relate the statistically significant factors contributing to power interruption costs. CDFs are the analytic engines that drive the ICE Calculator's results.

Separate CDFs were estimated for residential and non-residential customers. The residential CDF consists of a single equation, a generalized linear model (GLM). The non-residential CDF consists of two linked equations, one representing a Probit model and the other a GLM. The Probit model calculates the probability that a customer experiences a non-zero interruption cost while the GLM predicts the interruption cost of those customers with non-zero costs. ES Table 2 lists the explanatory variables included in the final CDFs for residential and non-residential customers, along with their statistical significance.

[3] Section 2.6 describes the validation procedures used to review surveys.

ES Table 2. Explanatory Variables Included in the CDFs

Residential	Non-residential	
	Probit	GLM
Duration of Interruption**	Duration of Interruption**	Duration of Interruption**
Annual kWh Usage**	Annual kWh Usage**	Annual kWh Usage**
Season**	Day of Week**	Percentage of Customers Given Advance Warning**
Percentage of Customers with Backup Generators**	Percentage of Customers Given Advance Warning**	Percentage of Customers in the Manufacturing Industry**
Percentage of Customers Working From Home*	-	Percentage of Customers in the Healthcare Industry**
Annual Household Income**	-	-

* significant at $p < 0.05$
** significant at $p < 0.01$

ES Figure 2 shows how interruption costs for residential and non-residential customers vary according to duration.

ES Figure 2. Predicted Residential (left) and Non-residential (right) Interruption Costs by Duration

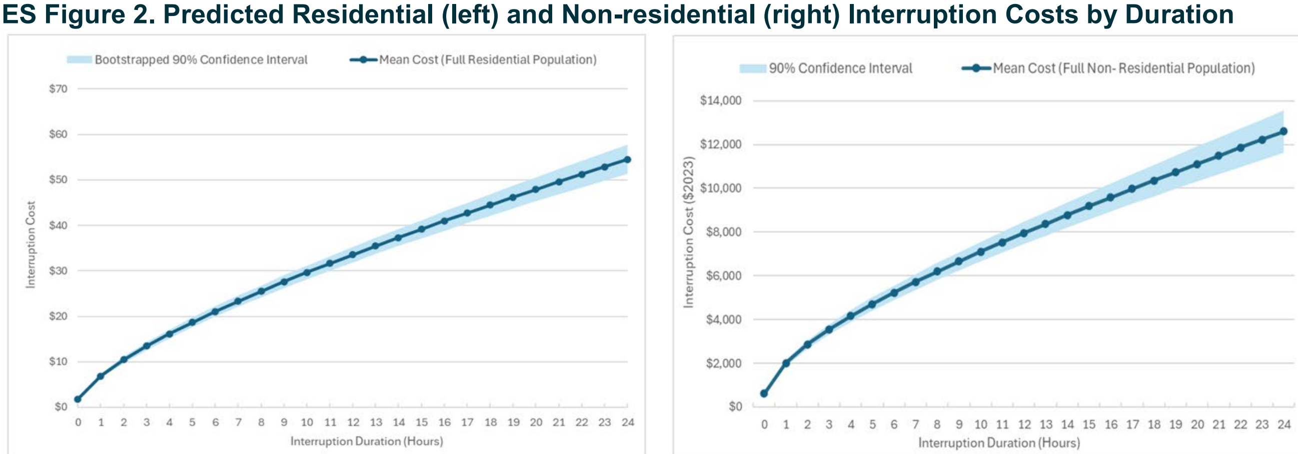

ES Table 3. ICE 2.0 Phase 1 Modeled Summary Results (2023$)

Duration of Power Interruption Event	Cost per Event	Cost per kW	Cost per Unserved kWh	Cost per CMI[4]
Residential				
Momentary	$1.80	$1.50	$18.03	$0.36
2 Hours	$10.49	$8.62	$4.31	$0.09
8 Hours	$25.55	$21.21	$2.65	$0.05
24 Hours	$54.52	$44.76	$1.86	$0.04
Non-residential				
Momentary	$609	$43	$521	$122
2 Hours	$2,839	$202	$101	$24
8 Hours	$6,172	$440	$55	$13
24 Hours	$12,646	$902	$38	$9

Additional survey questions were added to better understand customer resilience to power interruptions, the prevalence of customer-owned backup generation, and customer strategies for coping with interruptions of over 24 hours. The majority of residential customers indicated they would respond to an interruption of over 24 hours by temporarily relocating. Most non-residential customers reported they would respond to a multiday interruption by shutting down until power could be restored. These and related results are presented in Appendix C.

Next Steps

The CDFs documented in this report were used to update the ICE Calculator in late 2024 and early 2025. The functionalities of the online ICE Calculator have also been redesigned and enhanced in response to user feedback from members of our Project Executive and Advisory Committees.

At the time this document was prepared, five additional surveys were in various stages of completion (Phase 2), most of which involved customers from the U.S. West and Midwest. As a result, when these surveys are combined with the Phase 1 surveys analyzed in this report, we will be able to explore the effects of regional variations in power interruption costs. Finally, discussions are underway to survey customers in other U.S. regions that are not well represented, including those served by rural electricity cooperatives. The ICE Calculator will be updated following the completion and integration of each subsequent phase of the ICE 2.0 Initiative.

[4] Customer minutes interrupted.

1. Introduction

The ICE Calculator is a publicly available, online tool that estimates the economic costs electricity customers experience due to power interruptions.[5] It was first developed over 15 years ago for the U.S. Department of Energy by Lawrence Berkeley National Laboratory (Berkeley Lab) and Freeman, Sullivan, & Co. (Sullivan et al., 2009). The tool is used routinely by utility planners and decision makers to estimate the economic benefits of grid reliability and resilience improvements.

In 2021, Berkeley Lab and Resource Innovations, Inc. launched the “ICE 2.0 Initiative” – a national, multi-client study to refresh the underlying data and enhance the functionality of the ICE Calculator (Sullivan et al., 2018; Sullivan et al., 2019). The ICE 2.0 Initiative involves Berkeley Lab contracting with sponsoring utilities to administer identical, updated and expanded interruption cost surveys to statistically representative samples of each utility’s customers. Berkeley Lab and Resource Innovations then pool survey results across participating utilities and use them to update the analytical engines that drive the ICE Calculator.

The ICE 2.0 Initiative is being conducted in phases. Each phase involves the administration of interruption cost surveys to the customers of sponsoring utilities, followed by an update of the ICE Calculator based on analysis of pooled results from the surveys.

This report describes the activities and findings from Phase 1 of the Initiative, which was sponsored by eight utilities as shown in Figure 1.1: American Electric Power, Commonwealth Edison, Dominion Energy, Duke Energy, DTE Electric, Exelon, National Grid, and Puget Sound Energy. Phase 1 involved 11 customer interruption cost survey activities covering a total of 24 electricity distribution service territories, 23 of them located in the Eastern and Midwestern U.S. and one located in the Pacific Northwest.

[5] https://icecalculator.com/home.

Figure 1.1. ICE 2.0 Phase 1 Sponsoring Utilities and Distribution Service Territories

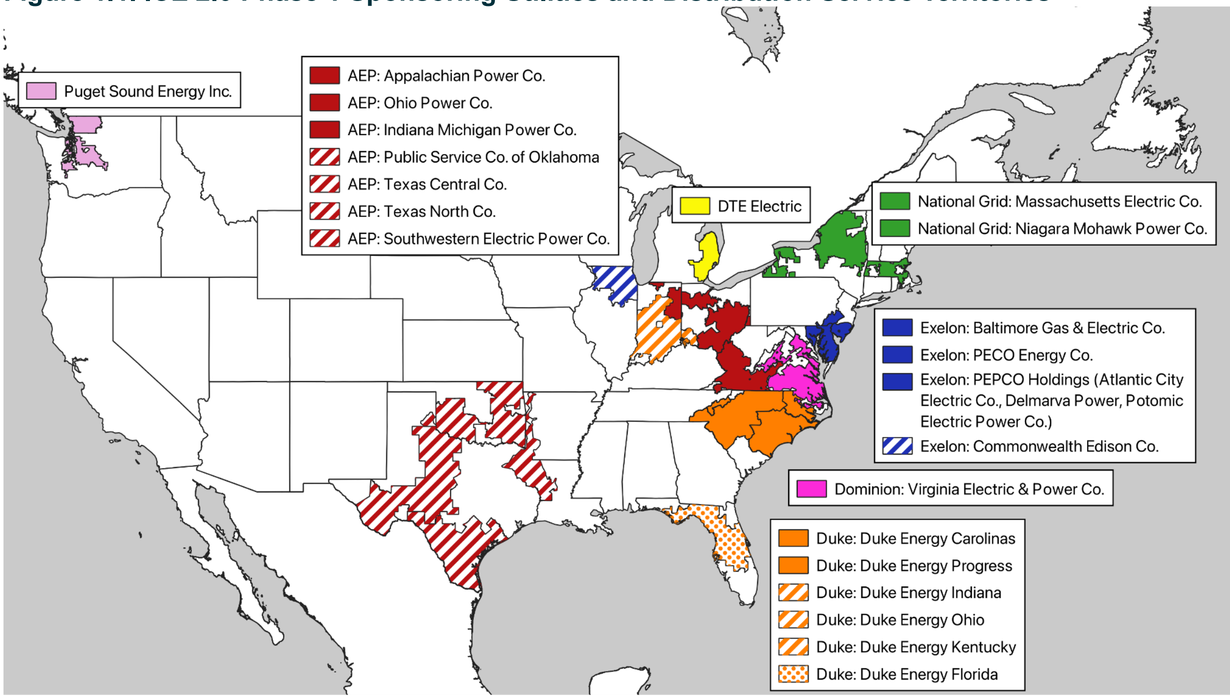

In addition to the support of sponsoring utilities, the ICE 2.0 Initiative received input from a Project Advisory Committee consisting of representatives from national organizations interested in the economic value of electricity reliability, including the U.S. Department of Energy (U.S. DOE), National Association of Regulatory Utility Commissioners (NARUC), National Association of State Utility Consumer Advocates (NASUCA), National Association of State Energy Offices (NASEO), Edison Electric Institute (EEI), American Public Power Association (APPA), National Rural Electric Cooperative Association (NRECA), and the Electric Power Research Institute (EPRI).

Rationale for Updating the ICE Calculator

The ICE 2.0 initiative was developed to address two concerns that have been raised regarding the continuing usefulness of the original ICE Calculator:

First, in view of the magnitude of the investments being considered to improve reliability and address resilience, information about the prospective economic benefits of these improvements must be current, robust, and easily accessible. The original ICE Calculator was developed based on surveys conducted from 1989 to 2012. As utility customer electricity use has changed, in some instances significantly (e.g., work from home, digital controls for manufacturing processes, etc.), it is reasonable to expect that the value customers place on reliable electric service may have also changed.

Second, it is critical that the ICE Calculator is truly national in scope, given that reliability-enhancing investments are being considered throughout the country. The original ICE Calculator was developed based on surveys conducted independently by utilities located primarily on the West Coast and in the Southeast. Surveying customers in all U.S. regions ensures that the ICE Calculator captures any regional differences in the value of reliable electric service.

This report is organized into the following sections:

- Section 2 describes the development and administration of the interruption cost surveys. We first explain the valuation methods employed in the residential and non-residential surveys. We then describe the design of the survey instruments. Next, we describe the procedures used to develop statistically representative samples of customers, as well as the procedures used to recruit customers. Finally, we present combined results from the administration of the surveys. Section 2 is supplemented by five Appendices with technical information on survey activities. Appendix A describes how the prices used in the residential valuation method were developed. Appendix B explains how response targets for customer recruitment were developed. Appendix C describes how invalid responses were identified and removed from the final set of responses used to update the customer damage functions (CDFs). Appendix D explains how customer responses were weighted to develop initial estimates of power interruption costs for later comparison to those produced by the updated CDFs. Appendix E presents findings from the specialized survey questions that were included to better understand customers' resilience to power interruptions, including the prevalence of customer-owned backup generation and the extent to which residential customers work from home.
- Section 3 describes the development of the updated residential CDF and presents selected results. We first describe the overall regression approach used to estimate the residential CDF, including the reasons for using a single-step estimation approach. We then describe the factors (i.e., explanatory variables) considered for inclusion in the

CDF and the methods used to develop and test candidate CDFs. Finally, we present the final set of factors included in the residential CDF and selected results from its application.

- Section 4 describes the development of the updated non-residential CDF and presents selected results. We first describe the overall regression approach used to estimate the non-residential CDF, including the reasons for using a two-step estimation approach. We then describe the factors considered for inclusion in the CDF and the methods used to develop and test candidate CDFs. Finally, we present the final set of factors included in the non-residential CDF and selected results from its application.
- Sections 3 and 4 are supplemented by five Appendices with supporting technical information on the development and findings from the updated CDFs. Appendix F describes the selection of the function form used to represent interruption costs over varying interruption durations. Appendix G provides more detail on the Least Absolute Shrinkage and Selection Operator (LASSO), which is the automated method used to select explanatory variables for inclusion in the CDFs. Appendix H explains how the confidence intervals associated with interruption cost estimates were developed. Appendix I describes the testing conducted to conclude that a single CDF was appropriate for the non-residential sector, rather than separate CDFs for small, medium, and large non-residential customers.
- Section 5 summarizes the Phase 1 activities described in this report, outlines planned next steps in the ICE 2.0 Initiative, and discusses caveats regarding the use of the updated ICE Calculator.

2. Survey Design, Administration, and Results

Phase 1 of the ICE 2.0 Initiative involved 11 customer interruption cost survey activities. This section describes the development of the interruption cost surveys and their administration.[6] Section 2.1 details the valuation methods employed by the residential and non-residential surveys. Section 2.2 presents the survey design. Section 2.3 describes the procedures used to develop statistically representative samples of customers, while Section 2.4 explains the customer recruitment approach. Section 2.5 presents the overall responses to the surveys. Section 2.6 describes the steps taken to prepare survey results for the development of customer damage functions. Section 2.7 presents the weighted interruption costs that result from direct analysis of the surveys.

2.1 Valuation Methods

Two valuation methods are used to measure interruption costs: (1) willingness-to-pay (WTP) for residential customers, and (2) direct cost measurement for non-residential customers. WTP measurement techniques involve estimating the amount residential customers would be willing to pay for a hypothetical backup service to avoid experiencing a given interruption. Direct cost measurement techniques involve asking non-residential customers to estimate the direct costs they would incur during a service interruption.

2.1.1 Willingness-to-pay Approach

Cost estimates for the residential segment are based on a WTP valuation approach because residential customers experience both tangible and intangible losses from power interruptions. During a short-duration interruption (i.e., up to 24 hours), a significant portion of the interruption cost for residential customers is likely the result of inconvenience, which is an intangible loss (Sullivan et al., 2019).

Rather than asking what an interruption would cost the customer, the WTP approach asks how much the customer would pay for a hypothetical backup service to avoid its occurrence. The WTP approach employs the concept of “compensating valuation”: customers are asked to estimate the economic value that would leave their welfare unchanged, compared to a situation in which no interruption occurred. This approach is especially useful when intangible costs are present, which by their nature are difficult to estimate using the direct cost measurement approach.

The specific WTP approach used to survey residential customers is called “one-and-one-half-bound dichotomous choice” (OHDC) contingent valuation (Cooper et al., 2002). OHDC is a type of choice experiment that is used to measure the value of a good or service that does not exist (e.g., perfectly reliable power).

---

[6] The survey activities described in this section were developed following a roadmap that was prepared in 2019 and updated in 2022 to guide the ICE 2.0 Initiative (Sullivan et al., 2019; Resource Innovations, 2022; Lawrence Berkeley National Laboratory, 2023).

The OHDC procedure is implemented in the survey by first describing a hypothetical but plausible backup service that could be purchased to avoid the interruption (Sullivan et al., 2019). Next, respondents are told that while the exact cost to purchase the service is unknown, it is believed to be within a range of two discrete prices. Respondents are then asked a series of questions based on their willingness to purchase the service starting from a pre-determined initial price.

OHDC contingent valuation is a form of closed-ended contingent valuation, where respondents are given a series of explicit prices and asked if they would be willing to pay one of them. This is in contrast to open-ended contingent valuation, where respondents are asked to state explicitly the price they would be willing to pay. In recent years, economists have preferred closed-ended to open-ended contingent valuation, as evidence indicates respondents understate their WTP when answering an open-ended framework (Cooper et al., 2002; Sullivan et al., 2019). Research has shown that, among closed-ended elicitation frameworks, the OHDC framework strikes a balance between increasing statistical efficiency relative to a single-bounded structure (e.g., asking only one WTP question) while also minimizing potential respondent bias relative to a double-bounded structure (e.g., asking all respondents two WTP questions).

Figure 2.1 illustrates the flow of questioning for an OHDC example scenario. As an example, the respondent is first asked if they would purchase the service for $5. If they agree to that price, then they are asked if they would be willing to purchase the service for $10. Alternatively, if they respond that they would not purchase the service for $5, then they are not asked any other questions. A similar logic applies if the respondent is first asked if they would purchase the service for $10. If they say they would purchase the service, there are no follow-up questions. If they respond “no”, then they are asked if they would purchase the service for $5.

Appendix A presents the prices used in the surveys and describes how they were tested and developed.

Figure 2.1. Survey Questioning Using One-and-one-half-bound Contingent Valuation

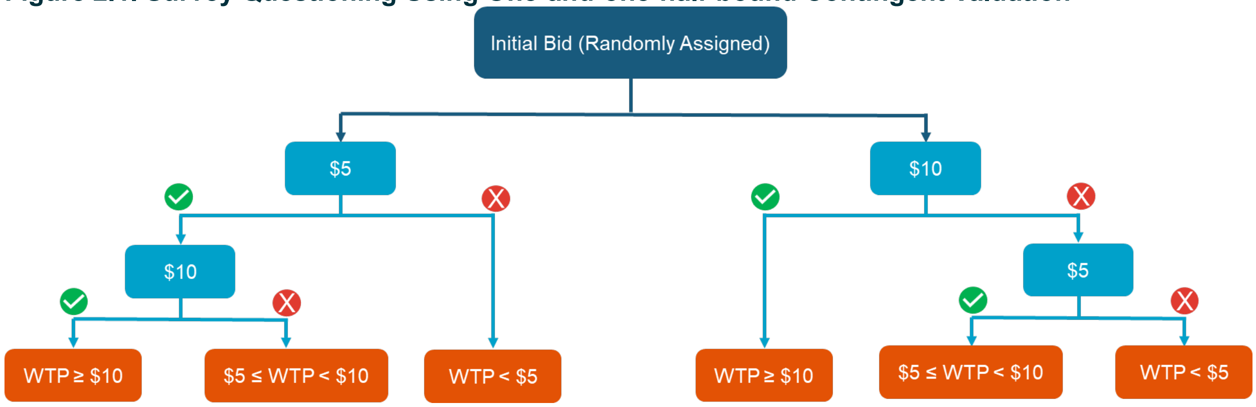

2.1.2 Direct Cost Measurement

For non-residential customers, direct cost measurement was used in this study because their interruption costs are more tangible and therefore easier to estimate than residential customers. The direct cost of an interruption is defined as follows in Equation 2.1 (Sullivan et al., 2018; Sullivan et al., 2019):

Equation 2.1. Direct Cost Measurement

$$\begin{aligned} Direct\ Cost = Value\ of\ Lost\ Production + \ Interruption\ Related\ Costs \\ - Interruption\ Related\ Savings \end{aligned}$$

The *Value of Lost Production* is the amount of revenue the customer's organization would have generated in the absence of the power interruption, minus the amount of revenue it was able to generate despite the power interruption. That is, it represents the net loss in economic value of production while accounting for the possibility of making up for lost production.

Interruption Related Costs are additional production costs directly incurred because of the interruption. These costs include:

- Labor costs to make up any lost production (if they can be made up);
- Labor costs to restart the production process;
- Material costs to restart the production process;
- Costs resulting from damage to input feedstocks;
- Costs of re-processing materials (if any); and
- Cost to operate backup generation equipment.

Interruption Related Savings are production cost savings resulting from the interruption. When production or sales cannot take place, economic savings result because inputs to the production or sales process cannot be used. For example, while electric power is interrupted, the enterprise cannot consume electricity and thus will experience savings on their electric bill. In many cases, savings resulting from interruptions are small and do not significantly affect interruption cost calculations. However, for manufacturing enterprises where energy and feedstock costs account for a significant fraction of production costs, these savings may be quite significant and should be subtracted from the other cost components to ensure interruption costs are not double-counted. These savings include:

- Savings from unpaid wages during the interruption (if any);
- Savings from the cost of raw materials not used because of the interruption;
- Savings from the cost of fuel not used; and
- Scrap value of any damaged materials.

In measuring interruption costs, only the incremental losses resulting from interruptions are included in the calculations. Incremental losses include only those costs deemed above and beyond normal production costs. If the customer can make up some percentage of their production loss at a later date (e.g., by running the production facility during times when it would otherwise be idle), the interruption cost does not include the full value of the production loss. Rather, it is calculated as the value of production not made up plus the cost of additional labor and materials required to make up the share of production eventually recovered.

To determine these costs, non-residential customers were asked to estimate a series of itemized costs, including lost revenue, restart costs, and damage to equipment, for the first interruption scenario they complete (an unexpected summer weekday "base" scenario at their first randomly assigned duration). Following this, respondents were presented with an estimated total cost based on their itemized costs, which they could confirm or adjust as needed. For all subsequent interruption scenarios, non-residential respondents were asked how their total cost would change relative to the first scenario. This format ensured that respondents were only required to answer questions relating to itemized costs once, saving time and increasing ease of taking the survey.

2.2 Survey Instrument Development

Three survey instruments were developed: one for residential customers; one for small- and medium-sized non-residential customers (SMNR), defined as having average hourly consumption of less than 200 kW; and one for large non-residential (LNR) customers, defined as having average hourly consumption of greater than 200 kW. This section describes the development of the interruption cost surveys, with a focus on the specification of the interruption scenarios.

Each survey began with a series of introductory questions asking respondents about their household's demographic or organization's firmographic characteristics, their experience with prior power interruptions, and any resilience measures they have in place (e.g., backup generators). Next, the surveys presented a series of scenarios varying in interruption duration. Each respondent was randomly assigned three of the following four durations[7], which span the range of interruption durations targeted by the ICE Calculator:

- Momentary (up to 5 minutes)

[7] Respondents are presented with only three interruption duration scenarios in order to reduce survey fatigue.

- 2 hours
- 8 hours
- 24 hours

For each of the three assigned durations, respondents were initially introduced to the interruption as being both unexpected and taking place on a summer weekday at a given onset time (either 9 AM, 2 PM, or 7 PM), labeled the “base” scenario. Following each base scenario, respondents were then randomly presented with one (for residential and SMNR customers) or two (for LNR customers) randomly assigned “pivots,” with each pivot changing a circumstance of the base scenario. The pivot conditions included:

1. The interruption still occurs during a summer weekday, but the customer receives advance warning.
2. The interruption still occurs during a weekday without warning, but in winter instead of summer.
3. The interruption still occurs during summer without warning, but on a weekend instead of a weekday.

Each residential and SMNR respondent was presented with only one pivot from the base scenario, so that all other circumstances of the new interruption scenario remained the same as the base scenario. For instance, if a respondent received the advance warning pivot, their second interruption scenario would still occur in the summer, on a weekday, at the same onset time. This was to ensure that respondents had to focus on only one change to their interruption circumstances at a time.

Respondents received the same pivot(s) for all three durations. For example, if they received the advance warning pivot for the first duration presented, they would receive it for the following two durations as well. Table 2.1 provides an example set of interruption scenarios.

Table 2.1. Example Set of Interruption Scenarios for an Individual Respondent

Scenario	Season	Time of Week	Onset Time	Advance Warning	Duration	Pivot
A	Summer	Weekday	2:00 PM	No	5 minutes or less	Weekend
B	Summer	Weekday	2:00 PM	No	2 hours	Weekend
C	Summer	Weekday	2:00 PM	No	24 hours	Weekend

Residential and SMNR respondents received questions involving three unique short-duration interruption scenarios of varying duration, each with a base scenario and a pivot scenario, resulting in six total interruption scenarios presented to each respondent. LNR respondents also received questions involving three unique interruption durations, but each included two randomly assigned pivots, resulting in nine total interruption scenarios.

At the end of the survey, each respondent was also presented with a hypothetical three-day interruption scenario that was randomly assigned to take place in summer or winter. Unlike the short-duration scenarios, the questions for this long-duration scenario asked customers how they would respond (i.e., what they would do), not what they would be willing to pay to avoid it or incur as a direct cost. Consequently, the resulting responses were not intended to support the development of updated CDFs for the ICE Calculator.[8]

Given that the OHDC method for structuring the WTP questions had not been implemented in prior interruption cost studies, the residential survey underwent cognitive testing with a random sample of more than 500 participants.[9] This testing helped ensure that survey questions were easy to understand, survey length was acceptable, and respondents could provide meaningful results to the OHDC WTP questions.

As a final step, the survey instruments were also provided to all members of both the Project Executive Committee (i.e., the participating Phase 1 utilities) and the Project Advisory Committee for review.

The final survey instruments are publicly available on the ICE website.[10]

2.3 Sample Design

The study aimed to collect the following numbers of completed surveys from each Phase 1 utility:

- 250 residential customers
- 250 SMNR customers[11]
- 67 LNR customers

These targets were set to provide samples sufficient in size to estimate interruption costs separately for each utility. Additionally, the sample sizes were established to provide an aggregate survey dataset for all Phase 1 utilities sufficient in size to estimate the CDFs used by the ICE Calculator.

Before describing the sample design methodology, it is important to note that a “customer” refers to an entire premises for SMNR and LNR customers, not an individual electric service

[8] See Appendix E.
[9] The non-residential survey was not subjected to cognitive pre-testing because direct cost measurement has been used successfully in many prior interruption cost studies of non-residential customers.
[10] https://icecalculator.com/home.
[11] LNR customers were defined as having an average hourly demand of greater than 200 kW, while non-residential customers below 200 kW were classified as SMNR. For Puget Sound Energy, the SMNR and LNR customer threshold was set at an average demand of 25 kW because its non-residential customer base had significantly lower usage compared to other Phase 1 utilities.

account. When SMNR and LNR business customers complete an interruption cost survey, they provide answers for all their accounts at a specified address. Both usage and customer contact information were aggregated across all accounts associated with each business at a premises, and it is these customers, so defined, that were sampled. For the residential segment, a “customer” refers to an individual account because it is rare that a residential customer has multiple accounts at a single address. A residential “customer” refers to an individual household at a specified address.

The first step in drawing samples from the total utility customer population was to divide the population into three segments: residential, SMNR, and LNR. Next, the population in each segment was further divided and grouped into a fixed number of strata according to average annual electricity demand (in kilowatts). A percentage of potential survey respondents in each stratum was then drawn based on the proportion of usage accounted for by the population in each stratum relative to the total usage of the segment.[12]

The goals of stratifying the sample were twofold: (1) to maximize the likelihood of receiving a targeted number of survey completions based on customer usage; and (2) to ensure that survey respondents represent the demographic and usage characteristics of the customers in the utility’s service area.

Drawing samples based on customer usage is necessary because the distribution of usage per customer is highly skewed. As shown in Figure 2.2, while the vast majority of customers are clustered towards the lower end of the usage distribution, there is a long tail of high-usage customers towards the upper end of the distribution. Considering that usage is a proxy for interruption costs, an objective of the sample design methodology was to ensure that a representative share of high-usage customers was included in the sample (Sullivan et al., 2019). A simple random sample would not accomplish this objective because high-usage customers account for a small percentage of the total number of customers and therefore would have a very low probability of being selected.

[12] See Appendix B for additional details regarding this sampling strategy.

Figure 2.2. Distribution of Average Hourly Usage by Customer Class (Top 5th Percentile for Each Customer Class Omitted)

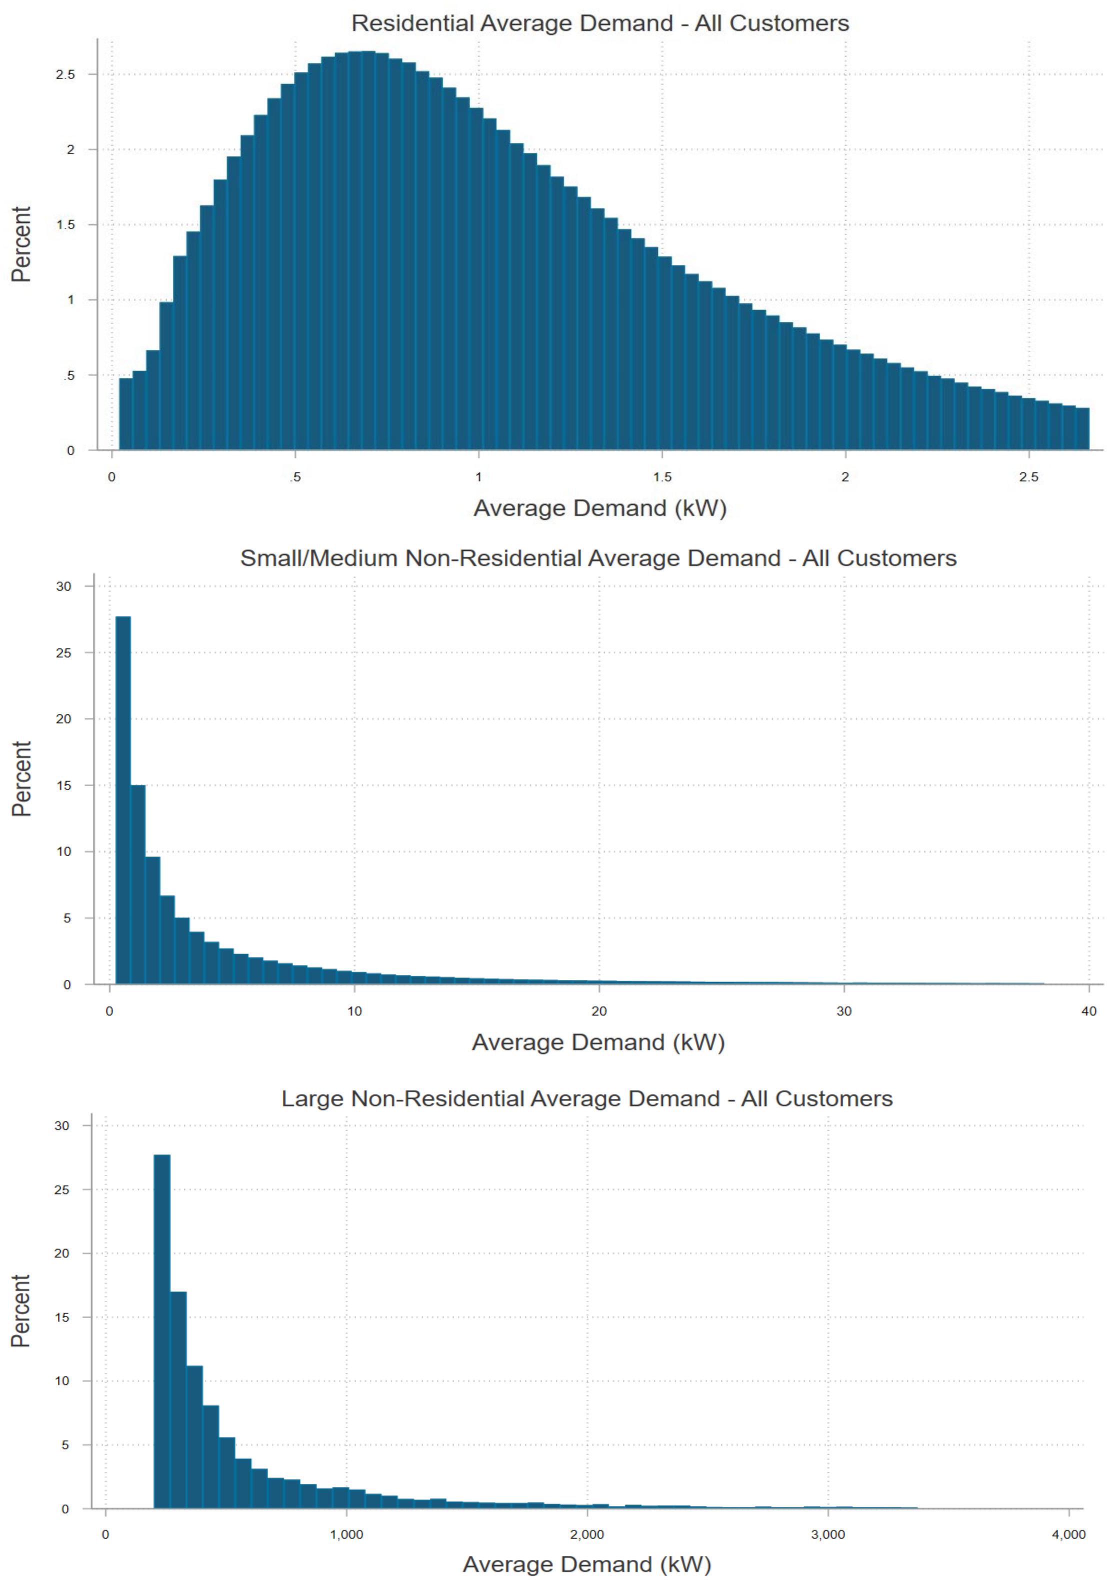

Table 2.2, Table 2.3, and Table 2.4 summarize the sample design for residential, SMNR, and LNR customers for the Phase 1 utilities. The customer population values represent the number of customers eligible to receive the survey, stratified into five usage categories. The percentage of customers sampled in each stratum was proportional to the total energy usage in that stratum as a percentage of total segment usage. For example, consider the 5 to 10 kW stratum in the residential segment. The total combined usage of respondents in this stratum represented approximately 2.2% of total electricity usage in the entire residential segment. Therefore, 2.2% of the total residential sample was selected from this stratum. This sample design ensured that a larger proportion of high-usage customers were included compared to a simple random sample. High-usage customers are more likely to experience higher and more variable interruption costs, making their representation in the sample essential. This sampling strategy allows for greater precision than a simple random sample of the entire population (Sullivan and Keane, 1995). This is because within-stratum variation is much smaller than the overall variation found in a random, non-stratified sample of the entire population (Sullivan et al., 2019).

Table 2.2. Residential Phase 1 Sample Design by Usage Category

Stratum	Usage Category (Average kW)	Customer Population	% of Population	Response Target	% of Sample
1	0–0.5	4,123,242	18.5%	105	3.8%
2	0.5–1	6,459,746	29.0%	524	19.1%
3	1–2	8,261,525	37.1%	1,196	43.5%
4	2–5	3,342,026	15.0%	866	31.5%
5	5–10	90,156	0.4%	59	2.2%
Total		**22,276,695**	**100%**	**2,750**	**100%**

Table 2.3. SMNR Phase 1 Sample Design by Usage Category

Stratum	Usage Category (Average kW)	Customer Population	% of Population	Response Target	% of Sample
1	0.25–2	1,035,733	49.1%	102	3.7%
2	2–10	699,857	33.1%	417	15.2%
3	10–50	293,371	13.9%	980	35.6%
4	50–100	53,597	2.5%	588	21.4%
5	100–200	28,804	1.4%	663	24.1%
Total		**2,111,362**	**100%**	**2,750**	**100%**

Table 2.4. LNR Phase 1 Sample Design by Usage Category

Stratum	Usage Category (Average kW)	Customer Population	% of Population	Response Target	% of Sample
1	200–400	15,973	52.9%	85	11.5%
2	400–1,000	8,829	29.2%	101	13.7%
3	1,000–2,000	2,866	9.5%	79	10.7%
4	2,000–5,000	1,510	5.0%	94	12.8%
5	>5,000	1,018	3.4%	378	51.3%
Total		**30,196**	**100%**	**737**	**100%**

2.4 Survey Recruitment Procedures

The Phase 1 surveys were conducted between December 2022 and April 2024. The majority of recruitment materials contained approved utility branding (e.g., logos). This section summarizes the survey recruitment procedures implemented for each customer segment.[13]

2.4.1 Residential Customers

Residential customers were initially recruited with a letter and an email containing a link to an online questionnaire. The letters and the email were coordinated to reach the customers around the same time. Approximately one week after the letters and emails were sent, customers who did not complete the online survey received a reminder email. All respondents who completed the survey were offered a $20 electronic gift card or check, which they could also decline.

2.4.2 Small and Medium Non-residential Customers

SMNR customers were also initially recruited with a letter and an email with a link to an online questionnaire where they could state their interest in participating and select a preference for completing the survey. SMNR customers were offered the opportunity to complete their surveys with the assistance of an interviewer (either over the phone or in an online meeting) or by completing the survey online without assistance. Approximately one week after the letters and emails were sent, customers who had not responded were sent a reminder email. An incentive of $50 was offered to respondents who completed the survey, which could also be declined.

2.4.3 Large Non-residential Customers

The customer recruitment strategy for LNR customers featured the active involvement of utility account representatives. Since LNR customers are often the hardest to contact, account representatives were asked to confirm customer contact information and provide additional outreach support by informing them about the survey.

[13] All survey administration process and procedures were first reviewed and approved under Pro00023294 to ensure adherence with all applicable Institutional Review Board guidelines.

As with SMNR customers, LNR customers were initially recruited with an email and letter. They were also given a choice of completing the survey with or without the assistance of an interviewer. One week after the survey link was emailed and the letter sent, respondents were given a reminder email. Additionally, for some utilities, customers with a valid phone number were contacted by a trained interviewer to complete the survey over the phone. An incentive of $100 was offered to respondents who completed the survey, which they could also decline.

Table 2.5 summarizes the survey implementation approaches for each customer class.

Table 2.5. Survey Implementation Approach by Customer Class

Customer Class	Sample Design Target per Utility	Recruitment Method	Data Collection Approach	Valuation Approach	Incentive Offered
Residential	250	Letter, Email	Online Survey	WTP	$20
SMNR	250	Letter, Email	Online Survey, Interview	Direct Cost	$50
LNR	67	Letter, Email, Phone	Online Survey, Interview	Direct Cost	$100

Certain customers were deemed ineligible for survey recruitment. These included customers missing more than two out of the most recent twelve months of electricity usage data, because calculating average demand (used to stratify customers by usage) would likely be inaccurate for these customers. In addition, residential customers with a calculated average demand of 10 kW or greater were not included because such high-usage residential accounts are likely to be non-residential facilities that have been misidentified. Non-residential customers with an average demand of less than 0.25 kW were not included because non-residential accounts with very low usage are often associated with customers who are not actively managing energy consumption and typically have much lower survey response rates. Finally, non-residential accounts with North American Industry Classification System (NAICS) codes indicating “Real Estate, Rental and Leasing” or with customer names including keywords such as “Apartments,” “Homes,” “Communities,” and “Condos” were not included to avoid inadvertently sampling residential tenants and landlords.

2.5 Overall Responses to Survey

Table 2.6 summarizes the responses to the Phase 1 surveys of residential customers. With 3,316 total completed residential surveys, the customer response exceeded the overall sample design target of 2,750. Overall, the survey had a 9.3% response rate, with the 2 to 5 kW usage category having the highest response rate, and the 5 to 10 kW usage category having the lowest response rate.

Table 2.6. Phase 1 Survey Response Summary (Residential)

Usage Category (Average kW)	Customer Sampling Population	Customers Sampled	Response Targets	Responses	Response Rate	Responses/ Target Responses
0–0.5	4,123,242	1,360	105	133	9.8%	126.7%
0.5–1	6,459,746	6,781	524	633	9.3%	120.8%
1–2	8,261,525	15,542	1,196	1,365	8.8%	114.1%
2–5	3,342,026	11,293	866	1,125	9.9%	129.9%
5–10	90,156	767	59	60	7.8%	101.7%
Total	**22,276,695**	**35,743**	**2,750**	**3,316**	**9.3%**	**120.6%**

Table 2.7 summarizes the responses to the Phase 1 surveys of non-residential customers. SMNR customers are reflected in the first five demand categories (0.25 to 200 kW) and LNR customers are reflected in the largest five demand categories (200 kW to >5,000 kW). With 4,579 total completed surveys, customer response was greater than the combined non-residential sample design target of 3,487. Overall, the non-residential surveys had a 5.1% response rate. By non-residential segment, response rates were highest in the 2,000 to 5,000 kW category at 7.3%. Response rates were lowest in the 100 to 200 kW categories at 4.1%.

Table 2.7. Phase 1 Survey Response Summary (Non-residential)

Usage Category (Average kW)	Customer Population	Customers Sampled	Response Targets	Responses	Response Rate	Responses/ Target Responses
0.25–2	1,035,733	5,456	102	303	5.6%	297.1%
2–10	699,857	19,438	417	1,062	5.5%	254.7%
10–25	293,371	31,611	980	1,525	4.8%	155.6%
25–100	53,597	12,961	588	595	4.6%	101.2%
100–200	28,804	7,820	663	324	4.1%	48.9%
200–400	15,973	6,437	85	381	5.9%	448.2%
400–1,000	8,829	3,833	101	221	5.8%	218.8%
1,000–2,000	2,866	1,475	79	75	5.1%	94.9%
2,000–5,000	1,510	885	94	65	7.3%	69.1%
> 5,000	1,018	548	378	28	5.1%	7.4%
Total	**2,141,558**	**90,464**	**3,487**	**4,579**	**5.1%**	**131.3%**

Although the total number of responses far exceeded the overall target, response targets for some strata were not met. For example, the largest usage strata (>5,000 kW) had a target of 378 responses but received only 28 responses. In this regard, it is important to note that response targets by strata were chosen to boost the accuracy of the customer damage functions with respect to the largest usage customers. Thus, even though the number of responses fell short of some strata target, the representation in the largest strata was greatly increased compared to a simple random sample. For instance, the number of non-residential customers in the population with an average demand greater than 5,000 kW is 1,018 out of 2,141,558, or 0.05%. However, the 28 responses received from this stratum represents 0.6% (as a percentage of the 4,579 responses received for the entire sample), which is over 12 times greater than the percentage of the total population in this stratum. See Appendix D for more information on how the survey responses were weighted to reflect the total population of utility customers.

2.6 Preparation of Results for Development of Customer Damage Functions

Before the Phase 1 survey results were combined into a residential and a single non-residential[14] dataset for use in updating the CDFs in the ICE Calculator, each individual response was reviewed to confirm its validity. Separate procedures were developed for each survey type. For the residential survey, the reviews focused on removing responses that were

[14] As discussed in Section 4, a single non-residential CDF was developed from combined SMNR and LNR survey responses. Appendix I documents the analysis conducted to determine that development of a single rather than two distinct CDFs for non-residential customers was warranted.

not based on economic considerations (e.g., unwilling to pay any price to avoid an interruption for reasons other than economic ones) or that were internally inconsistent (e.g., a willingness to pay more to avoid a shorter interruption than a longer one). For the non-residential sector, reviews focused on removing responses that were internally inconsistent (in a manner akin to that used to remove inconsistent residential responses) or represented extreme outliers (i.e., reporting costs that, on a normalized basis, were statistically significant and exceeded the 75th interquartile range). Appendix C provides more details on how these criteria were used to identify and remove individual responses from the final datasets that were used to update the CDFs.

Table 2.8 displays the number of validated responses included in the final datasets per Phase 1 sponsoring utility. Following the review process, the final residential dataset included 3,026 responses. The non-residential dataset included 3,874 responses. Both final datasets exceeded the original targets by more than 10%.

Table 2.8. Phase 1 ICE Calculator Validated Residential and Non-residential Responses by Utility

Utility	Validated Residential Responses	Validated Non-residential Responses
AEP East	314	342
AEP West	263	301
ComEd	259	369
Duke Energy Carolinas	270	404
Duke Energy Florida	267	367
Duke Energy Midwest	280	384
DTE Electric	271	351
Dominion Energy	281	288
Exelon	270	294
National Grid	275	350
PSE	276	424
Total	**3,026**	**3,874**

2.7 Survey-based Customer Interruption Costs

In this section, we present estimates of customer interruption costs that emerge directly from analysis of the survey results. These estimates will be used in Sections 3 and 4 to support the development of aspects of the CDFs and as benchmarks against which to compare the results from the final CDFs.

Table 2.9 presents estimates of residential and non-residential customer interruption costs that were calculated directly from analysis of the surveys for each of the four interruption event durations. Survey responses were first weighted to be representative of the overall population of customers following procedures that are described in Appendix D.[15]

Table 2.9. Survey-based Interruption Costs (mean and +/- 90% confidence interval bounds)

Duration	Residential	Non-residential
Momentary	$2.17 +/- $0.17	$772 +/- $75
2 hours	$7.50 +/- $0.47	$2,771 +/- $153
8 hours	$30.04 +/- $1.60	$6,568 +/- $310
24 hours	$52.47 +/- $2.72	$11,677 +/- $500

[15] For the residential findings, a bootstrapping methodology was used to estimate the standard error associated with each mean value. “Bootstrapping” is a method that involves repeated sampling of a dataset and is commonly implemented for data where the population statistics are unknown.
For the non-residential findings, confidence intervals were calculated using the standard error estimated directly from the survey responses. The standard error of the weighted average interruption costs is clustered by respondent to account for serial correlation between responses from a given customer, which would otherwise falsely increase the precision of the confidence intervals. Two-sided confidence intervals are then calculated from the standard error to represent the 90% confidence interval.

3. Residential Customer Damage Function

This section describes the development of the updated residential CDF and presents selected results. The overall process to develop the residential CDF is presented in Section 3.1, including the rationale for using a single-step estimation approach. Section 3.2 describes the factors (i.e., explanatory variables) considered for inclusion in the residential CDF and the methods used to develop and test candidate CDFs. Section 3.3 presents the final set of factors selected for the residential CDF, including a variety of results from its application.

3.1 Overview of the Development Process

The residential CDF was developed from the survey responses in four steps. The first involved translating results from the OHDC contingent valuation questions into a single interruption cost for each of the four interruption event durations (momentary, 2 hours, 8 hours, and 24 hours). The second step involved specifying a continuous form anchored by these point estimates to produce interruption costs for any interruption duration lasting up to 24 hours. The third step was to select a functional form for the CDF. The fourth step was to choose from among the available explanatory factors to develop a final specification for the CDF. The first three steps are discussed in this subsection. The fourth step is discussed in Section 3.2.

The first step in identifying the regression formula was to translate responses from the OHDC questions into a single interruption cost. The method used was based on previous studies, which have shown that the response (yes or no) to a given price can be translated to a lower-bound interruption cost estimate by dividing it by the probability density associated with receiving that price (Watanabe, 2010).[16]

Specifically, due the nature of the OHDC framework, WTP values are expressed as the percentage of respondents who accepted or rejected a given price. Figures 3.1 through 3.4 display the percentage of residential respondents who indicated they would pay a given price to avoid the interruption. In general, the amount respondents are willing to pay to avoid an interruption steadily decreases as the price increases.[17] For example, during a 24-hour interruption, 85.1% of respondents are willing to pay $5 to avoid the interruption but only 26.4% are willing to pay $120.

[16] Note that these values were presented in Table 2.9.

[17] In some cases, the percentage of customers willing to pay for a higher bid is greater than the percentage willing to pay for a lower bid for the same duration interruption (e.g., the percentage willing to pay $0.20 in Figure 3.1 is less than the percent willing to pay $0.25 for the same interruption). These inconsistencies are likely due to random variation in the respondents who received each bid, and not an increasing willingness-to-pay with price.

Figure 3.1. Momentary Interruption: Percent of Respondents Willing to Pay to Avoid the Interruption by Price

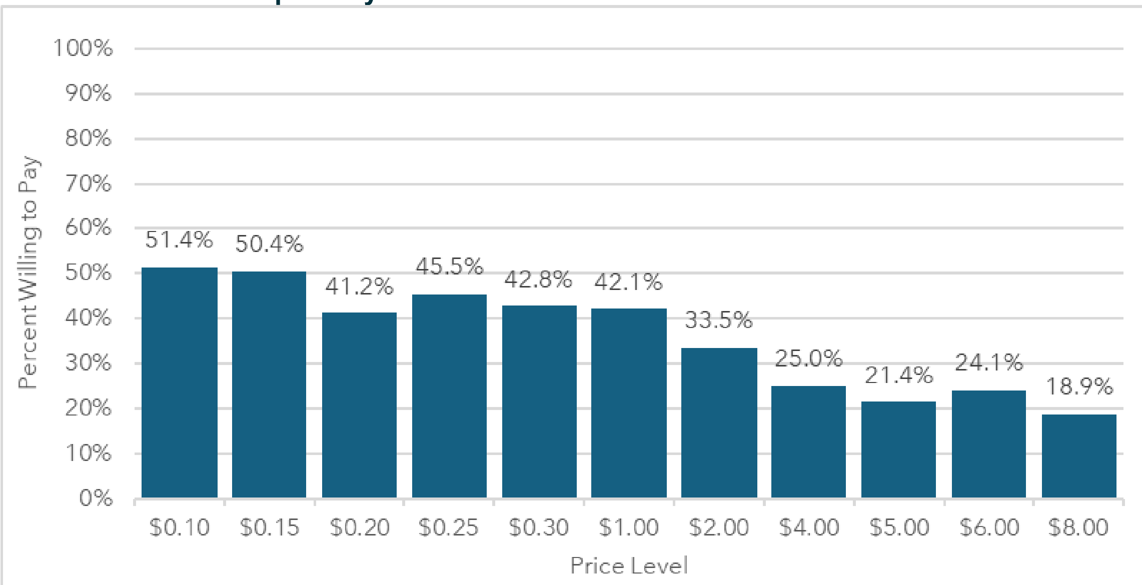

Figure 3.2. Two-hour Interruption: Percent of Respondents Willing to Pay to Avoid the Interruption by Price

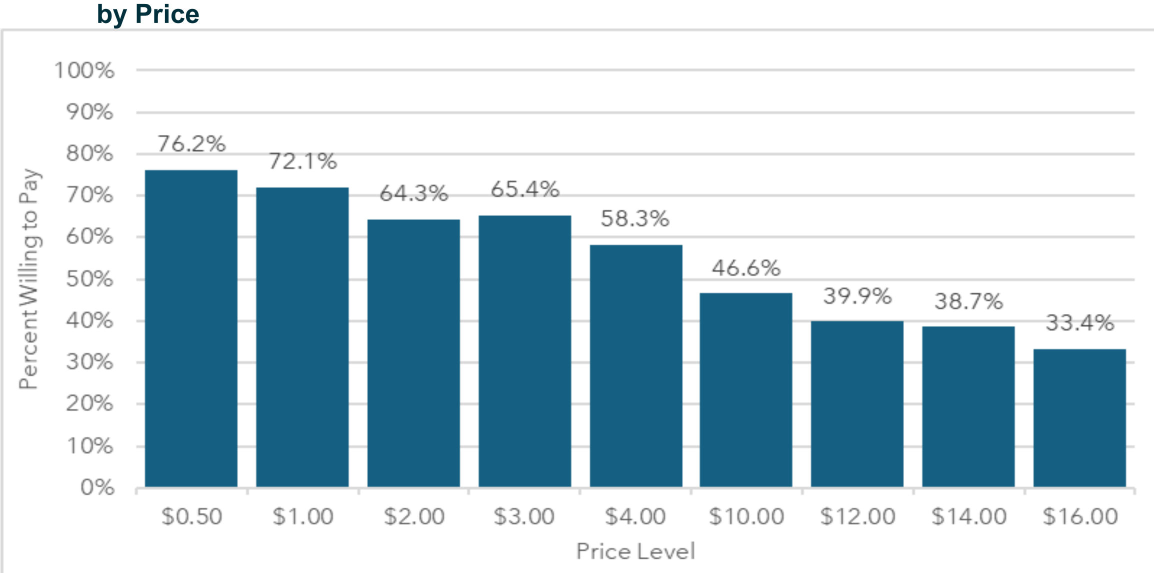

Figure 3.3. Eight-hour Interruption: Percent of Respondents Willing to Pay to Avoid the Interruption by Price

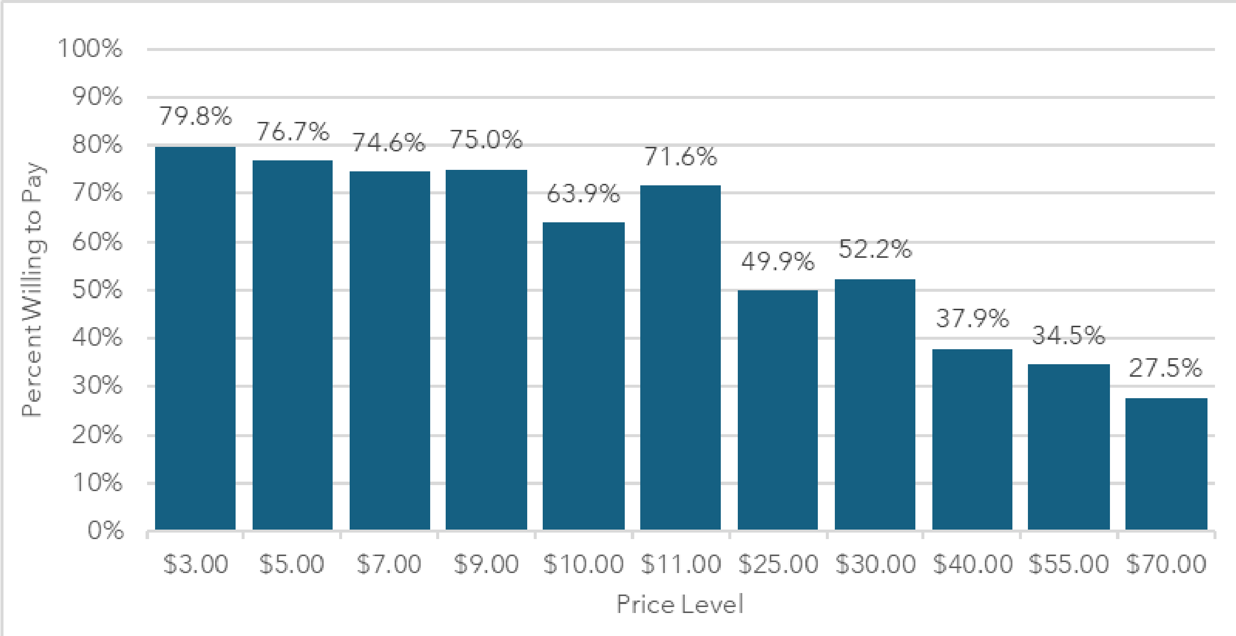

Figure 3.4. 24-hour Interruption: Percent of Respondents Willing to Pay to Avoid the Interruption by Price

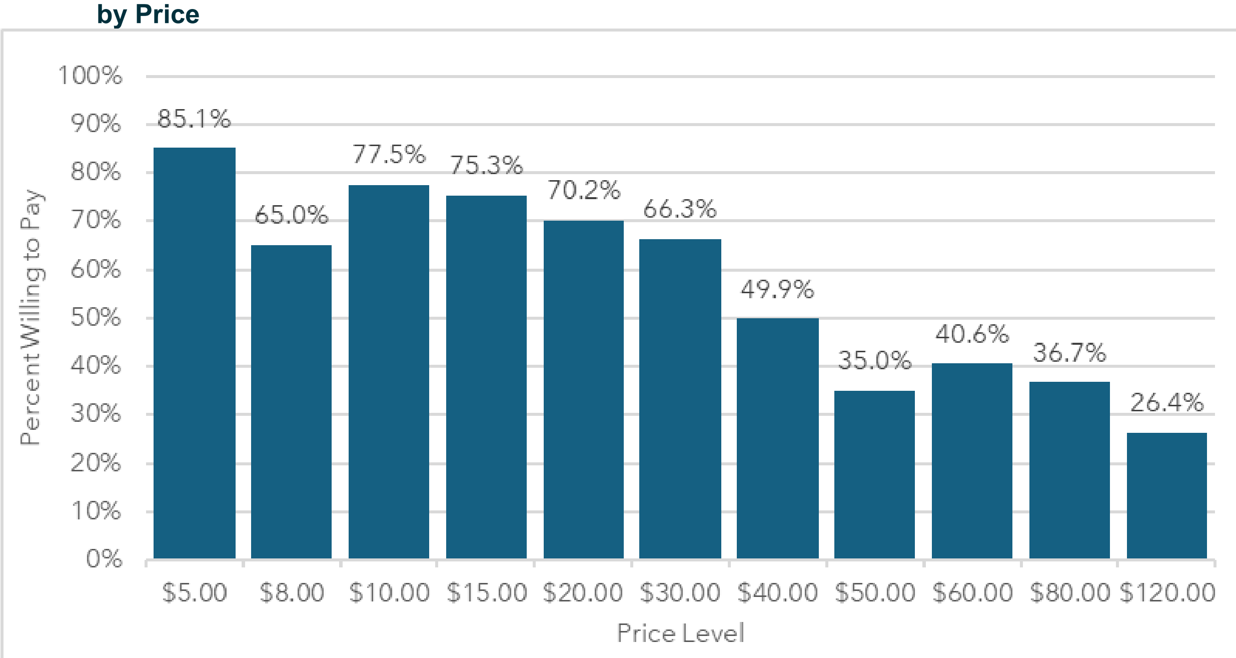

In examining the WTP by price in Figures 3.1 through 3.4, only Figure 3.1 shows a substantial percentage of respondents unwilling to pay the lowest price, indicating that for durations exceeding a momentary interruption, WTP is consistently non-zero. Even among respondents who reject the lowest price, it is plausible that their WTP is greater than zero. A true WTP of zero indicates that the consumer derives no utility from having power and is therefore completely indifferent to an outage occurring or not occurring, which is very unlikely. As a result of these observations, we selected a single-part regression model for the residential CDF.[18]

The second step involved developing a continuous form to express interruption costs for any duration ranging from a momentary interruption to a 24-hour interruption. For the residential customer segment, duration was plotted against average costs to determine the most appropriate model functional form. The possible functional forms were restricted to those that monotonically increased in cost with respect to duration. In other words, the functional form must not predict cost to decrease as duration increases in any region from 0 to 24 hours. The resulting plots indicated that a functional form including the natural log of duration and duration squared best approximates the relationship between interruption duration and cost. As such, the natural logarithm of duration to the first and second power is used. Appendix F provides additional details on the determination of this functional form.

The third step involved selecting a functional form for the residential CDF. In modeling interruption cost within a regression framework, the outcome variable is expressed as a function of a vector of explanatory variables, typically denoted as X'. Instead of using a typical ordinary least squares regression, Poisson Quasi-Maximum Likelihood Estimator (QMLE) was employed. This estimator has the advantage of yielding consistent estimates of the mean interruption cost without requiring the modeler to know the true distribution of interruption costs (Wooldridge, 1999). This estimator also ensures that outcomes will always be positive in value, which is appropriate for WTP.[19]

As opposed to an ordinary least squares regression, where the outcome variable Y is modeled as a linear combination of a vector of covariates X' plus an error term and the vector of coefficients β is selected by minimizing the sum of squared errors, the Poisson QMLE models the *natural log* of Y as a function of X' (Wooldridge, 1999).

As a quasi-maximum likelihood model, the regression runs a series of iterations with different values of the vector of coefficients β and selects the coefficients β that maximize the log-likelihood of generating the distribution of WTP observed in the sample. For this reason, the outcome variable is the expected value of Y (conditional on the covariates X') and, as a generalized linear model, the regression does not include an error term (as seen with OLS). Equation 3.1 shows the general regression specification for the Poisson QMLE. This could also be expressed as estimating $E(Y|X')$ as a function of $e^{X_i'\beta}$.

Equation 3.1. Regression Specification for Poisson Quasi-maximum Likelihood Estimator

$$log(E(Y|X')) = X_i'\beta$$

[18] The non-residential CDF is a two-part regression model because the non-residential survey yielded a significant number of valid zero interruption costs. For more information, see Section 4.
[19] A negative WTP for electric power would imply that households and firms would be willing to pay to not consume electricity.

Equation 3.2 presents the general functional form for the ICE Calculator residential CDF, where the expected value of the cost of interruption *d* for customer *i* ($E(Cost)_{i,d}$) is a function of D, a vector of duration-specific terms; α, the vector of regression coefficients associated with each of those terms; X, a vector of household- and interruption-specific characteristics; and β, the coefficients associated with each of those characteristics.

Equation 3.2. General Residential Regression Specification for ICE Calculator Model

$$E(Cost)_{i,d} = e^{D'_{i,d}\alpha + X'_i\beta}$$

3.2 Selection of Explanatory Variables

The explanatory factors or variables considered for inclusion in the residential CDFs were selected from the interruption- and household-specific characteristics used in or collected through the surveys. Table 3.1 lists the variables that were tested for inclusion in the residential model.

Table 3.1. Residential Potential Model Variables

<table>
<tr><th colspan="2">Continuous Variables</th></tr>
<tr><td colspan="2">● Interruption duration (in minutes)
● Annual electricity usage (in kWh)
● GDP per kWh (collected at the state level)</td></tr>
<tr><th colspan="2">Categorical Variables</th></tr>
<tr><td>Interruption Onset Time
● Morning
● Midday
● Evening
Season
● Summer
● Winter
Day of Week
● Weekday
● Weekend
Advance Warning
● Yes
● No
Previous Interruption in Last 12 Months
● Yes
● No</td><td>Persons in Household
● 1-2 people
● 3+ people
Ownership of Backup Generation
● Yes
● No
Work from Home
● Yes
● No
Age of Respondent
● Under 40 years
● 40-70 years
● 70+ years
Total Household Income
● Under $50,000 per year
● $50,000-$100,000 per year
● $100,000-$150,000 per year
● Over $150,000 per year
Housing Type
● Apartment/Condominium
● Attached Single-Family
● Detached Single-Family
● Mobile Home
● Unknown/Other</td></tr>
</table>

The selection of explanatory variables for the residential CDF was implemented as follows. First, the Least Absolute Shrinkage and Selection Operator (LASSO) regression was used to initially select from among the list of potential variables. Next, the selected LASSO candidate models were then individually tested using 100-fold cross-validation, which is a machine learning technique often used to assess model prediction accuracy. Lastly, the cross-validation results were reviewed in light of other design considerations in order to select a final model.

The first step in residential model selection used LASSO regression as an initial guide for selecting variables for inclusion in the CDF. LASSO regression is a regression method commonly used for variable selection that is designed to avoid overfitting, and results in accurate but parsimonious models (Desboulets, 2018).

The LASSO regression operates by using a penalty term, λ, to include or exclude potential variables according to their contribution to the overall explanatory power of the regression. LASSO regressions were run with a series of decreasing penalty terms, λ. As λ is lowered (decreasing the penalty term), the number of variables included in the regression model increases. This process revealed which variables were most important by the order in which they appeared in the regression as λ decreased. Appendix G provides additional details on the LASSO regression.

Table 3.2 presents the results of the LASSO regressions. The highest four λ values (1, 0.5, 0.25, and 0.1) were specified manually, and the lowest λ value, 0.03, was selected by statistical software that minimized the out-of-sample deviance, which is a metric for goodness-of-fit for non-linear models.

Generally speaking, the regressions show that continuous variables for usage (in kWh), a binary variable for the highest income class (household income of $150,000 or more), and a binary variable for backup generation were the strongest predictors of interruption cost.

Table 3.2. Residential Explanatory Variables as Selected by LASSO (Duration Terms Always Included)[20]

Model #	λ	Usage	Income Category			Backup Gen	Work From Home	Interruption Start Time		Season	House-hold Size	Day of Week	Warning	Age		Housing Type
			$150k	$100k–150k	$50k–100k			After-noon	Evening					>70 years	40–70 years	
Base-line	--	*x*														
1	1	*x*	*x*			*x*										
2	0.5	*x*	*x*	*x*		*x*	*x*			*x*						
3	0.25	*x*	*x*	*x*		*x*	*x*	*x*		*x*	*x*					
4	0.1	*x*	*x*	*x*	*x*	*x*	*x*	*x*		*x*	*x*	*x*	*x*	*x*		*x*
5	0.03*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*	*x*

*This is the λ value optimized by minimizing the out-of-sample deviance.

[20] Interaction terms between variables were tested, both in the LASSO model selection and in the second-stage model testing described below, but these terms were not found to improve out-of-sample model performance.

The results of the LASSO regressions were used as a starting point for the development of a series of final candidate models. The final candidate models were then tested using a cross-validation process that closely emulated the way the final residential CDF would be used in the ICE Calculator.

In the OHDC methodology, the average of the outcome variable, the weighted yes/no vote, identifies the mean WTP. However, the individual vote values do not have an intuitive interpretation, as many of these individual values will be zero if the respondent rejected a given bid, even if WTP is greater than zero. Furthermore, in the final ICE Calculator, the residential CDF predicts the mean WTP as a function of a series of average inputs. For instance, instead of receiving individual 1 (“Yes”) or 0 (“No”) values for the backup generation binary variable, the input for the final ICE Calculator is the percentage of backup generation in the population for which the user wishes to estimate interruption costs.

For the above reasons, the final candidate models were tested using an iterative cross-validation process that compared the mean predicted outcome variable to the mean actual outcome variable using a series of training-testing splits. The cross-validation process was implemented as follows (see also Figure 3.5):

1. The data were divided randomly into a training and testing dataset using a 90% training and 10% testing split.
2. For the test dataset, the mean weighted yes/no vote variable (equal to the mean WTP) and mean values of all the predictor variables were stored.
3. Each of the candidate models was trained on the training dataset, then fed the average inputs of the test set as described above. The mean WTP predicted using these average inputs was then compared to the actual mean WTP for the test set.
4. This process was repeated 100 times, with 100 distinct training/testing splits for each final candidate model.

Figure 3.5. Residential Cross-validation Process

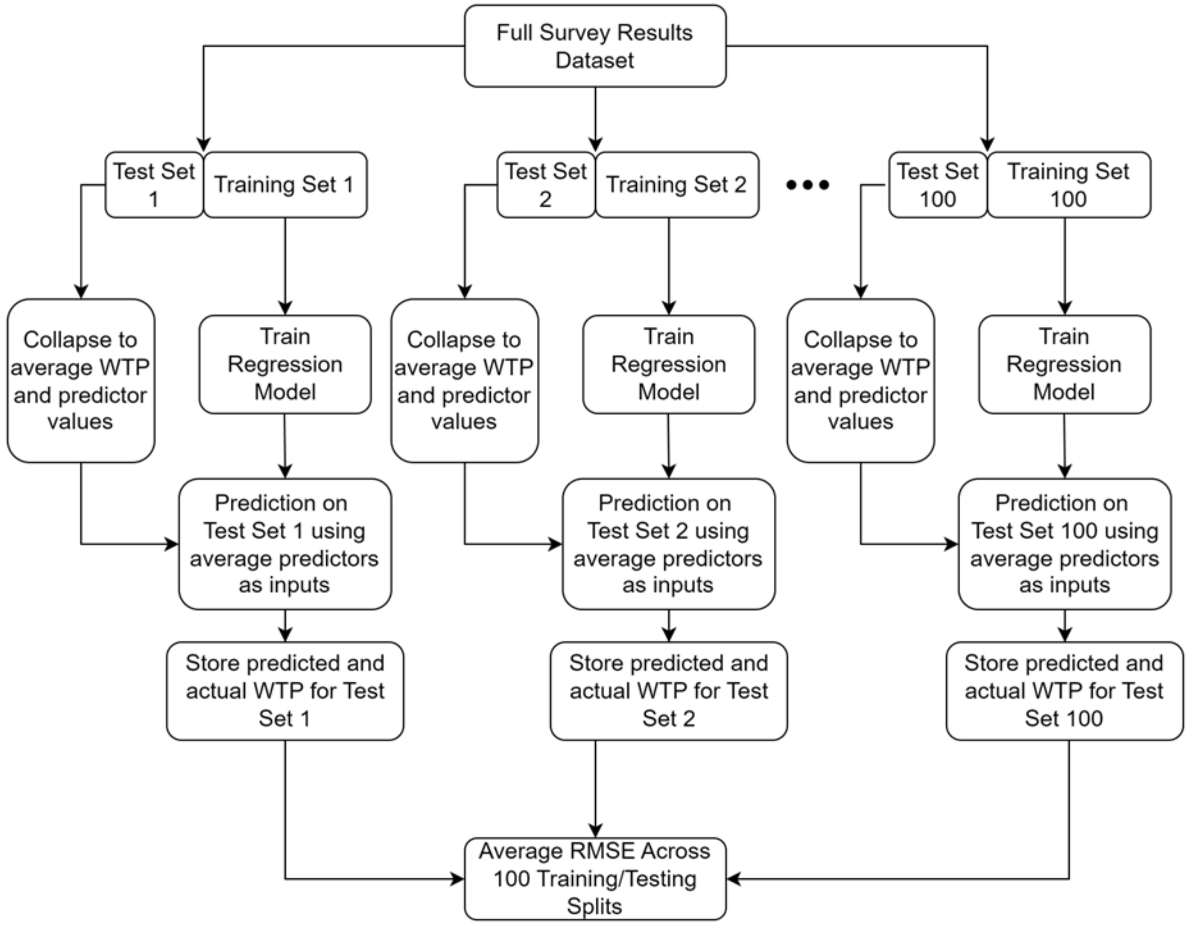

The cross-validation process results in 100 distinct actual and predicted WTP values for each model tested. The process has two important advantages over individual-level cross-validation for the purpose of testing for the ICE Calculator. First, using a series of mean WTP values from 100 different trials provides more interpretable values for metrics such as Root Mean Squared Error (RMSE) and mean absolute percentage error (MAPE), relative to using individual-level data. In this process, RMSE and MAPE would be calculated using the distribution of actual and predicted average WTP across the 100 trials, as opposed to using the difference between individual-level predicted WTP and the individual-level weighted vote values. As noted above, these individual vote values do not have an intuitive interpretation. While it would be possible to calculate RMSE and MAPE using the individual-level vote values, the wide variation in these values, relative to mean WTP itself, makes the interpretation more difficult. Second, testing 100 distinct training/testing splits ensures that any given model performance is not dependent on any single training/testing split, and instead reflects the average performance of the model.

Table 3.3 lists the models tested and the results of the cross-validation. Model 2 in Table 3.3 was initially identified by the team as an option that balanced predictive accuracy and simplicity. However, this model, while including categorical variables for the highest two income groups, did not include a categorical variable for the lowest income group (annual household income less than $50,000). Recognizing that estimating interruption costs for low-income households was an important area of interest, we included this variable to ensure that users would be able model costs separately for customers in this income group. Therefore, a variation of Model 2, Model 2a, which included this lowest income category, was also evaluated in the analysis.

The models were evaluated using several metrics, including MAPE across all durations, Akaike Information Criteria (AIC), RMSE across all four durations (momentary, 2 hours, 8 hours, and 24 hours), and RMSE for each individual duration. As MAPE and RMSE are error metrics, smaller values for these metrics indicate better-performing models, as they indicate that the models' predictions were closer to the actual cost values. RMSE was evaluated for each individual duration to ensure that the final model selected performed well in predicting interruptions of varying lengths. For simplicity, only RMSE for all durations combined is shown in Table 3.3.

Table 3.3. Residential Model Variable Selection

Model #	Rationale for Inclusion	Model	RMSE (All Durations)	RMSE (Momentary)	RMSE (2-hour)	RMSE (8-hour)	RMSE (24-hour)
Base-line	Includes just duration terms	duration	4.426	0.511	3.187	5.715	5.940
1	Lambda = 1	duration, annual kWh, backup gen , >$150k	4.356	0.541	2.957	6.142	5.398
2	Lambda = 0.5	duration, annual kWh, backup gen, >$150k, $100k - $150k, season, work from home (WFH)	4.357	0.549	2.874	6.279	5.285
2a	Model 2 with lowest level of income added	duration, annual kWh, backup gen, 4 income levels, season, WFH	4.357	0.550	2.856	6.288	5.285
3	Lambda = 0.25	duration, annual kWh, backup gen, >$150k, $100k - $150k, season, WFH, afternoon, household size	4.361	0.554	2.824	6.343	5.248
4	Lambda = 0.1	duration, annual kWh, backup gen, 4 income levels, season, WFH, afternoon, household size, weekend, warning, >70 years old, housing type (detached)	4.369	0.555	2.814	6.369	5.251
5	LASSO optimized Lambda (0.03)	duration, annual kWh, backup gen, 4 income levels, season, WFH, 3 start times, household size, weekend, warning, all age levels, housing type	4.384	0.558	2.798	6.412	5.254

*In the table, duration and annual kWh are simplified from ln(duration) + $\ln(\text{duration})^2$ and annual kWh + annual kWh^2 + annual kWh^3 for readability.

.

The models showed relatively low variation in their average out-of-sample performance in the metrics tested. The lowest RMSE value for all durations was 4.356 (Model 1) and the highest RMSE value for all durations was 4.426 (Baseline Model). With this similarity in model performance, Model 2a was selected as the final residential CDF to strike a balance between predictive accuracy and model simplicity.

In selecting the final residential CDF, we first sought to avoid making the model unnecessarily complicated, especially given that complexity yields diminishing returns in predictive accuracy, which would make the ICE Calculator less user-friendly. Second, we prioritized the inclusion of categorical variables not selected by LASSO if they represented an important segment of the population and if including them had minimal impacts on model accuracy. For example, we ultimately chose to include the lowest-earning income group categorical variable in the residential model to ensure that users could examine the impacts of interruptions on low-income households.

3.3 Final Specification of the Residential CDF and Selected Results

Table 3.4 presents the final variables selected for the residential CDF. Both continuous variables, such as duration (in minutes) and annual usage (in kWh), and categorical variables such as season and income level were selected.

Table 3.4. Final Residential Model Specifications

Continuous Variables
● Interruption duration (in minutes) ● Annual electricity usage (in kWh)
Categorical Variables
Season ● Summer ● Winter **Ownership of Backup Generation** ● Yes ● No **Work from Home Status** ● Yes ● No **Total Household Income** ● Under $50,000 per year ● $50,000-$100,000 per year ● $100,000-$150,000 per year ● Over $150,000 per year

*The functional form for duration and usage in the final model are $\ln(\text{duration}) + \ln(\text{duration})^2$ and annual kWh + annual kWh^2 + annual kWh^3.

Figure 3.2 shows the residential interruption costs estimated by the residential CDF for all durations, from a momentary interruption (up to 5 minutes) to 24 hours. These results were produced by using the average characteristics from the full set of residential customer responses described in Section 2. As a result, these results can be compared directly to the interruption costs derived from the survey responses (see Table 2.9).

Figure 3.6. Residential CDF vs. Survey-based Interruption Costs

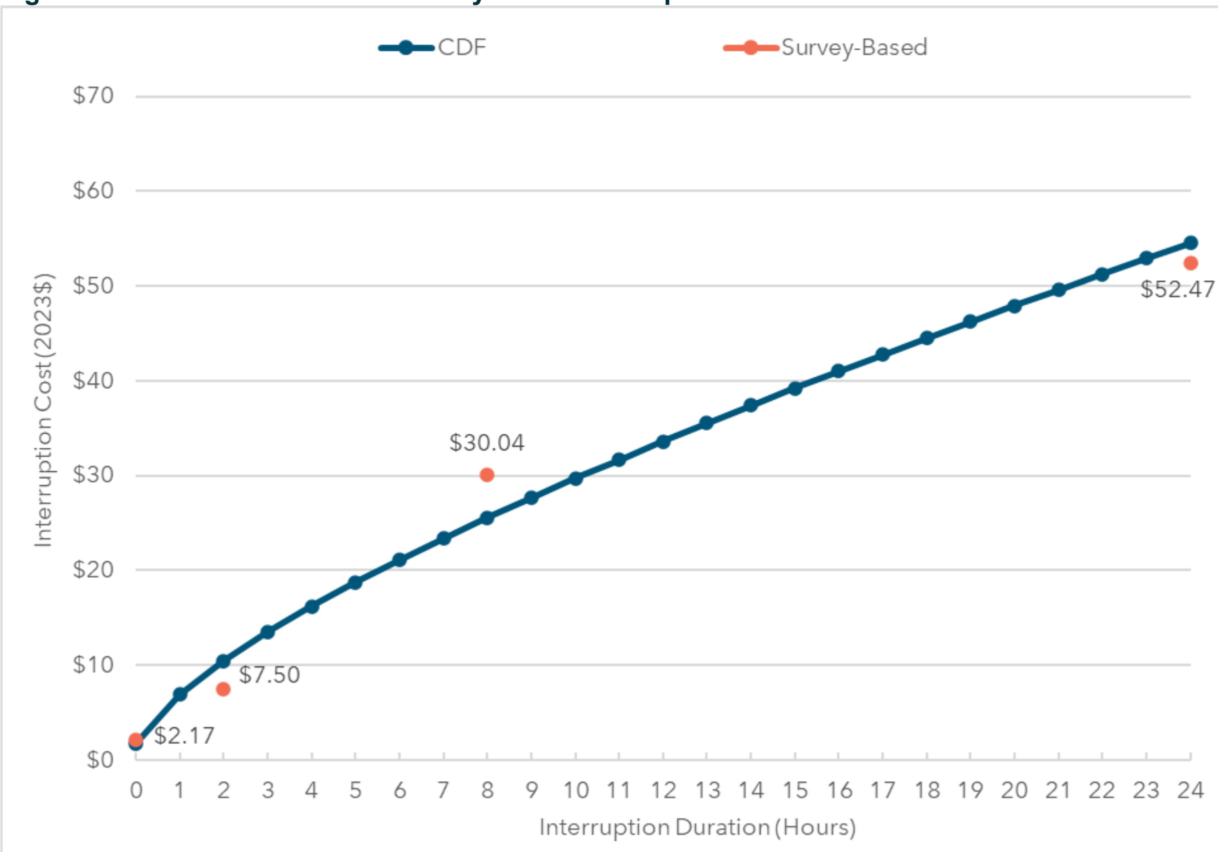

Table 3.5. Residential Interruption Costs (2023$)

Duration of Power Interruption Event	Cost per Event	Cost per kW	Cost per Unserved kWh	Cost per CMI
Momentary	$1.80	$1.50	$18.03	$0.36
2 Hours	$10.49	$8.62	$4.31	$0.09
8 Hours	$25.55	$21.21	$2.65	$0.05
24 Hours	$54.52	$44.76	$1.86	$0.04

Figure 3.7 through Figure 3.11 represent how residential interruption costs vary for selected explanatory variables. The figures were developed by holding all explanatory variables at their population-average values.

Figure 3.7 shows how interruption costs vary by annual energy usage. Generally, residential customers with higher usage experience higher interruption costs. A customer in the 5th percentile of annual usage (1,959 kWh per year) has an estimated interruption cost of $52.19

for a 24-hour interruption, whereas a customer in the 95th percentile of annual usage (23,653 kWh per year) has an estimated interruption cost of $57.82 for a 24-hour interruption.

Figure 3.7. Residential Interruption Cost by Usage

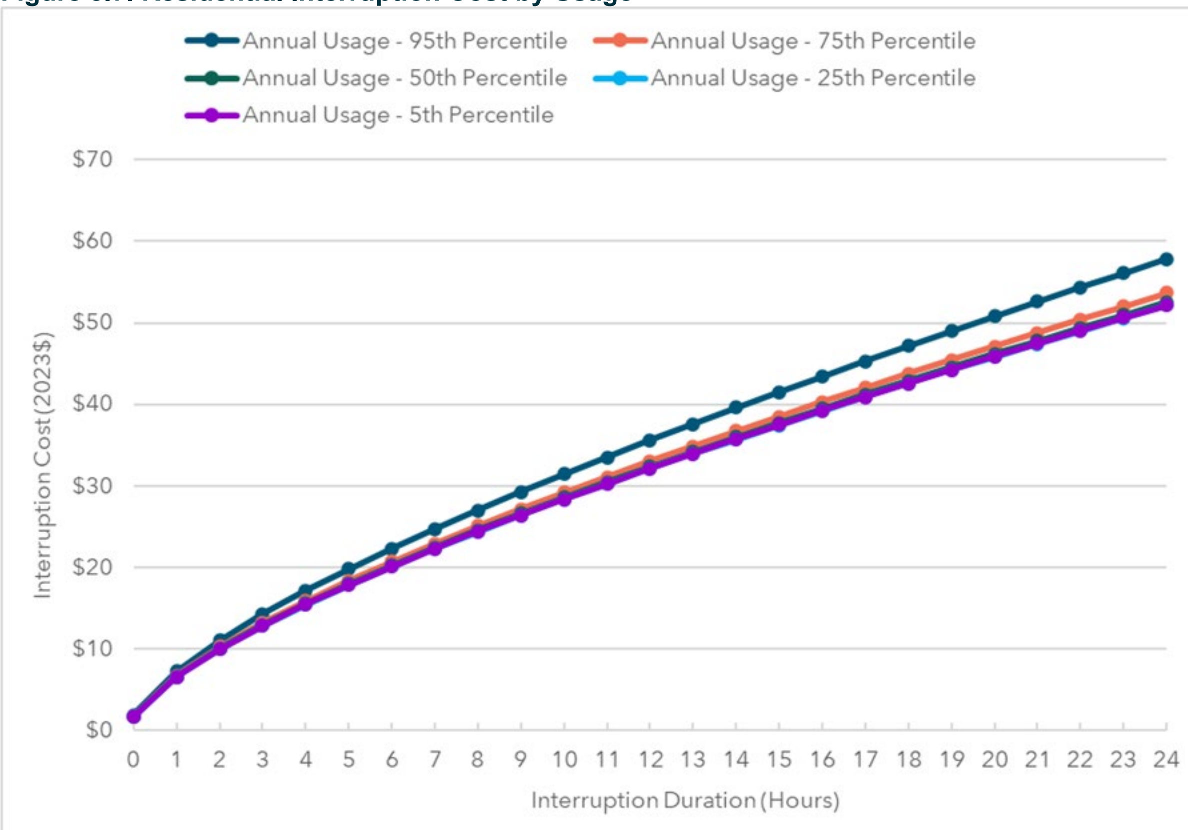

Figure 3.8 shows how interruption costs vary by season. On average, costs are estimated to be higher for winter interruptions compared to summer interruptions. For example, the estimated cost of a 24-hour winter interruption is $60.28, while the estimated cost of a 24-hour summer interruption is $49.46.

Figure 3.8. Residential Interruption Cost by Season

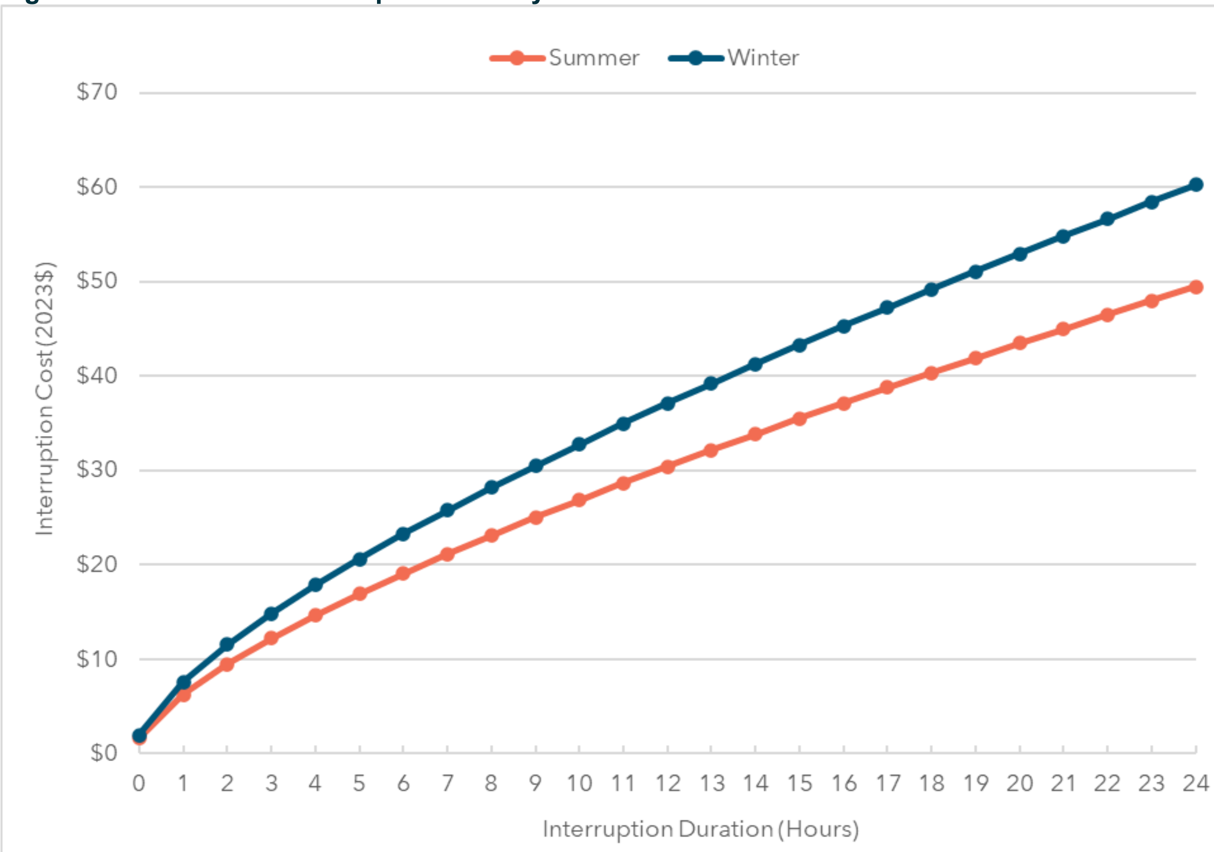

Figure 3.9 shows interruption costs vary by income category. On average, interruption costs increase as the income level increases. For example, a 24-hour interruption has an estimated cost of $47.48 for the low-income category (under $50,000 per year), while it has an estimated cost of $62.02 for the highest income category (over $150,000 per year).

Figure 3.9. Residential Interruption Cost by Income Category

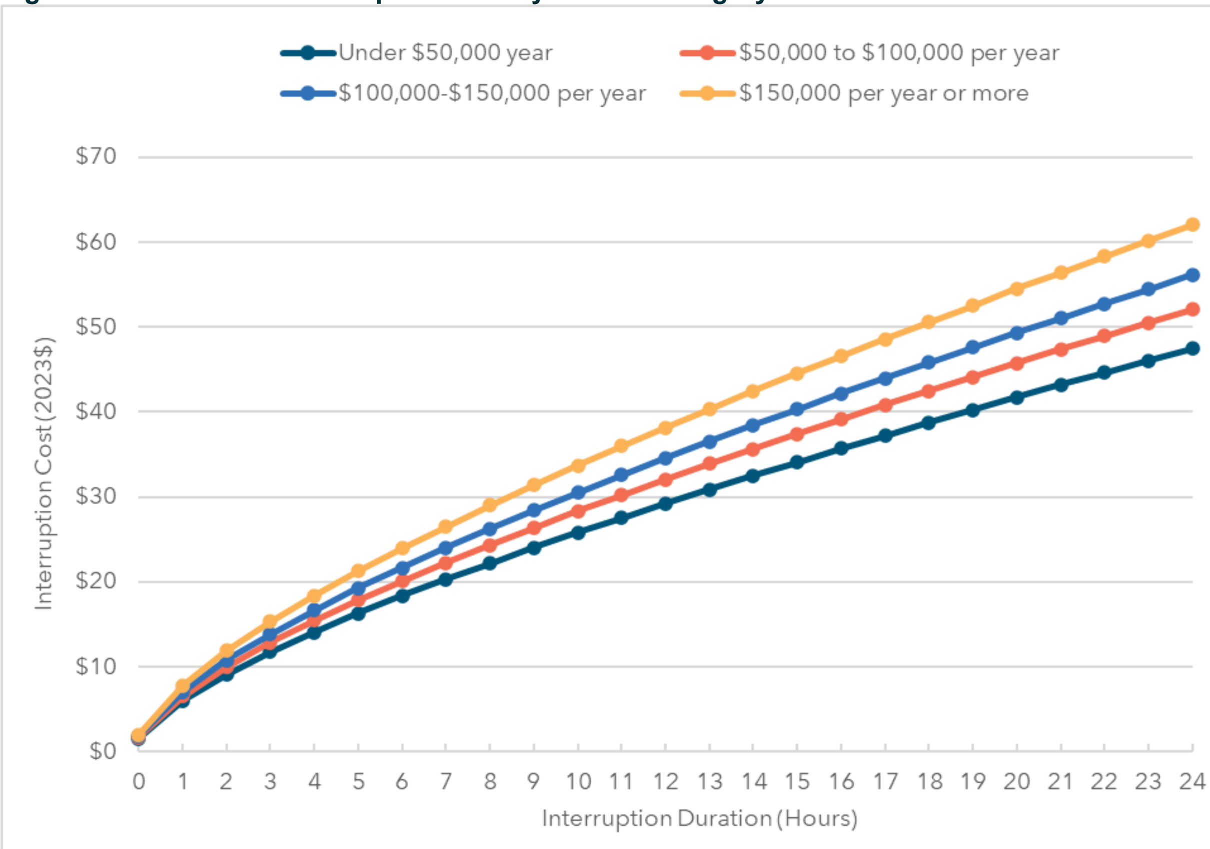

Figure 3.10 presents how interruption costs vary according to the presence or absence of a backup generator. As expected, interruption costs are higher for households without backup generators. For example, the interruption cost for a 24-hour interruption for households without a generator is estimated to be $56.40, while for households with a generator it is estimated to be $40.53.

Figure 3.10. Residential Interruption Cost by Backup Generator Ownership

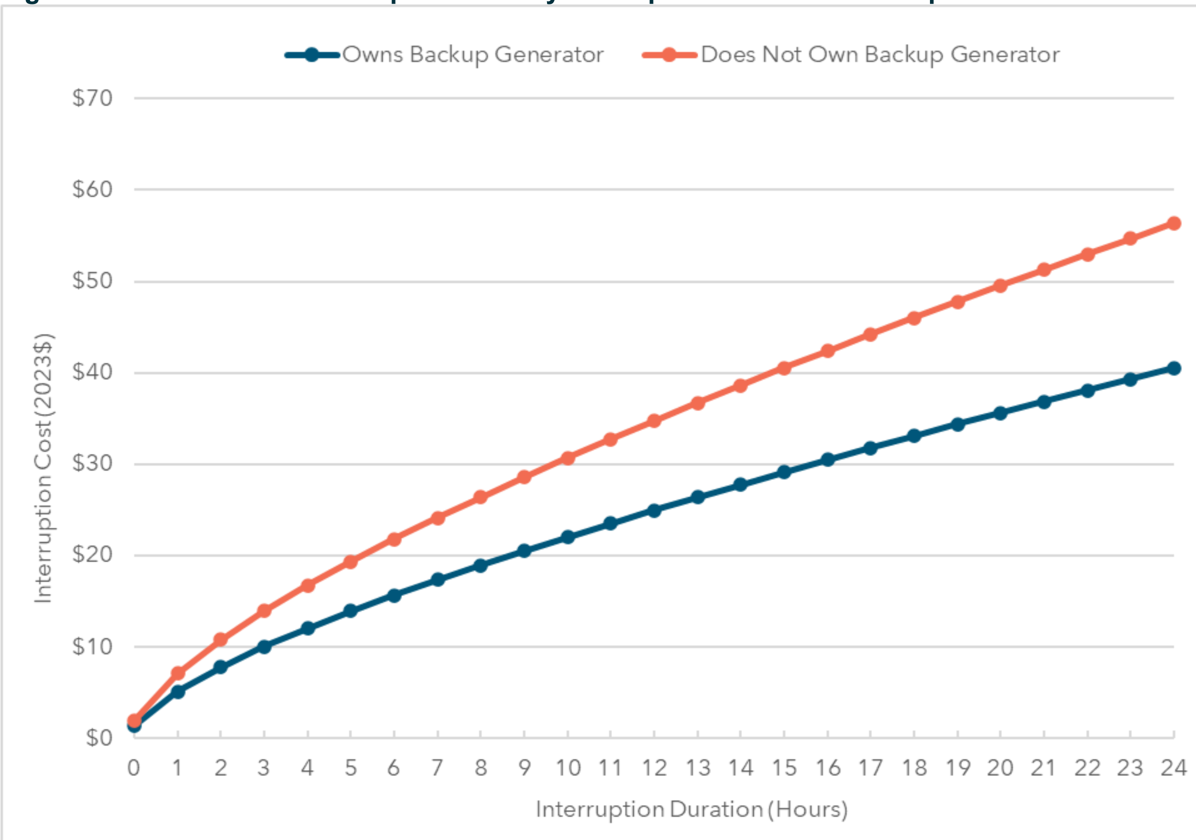

Figure 3.11 shows how interruption costs vary depending on whether a member of a residential household works from home. Costs are higher when at least one household member works from home. For example, the cost of a 24-hour interruption for a household in which at least one member works from home is $57.25, while for households in which no member works from home it is $50.39.

Figure 3.11. Residential Interruption Cost by Work from Home Status

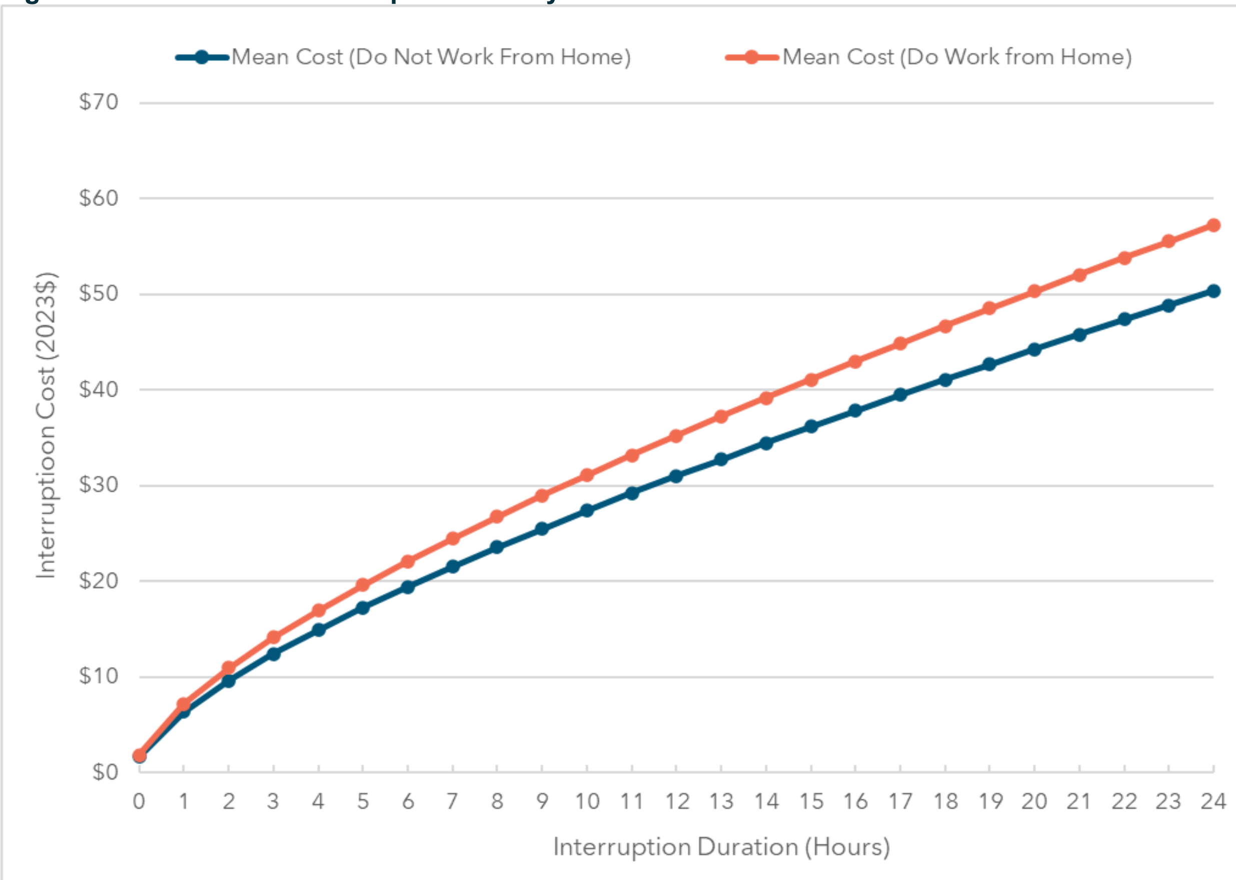

4. Non-residential Customer Damage Function

This section describes the development of and presents selected results from the updated non-residential CDF. Section 4.1 describes the overall process for developing the non-residential CDF, including the reasons for using a two-step estimation approach. Section 4.2 details the factors (i.e., explanatory variables) considered for inclusion in the non-residential CDF and the methods used to develop and test candidate CDFs. Section 4.3 presents the final set of factors selected for the non-residential CDF, along with a variety of results from its application.

4.1 Overview of the Development Process

The non-residential CDF was developed from the survey responses in three steps.[21] The first step involved specifying a continuous form anchored by the point estimates for a momentary, 2-hour, 8-hour, and 24-hour interruption to produce interruption costs for any interruption duration lasting up to 24 hours. The second step was to select a functional form for the CDF. The third step was to choose among the available explanatory factors to develop a final specification for the CDF.[22]

The first step follows the same methods used to develop the residential CDF described in Section 3.1, with additional details provided in Appendix A. The second step is described in this subsection, and the third and final step is described in Section 4.2.

The non-residential CDF was developed using a two-step estimation approach because the distribution of interruption costs from the survey responses is highly skewed. While the majority of respondents reported interruption costs that varied over a wide range, a significant number of respondents reported interruption cost of $0. This was not an unexpected finding. Depending on the business type, non-residential customers can incur a wide range of costs due to power interruptions. For example, as a result of a momentary interruption, a school may have zero costs, while a large manufacturing facility may have large costs associated with delays while restarting equipment and processes.

Table 4.1 presents the percentages of non-residential customers reporting either zero or positive interruption costs for each interruption duration. Across all durations, approximately 24% of respondents reported zero costs.[23] As expected, this percentage decreases as interruption duration increases.

---

[21] The first step in the development of the residential CDF (i.e., translation of OHDC responses into a single interruption cost for each duration) was not required for the development of the non-residential CDF because the responses to the non-residential survey are point estimates of interruption costs.
[22] Note that the original ICE Calculator relied on two separate sets of CDFs – one set for small non-residential customers, and a second set for medium and large non-residential customers. As a result of testing that is documented in Appendix I, which confirmed its appropriateness, only a single set of CDFs was developed for the non-residential sector.
[23] This value is generally consistent with the survey responses used to develop ICE 1.0, which had 33% zero values.

Table 4.1. Percentage of Non-residential Customers with Zero Interruption Costs

Cost Type	Interruption Duration				Total
	Momentary	2 Hours	8 Hours	24 Hours	
Zero Cost	49%	21%	15%	8%	**24%**
Non-zero Cost	51%	79%	85%	92%	**76%**

Figure 4.1 through Figure 4.4 show the distribution of reported non-zero interruption costs for each interruption duration. For ease of presentation, the figures do not show interruption costs greater than $10,000. As expected, the distribution becomes flatter (i.e., higher costs are reported more frequently) as interruptions increase in duration. In addition to the large concentration of reported costs at $0, there is a noticeable and consistent spike in reported costs at $5,000, which also increases with interruption duration. We speculate that this is due to respondents who perceive $5,000 as both a simplifying round number and as a reasonable approximation of their costs.

Figure 4.1. Distribution of Non-residential Interruption Costs (Momentary Interruption, 2023$)

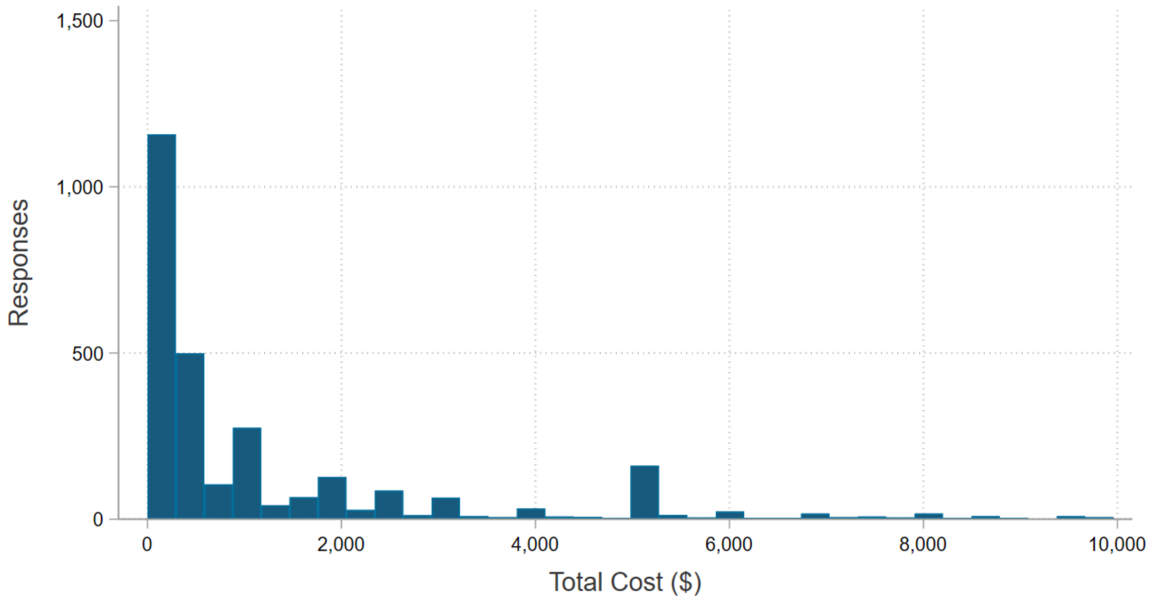

Figure 4.2. Distribution of Non-residential Interruption Costs (2-hour Interruption, 2023$)

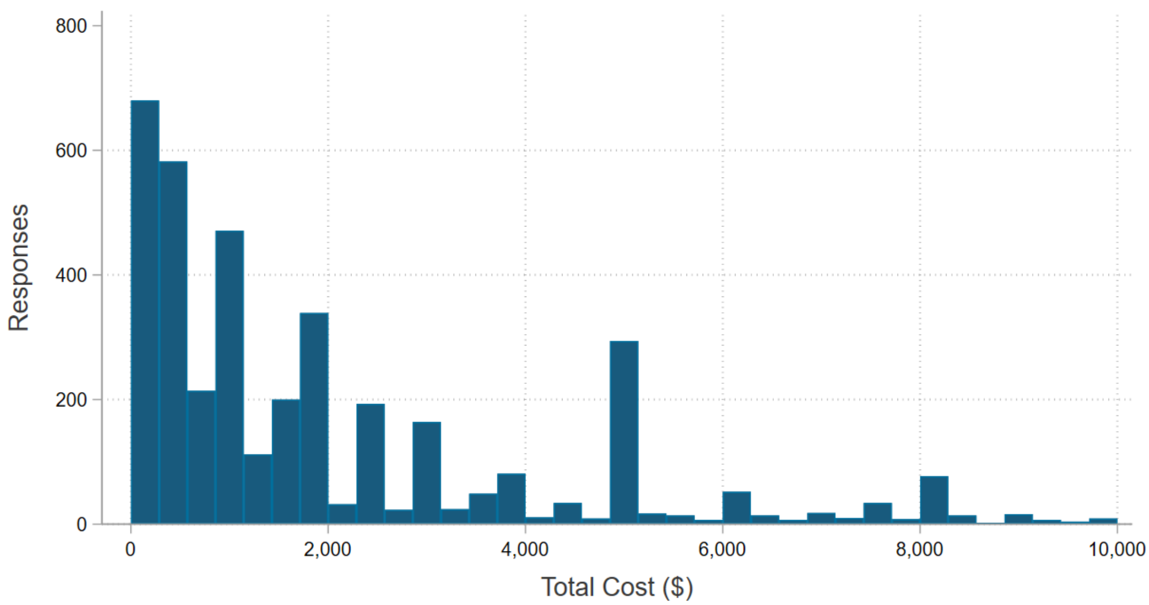

Figure 4.3. Distribution of Non-residential Interruption Costs (8-hour Interruption, 2023$)

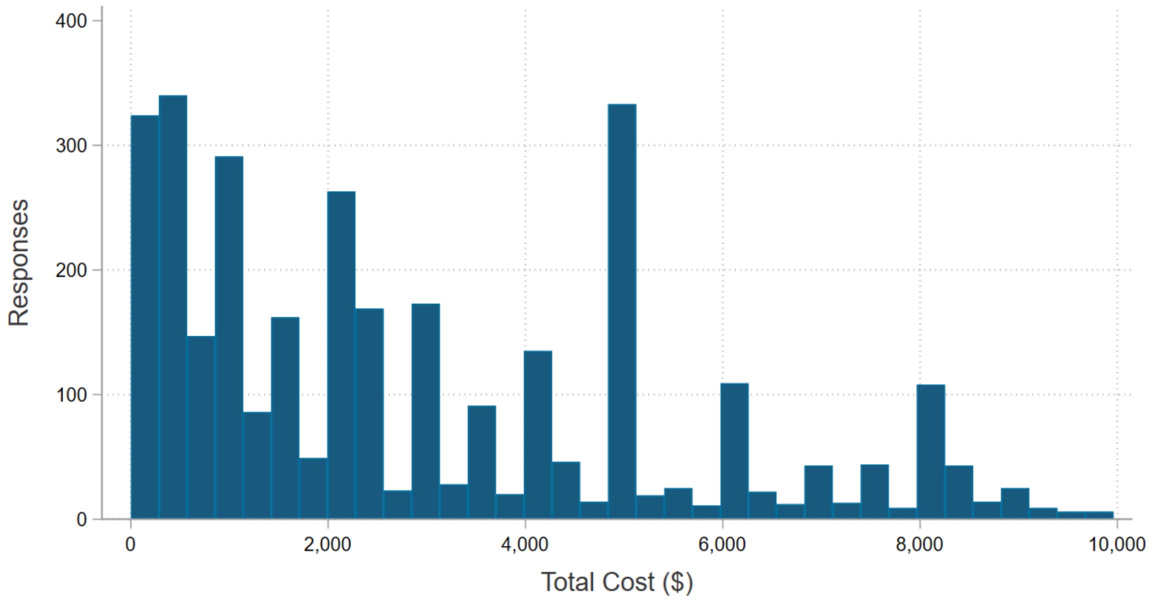

Figure 4.4. Distribution of Non-residential Interruption Costs (24-hour Interruption, 2023$)

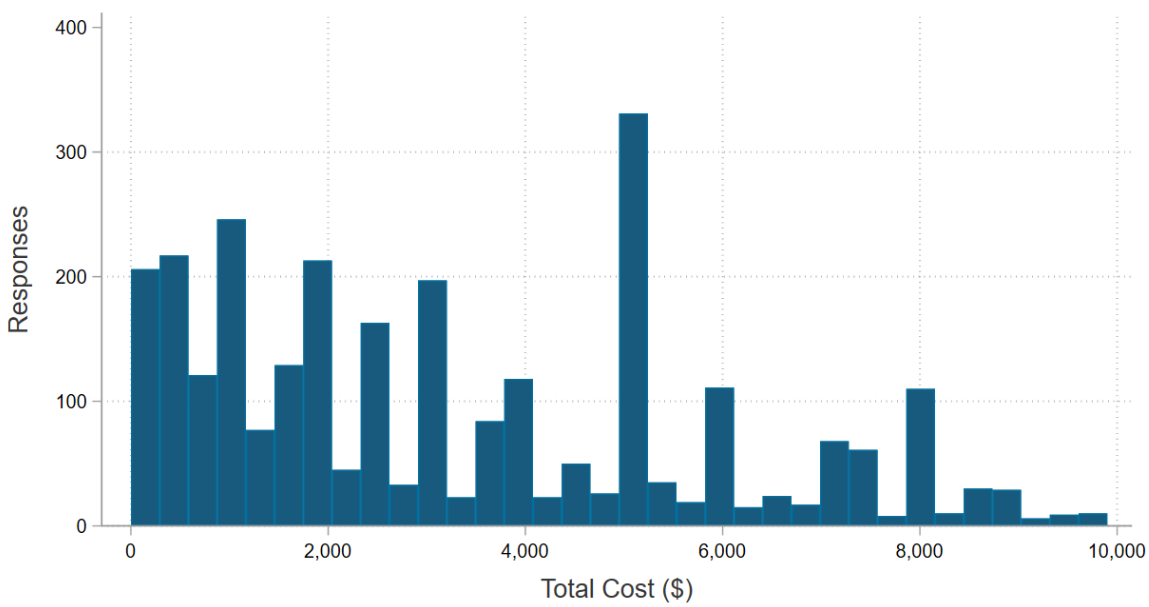

Two-part estimation approaches are commonly used to analyze highly skewed data involving a mixture of zero values and wide range of continuously distributed non-zero values (Farewell et al., 2017). For example, two-part approaches are often used in other applications with skewed data (e.g., healthcare costs) because they have been shown to produce accurate predictions while accounting for large concentrations of zero values (Deb and Norton, 2018).

A two-part estimation approach involves the specification of two independent composite models. The first model assesses the probability that a customer will report a cost of zero versus a non-zero (and positive) cost. This model is based on a set of explanatory variables that describe the nature of the interruption (e.g., duration) as well as customer characteristics. The second model is used to estimate interruption costs for only those customers who reported a non-zero interruption cost. As with the first model, the second model is also based on explanatory variables that describe the nature of the interruption (e.g., duration) as well as customer characteristics. In some instances, the explanatory variables will be common to both models, such as interruption duration, but they can and will often differ. When implemented in the ICE Calculator, the predicted probabilities from the first model are multiplied by the estimated interruption costs from the second model to produce the estimate of interruption cost.

The first model estimates the probability that a customer will experience (i.e., will have reported) a non-zero interruption cost. Either a probit or a logit model can be used to estimate a binary outcome. Both models generally produce similar results (Vasisht, 2007). For consistency with the development of the original ICE Calculator (Sullivan et al., 2009; Sullivan et al., 2015), a probit model was used for the ICE 2.0 modeling.

The second model estimates the interruption costs for customers who have reported a non-zero cost. These costs follow a continuous distribution, so a generalized linear model (GLM) was used to model them. A GLM model is defined by three main components: the distribution family of the response variable (reported interruption costs), a linear regression function, and a link function that relates the output of the linear regression to the mean of the interruption cost distribution (Khuri et al., 2006). A distribution family is a set of statistical models that best describe the distribution of the reported interruption costs. The linear regression function comprises the variables that affect the interruption cost. The link function considers how the model variables need to be transformed in order to describe the interruption costs.

The development of linear regression function is decribed in Section 4.2. The selection process for the distribution family and link function is described in the remainder of this subsection.

The distribution family of the response variable was identified by using the Modified Park Test (Park, 1966; Manning and Mullahy, 2001). The Modified Park Test finds the most appropriate distribution based on how the mean is related to the variance in the context of different power functions. The test found that the non-residential interruption cost variance is not statistically different from the mean squared. This suggests that the data follows a gamma distribution as this family of distributions has a variance proportional to their mean squared. The choice of a gamma distribution is also consistent with that used to develop the original ICE Calculator.

The link function was chosen based on the parameter for the family of transforms proposed by Box and Cox (Box and Cox, 1964). The Box-Cox transformation technique uses a single parameter to make the values in the dataset more closely follow a normal distribution. Optimizing the associated parameter can suggest what transformation is most appropriate for a dataset. For example, if the parameter is equal to 1 no transformation is needed; a parameter value of 0.5 suggests a square root transformation; and a parameter value of 0 suggests a natural log transformation. The maximum likelihood estimate for the transform parameter was calculated to be 0.009. Following the suggestion outlined by Sheather, the estimate was rounded to the closest interpretable value (Sheather, 2009). For the non-residential dataset, this yields a transform parameter of zero, which suggests a natural log link function. This is consistent with choices made in developing the original ICE Calculator.

The two-part model is described in Equation 4.1. Note, X and Y represent the covariate vector for the probit and GLM models respectively. Furthermore, γ and β represent the respective coefficient vectors. Finally, let ϕ represent the cumulative distribution function of the standard normal function.

Equation 4.1. General Non-residential Regression Specification for ICE Calculator Model

$$Pr(Cost > 0|X) \ = \ \phi(X^T\gamma)$$

$$E(Cost|Cost > 0) \ = exp(Y^T\beta)$$

$$C(X,Y) = \ \phi(X^T\gamma) \times exp(Y^T\beta)$$

4.2 Selection of Explanatory Variables

Variables considered for inclusion in the non-residential CDFs were selected from a series of interruption- and firm-specific characteristics (Table 4.2).

Table 4.2. Non-residential Potential Model Variables

Continuous Variables	
• Interruption duration (in minutes) • Annual electricity usage (in kWh) • GDP per kWh (collected at the state level)	
Categorical Variables	
Interruption Onset Time • Morning • Midday • Evening **Season** • Summer • Winter **Day of Week** • Weekday • Weekend **Advance Warning** • Yes • No **Previous Interruption in Last 12 Months** • Yes • No	**Ownership of Backup Generation** • Yes • No **Industry** • Accommodation and Food Services • Administrative and Support and Waste Management and Remediation Service • Agriculture, Forestry, Fishing and Hunting • Arts, Entertainment, and Recreation • Construction • Educational Services • Finance and Insurance • Health Care and Social Assistance • Information (e.g., Data Centers) • Management of Companies and Enterprises • Manufacturing • Mining, Quarrying, and Oil and Gas Extraction • Other Services • Professional, Scientific, and Technical Services • Public Administration • Real Estate and Rental and Leasing • Retail Trade • Transportation and Warehousing • Utilities • Wholesale Trade

The variable selection for the non-residential CDF was implemented following the same procedures used to develop the residential CDF. First, LASSO regression was used to initially trim the list of potential variables.[24] Next, the candidate models selected using LASSO were then individually tested using 10-fold cross-validation, which is a technique often used to gauge model prediction accuracy. [25] Finally, the cross-validation results were reviewed in light of other design considerations in order to select a final model.

[24] See Appendix G for more details on the LASSO regression.
[25] The amount of individual-level non-residential data available allowed for only 10 folds to be used. The residential data required the individual-level responses to be collapsed into a sample average resulting in a higher number of folds (100) being needed.

In selecting the final non-residential CDF, we sought to avoid making the model unnecessarily complicated, especially given that complexity yields diminishing returns in predictive accuracy. We prioritized parsimony, aiming to retain only those variables that offered meaningful improvements in model performance while keeping the tool accessible for end users.

At the start of the modeling process, the non-residential data included over 30 potential covariates. An ideal model would only use covariates that significantly impact the response variable. However, exploring every possible model permutation to find the optimal set of covariates would be computationally burdensome. Therefore, as with the development of the residential CDF, LASSO regression was used to select variables for inclusion in candidate CDF models.

LASSO regression was used separately to select variables for both the probit and GLM models estimated for the non-residential sector. As a result, the selected variables differed for each set of candidate models. The overall process is outlined in Figure 4.5. For the probit model, we used a non-zero cost flag to delineate which customers experienced a positive cost. Using the non-zero cost flag as a response variable, a LASSO regression was performed over a range of n values for the penalty term. Each value, λ_i produces a set of covariates, thus the LASSO regression yields n potential models. It is worth noting that decreasing λ results in additional variables being selected by the LASSO. Over the range considered, the set of covariates that correspond to a given λ will always be a subset of the covariates selected by smaller values of λ.

Figure 4.5. Non-residential Variable Selection Process Overview

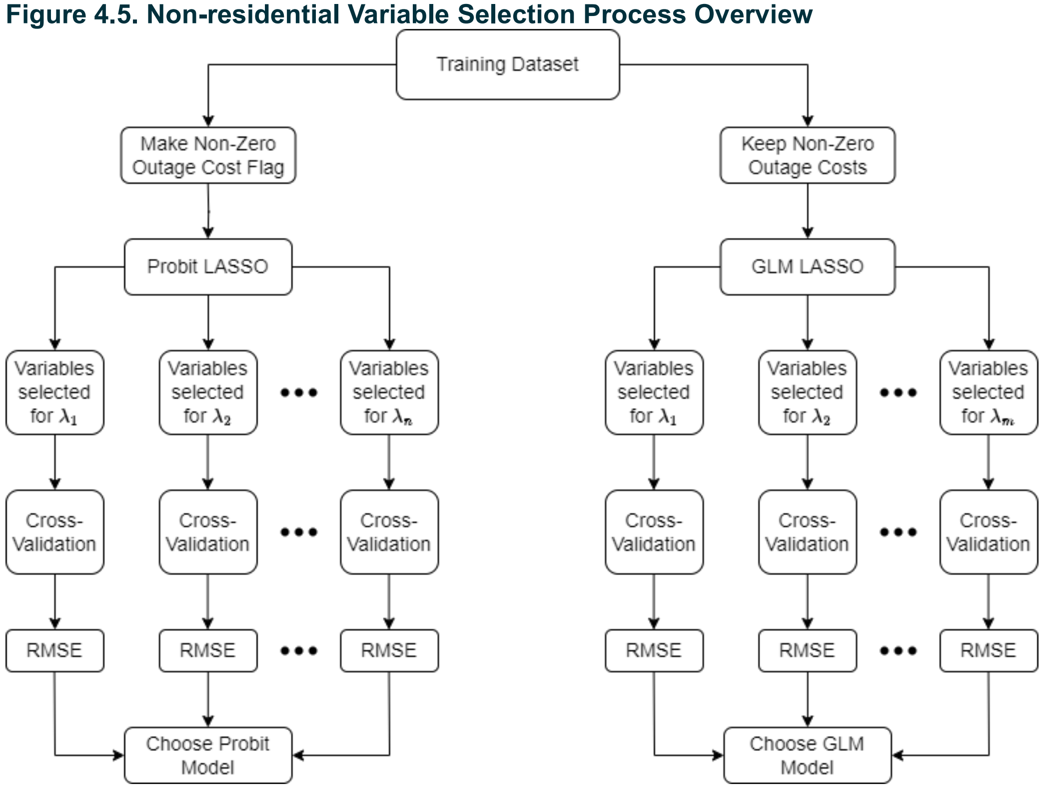

Following the selection of candidate models, we used a 10-fold cross-validation process to calculate each model's respective RMSE. This process is shown in Figure 4.6. Cross-validation involves the training dataset being split into 10 distinct subsets (folds). A single fold is then reserved as a testing dataset while a model is trained with the other nine folds. To avoid overfitting when assessing accuracy, the trained model makes predictions on the testing dataset, which was not included in the training. The RMSE is calculated based on the differences between the model predictions and the testing dataset. This process is repeated until all 10 folds have served as the training dataset (nine times) and the testing dataset (one time). The overall model RMSE is calculated by finding the mean in the testing dataset across the individual RMSEs.

Figure 4.6. Non-residential Cross-validation Process

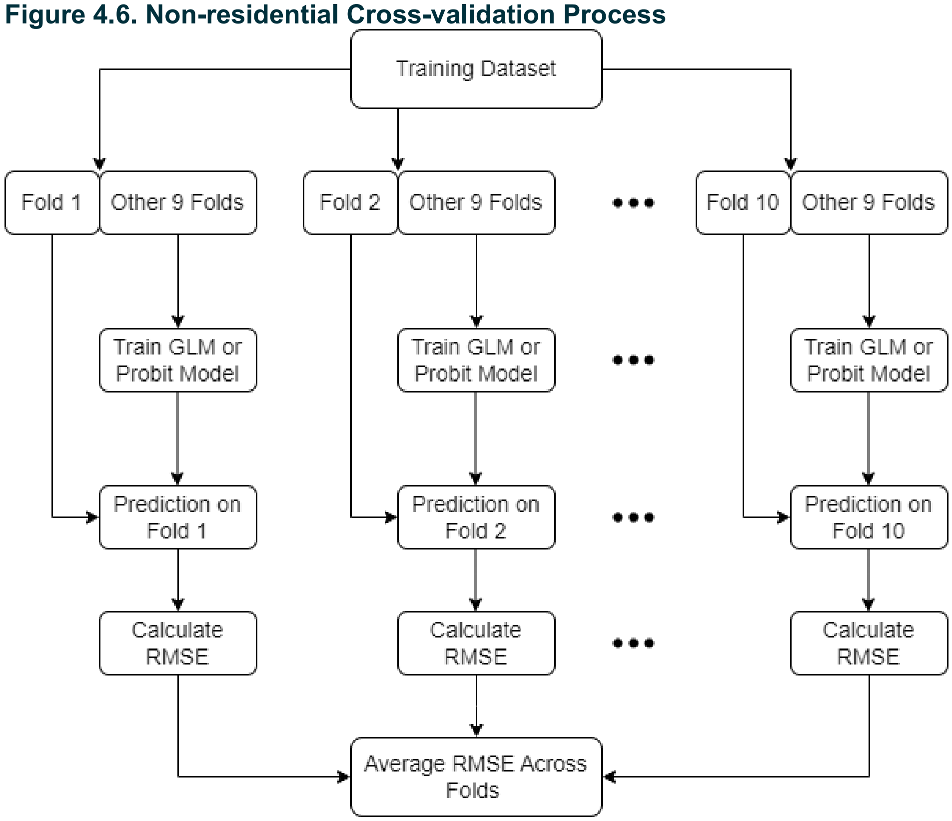

The final probit model was selected by finding the point at which adding new variables offered diminishing reductions to the RMSE.

The results of the probit variable selection process are presented in Table 4.3. The percent change in RMSE measures the effect of adding a new variable to the previous model iteration. For the first few models, adding covariates initially reduces the RMSE by more than one percentage point. However, the error reduction quickly drops to around 0.2%. The final model (Model 3) was chosen along this boundary because it yields the simplest model that produces accurate results. More specifically, the variables added to Models 4 and 5 reduce error by less than a percent and were therefore determined to be not worth the added complexity in the model.

Table 4.3. Non-residential Probit Model Variable Selection

Model #	Lambda	Model	RMSE	% Change in RMSE
1	exp(-2.5)	duration* annual kWh	0.392	NA
2	exp(-2.6)	duration* annual kWh, annual kWh	0.386	-1.55%
3	exp(-2.7)	duration* annual kWh, annual kWh, weekday, warning	0.381	-1.22%
4	exp(-2.8)	duration* annual kWh, annual kWh, weekday, warning, afternoon	0.381	-0.18%
5	exp(-3.0)	duration* annual kWh, annual kWh, weekday, warning, afternoon, season	0.380	-0.09%

In the table, duration annual kWh and annual kWh are simplified from log(annual kWh)*log(duration) and log(annual kWh) for readability.

The GLM variable selection process is similar to the probit process, except non-zero interruption costs are used as the response variable. Table 4.4 highlights the GLM variable selection results. Like the probit model, adding new covariates offers diminishing improvements to the RMSE. Based on the change in RMSE alone, Model 2 seemed like a reasonable final model choice. However, after comparing the function forms of Models 2 and 3, we observed that models with a log duration cubed term perform better across interruption durations from 0 to 24 hours than the models that only include the log duration term. Out of the three models that include the cubic duration, only Model 4 had a reduction in error at the magnitude of a tenth of a percent. Thus, Model 4 was selected as it is the simplest model that offers a relatively high reduction in error across all models that included the cubic term.

Model 4 includes manufacturing and healthcare terms, suggesting these are the most meaningful variables related to industry classification. Additional specific industries could be included in the model, such as retail trade. However, since there are 20 different industry categories, including all 20 components would dramatically increase model complexity. The final model balances simplicity and the user experience over a complex model by including only the two most influential industries: manufacturing and healthcare. As a result, the model describes all other industries as the reference case, which can be adjusted by the presence of manufacturing and healthcare indicator variables. In other words, if the manufacturing and healthcare indicators are zero, then the model describes cost for all other industries.

Table 4.4. Non-residential GLM Variable Selection

Model #	Lambda	Model	RMSE	% Change in RMSE
1	exp(-1.0)	duration*annual kWh, annual kWh	349,796	NA
2	exp(-1.3)	duration*annual kWh, annual kWh, manufacturing	346,254	-1.01%
3	exp(-1.6)	duration*annual kWh, annual kWh, manufacturing, healthcare, duration^3	346,210	-0.01%
4	exp(-1.9)	duration* annual kWh, annual kWh, manufacturing, healthcare, duration^3, warning	345,515	-0.20%
5	exp(-2.3)	duration* annual kWh, annual kWh, manufacturing, healthcare, duration^3, warning, construction	345, 254	-0.08%

*In the table, duration*annual kWh and annual kWh are simplified from log(annual kWh)*log(duration) and log(annual kWh) for readability

4.3 Final Specification of the Non-residential CDF and Selected Results

Table 4.5 presents the final factors (or explanatory variables) selected for the non-residential segment. Both continuous variables, such as duration (in minutes) and annual usage (in kWh), and categorical variables (e.g., advance warning of the interruption, day of week, and industry) were selected.

Table 4.5. Final Non-residential Model Specification

Continuous Variables
• Interruption duration (in minutes) • Annual electricity usage (in kWh)
Categorical Variables
Advance Warning • Yes • No **Day of Week (Probit Only)** • Weekday • Weekend **Industry (GLM Model Only)**[26] • Health Care and Social Assistance • Manufacturing

*The functional form for duration and usage in the final probit model are log(annual kWh)xlog(duration) and log(annual kWh). The functional form of the final GLM is log(annual kWh)xlog(duration), log(annual kWh), log(duration)^3.

Figure 4.7 shows the non-residential interruption costs estimated by the non-residential CDF for all durations from a momentary interruption (lasting up to 5 minutes) to 24 hours. These results were produced by using the average characteristics from the full set of non-residential

[26] Customers in all other industries are also included in the model, but are not shown in the table as a separate variable.

customer responses described in Section 2. As a result, these results can be compared directly to the interruption costs derived from the survey responses (see Table 2.9).

Figure 4.7. Non-residential CDF-based vs. Survey-based Interruption Costs

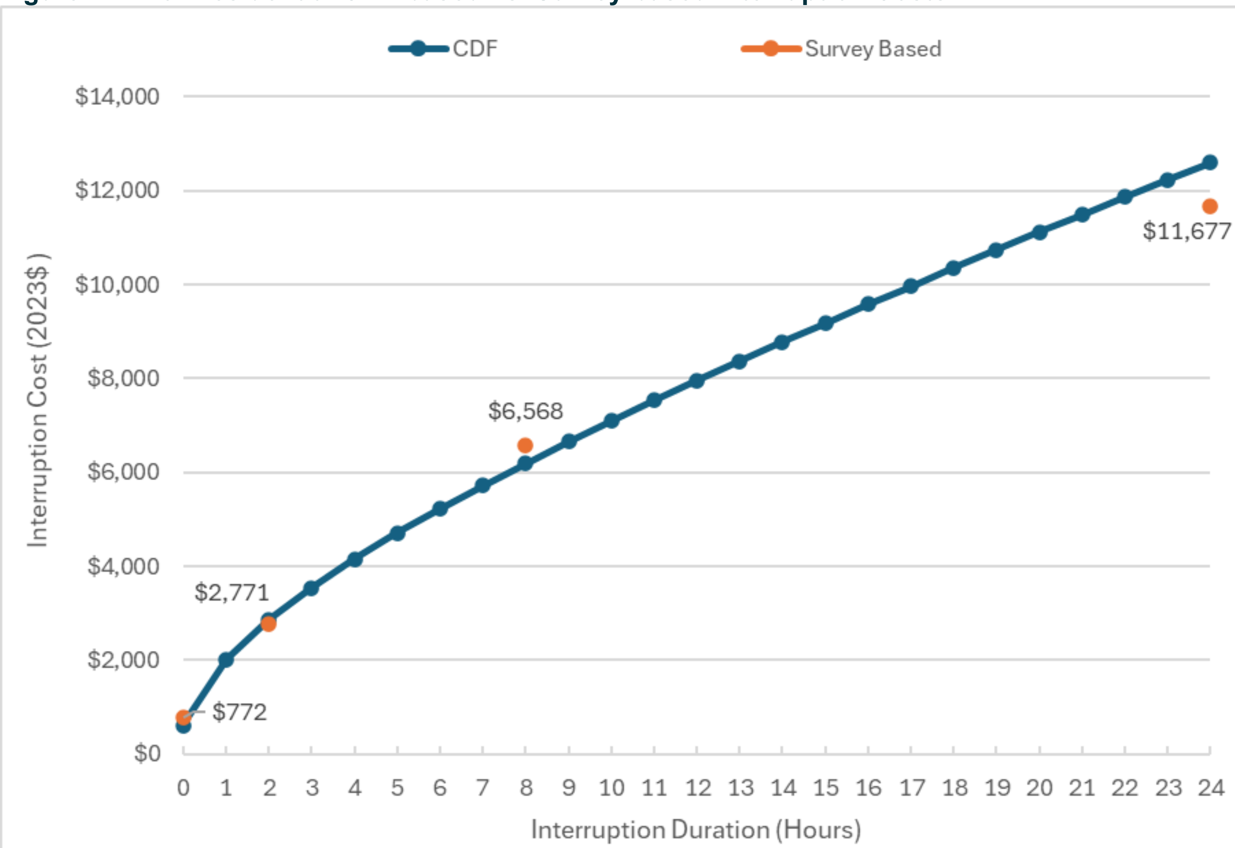

Table 4.6 presents interruption costs estimated using the non-residential CDF for the interruption durations used in the surveys. In addition, summary indices based on the interruption cost are also presented.

Table 4.6. Non-residential Interruption Costs (2023$)

Duration of Power Interruption Event	Cost per Event	Cost per kW	Cost per Unserved kWh	Cost per CMI
Momentary	$609	$43	$521	$122
2 Hours	$2,839	$202	$101	$24
8 Hours	$6,172	$440	$55	$13
24 Hours	$12,646	$902	$38	$9

*Cost per kW (and kWh) were calculated by dividing the average interruption cost by average kW (and kWh).

Figures 4.8 through Figure 4.11 show how non-residential interruption costs vary for selected explanatory variables. The figures were developed by holding all explanatory variables at their population-average values.

Figure 4.8 shows how interruption costs vary widely by annual energy usage. Generally, non-residential customers with higher usage experience higher interruption costs. For example, for a 24-hour interruption, a customer in the 5th percentile of annual usage (3,518 kWh per year) has an estimated interruption cost of $1,552, while a customer in the 95th percentile of annual usage (9,273,992 kWh per year) has an estimated interruption cost of $187,690.

Figure 4.8. Non-residential Interruption Cost by Usage

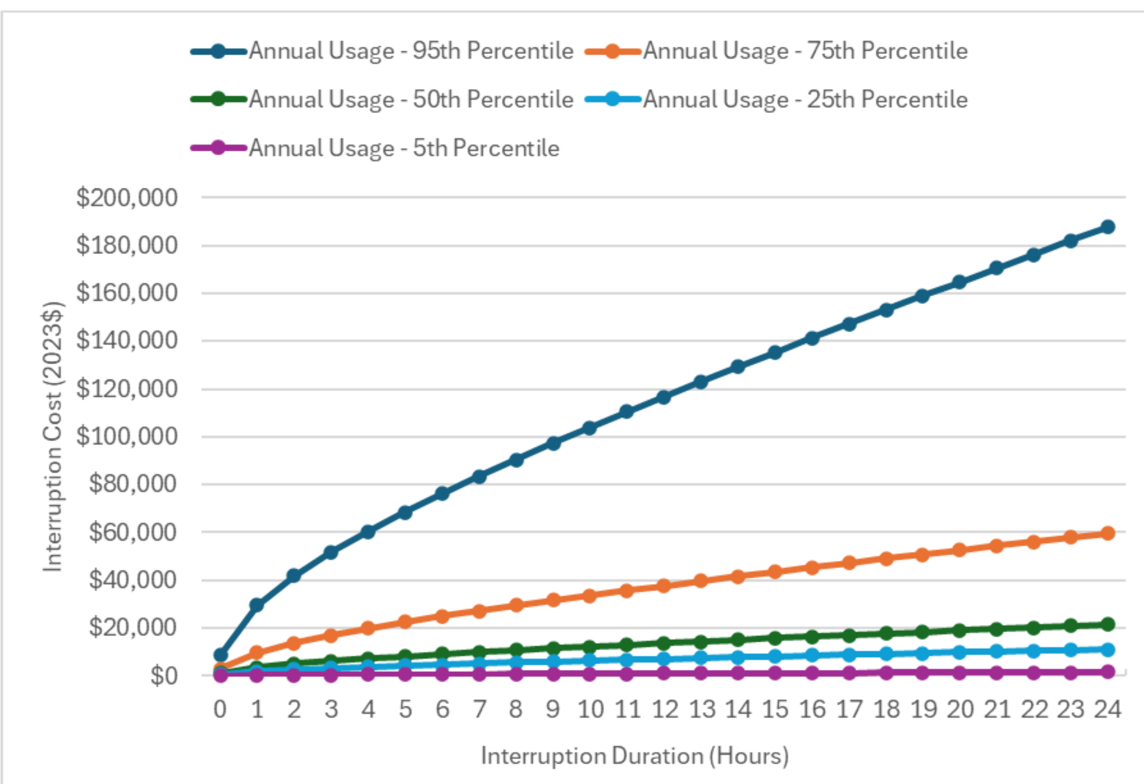

Figure 4.9 shows how interruption costs vary by the day of week. Overall, interruption costs are estimated to be higher for interruptions that take place during weekdays. This is likely due to the fact that some organizations do not operate on weekends.

Figure 4.9. Non-residential Interruption Cost by Day of the Week

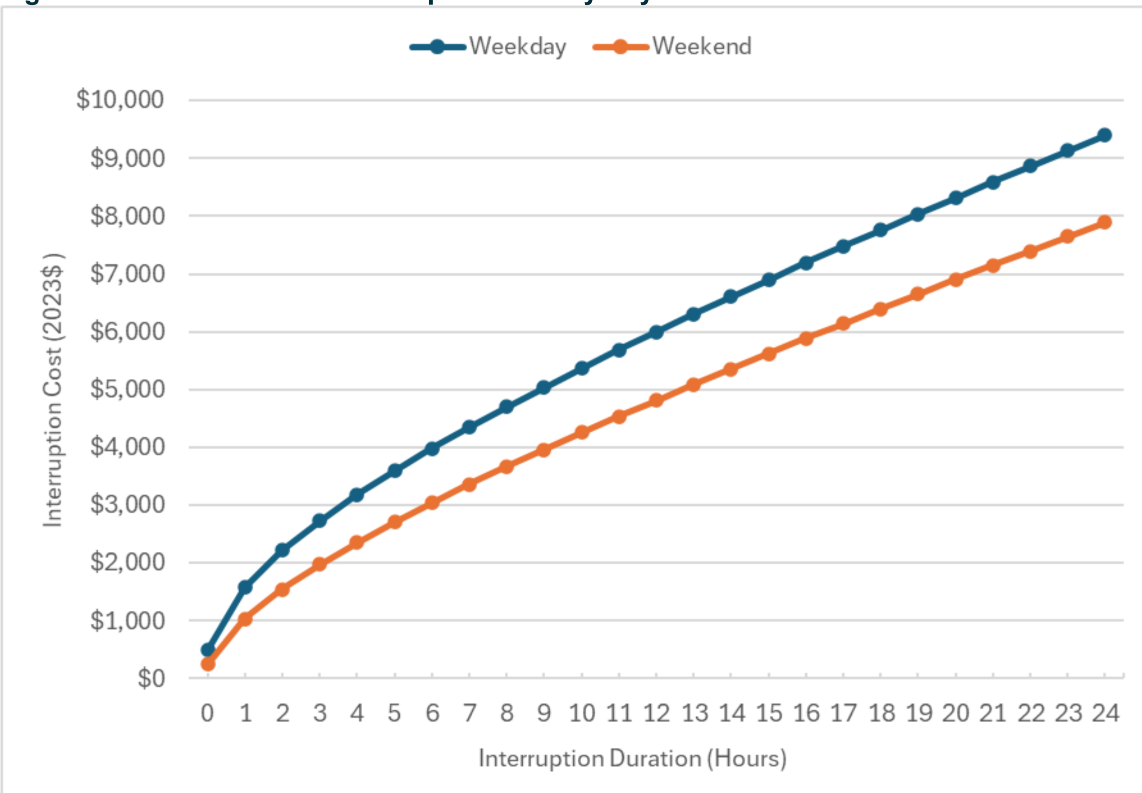

Figure 4.10 shows how interruptions costs vary for different types of firms. Customers in the manufacturing and healthcare industry are estimated to incur higher interruption costs compared to customers in other industries.

Figure 4.10. Non-residential Interruption Cost by Industry

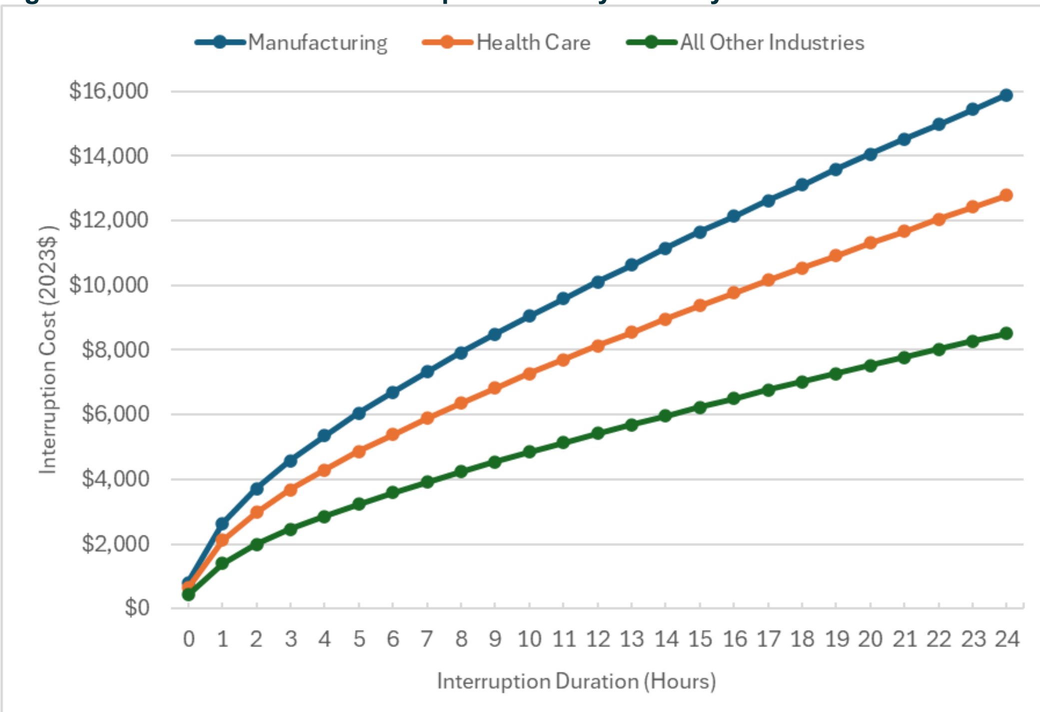

Figure 4.11 shows how interruption costs vary depending on whether non-residential customers are provided with an advance warning that an interruption will occur. Estimated interruption costs are lower when advance warnings are provided. The percentage reductions increase as the duration of an interruption increases.

Figure 4.11. Non-residential Interruption Cost With and Without Advance Warning

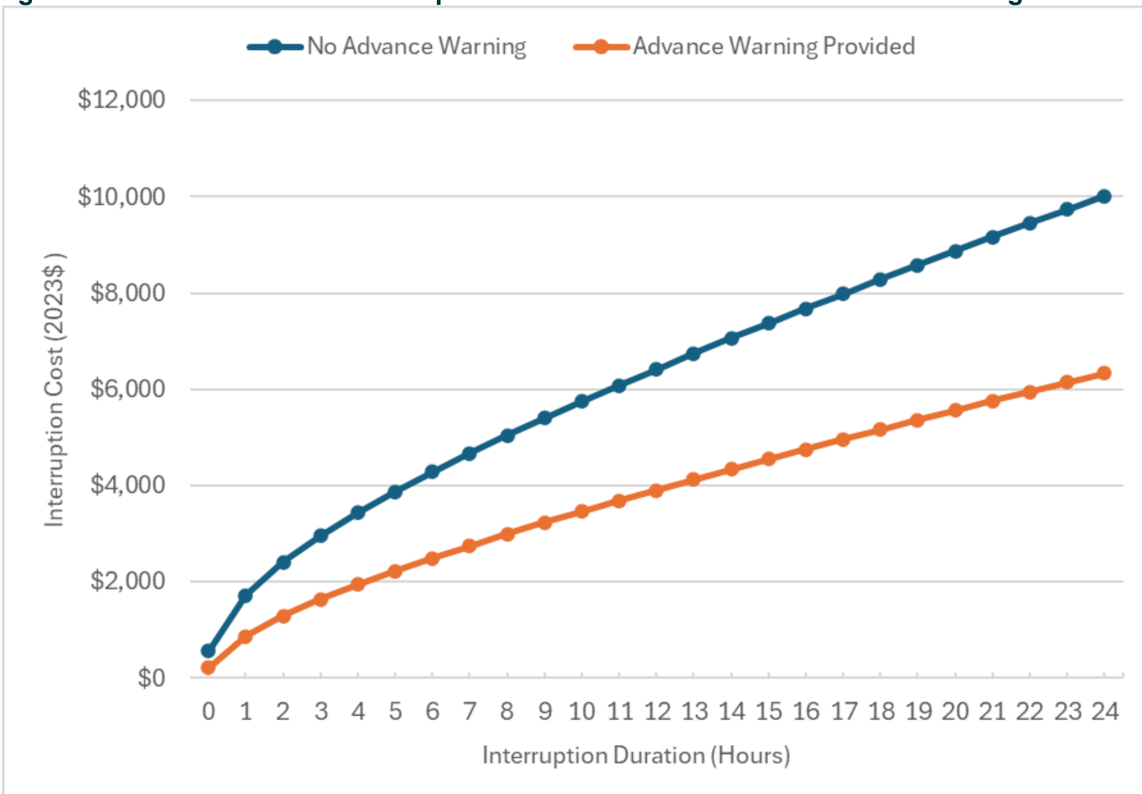

5. Conclusion, Next Steps, and Caveats

The ICE 2.0 Initiative is a national, multi-client study to update the underlying data and enhance the functionality of the ICE Calculator. The Initiative involves Berkeley Lab contracting with sponsoring utilities to administer identical, modernized interruption cost surveys to statistically representative samples of each sponsoring utility's customers. Berkeley Lab and Resource Innovations then pool the survey results across the utilities and use them to update the CDFs that drive the ICE Calculator.

This report documented the activities and findings from Phase 1 of the Initiative. Phase 1 was sponsored by eight utilities: American Electric Power, Commonwealth Edison, Dominion Energy, Duke Energy, DTE Electric, Exelon, National Grid, and Puget Sound Energy. In this phase, 11 customer interruption cost surveys were administered representing a total of 24 electricity distribution service territories, 23 of them located in the Eastern and Midwestern regions of the U.S. and one located in the Pacific Northwest.

The CDFs described in this report were used to update the ICE Calculator in early 2025. In addition, the functionalities of the online ICE Calculator have been redesigned and enhanced in response to user feedback.

At the time this document was prepared, surveys of Phase 2 sponsoring utilities were in various stages of completion, and planning was in progress for the utilities that will be included in Phase 3.[27] As a result, when these surveys are combined with the Phase 1 surveys analyzed in this report, we will be able to explore the effects of regional variations in power interruption costs. Finally, discussions are currently underway to survey customers in other regions of the U.S. that are not well-represented.

The ICE Calculator will be updated following the completion and integration of each subsequent phase of the ICE 2.0 Initiative. In particular, we will revisit the development of the CDFs in order to consider additional variables that are not currently included. As the overall sample size and geographic representation improve in Phase 2, variables that previously did not influence model accuracy may prove to be important. This will ensure that the final ICE 2.0 models are the most accurate and reliable models of interruption costs available, based on reassessing all data and variables after each phase.

It is important to bear in mind that the ICE Calculator focuses on estimating costs associated with localized, short-duration power interruptions. The surveys that form the basis for the ICE Calculator were not designed to support the estimation of costs associated with power interruptions of over 24 hours. Moreover, they did not consider the size or geographic scope of power interruptions (i.e., they did not consider whether the total number and type of customers affected by a power interruption might affect an individual customer's costs).

[27] Phase 2 of this Initiative involves surveys of Pacific Gas & Electric, San Diego Gas & Electric, Southern California Edison, and Ameren Missouri customers. We are now recruiting utilities for Phase 3.

Our focus on estimating the costs of localized, shorter-duration power interruptions stems from two practical considerations. First and foremost, the vast majority of power interruptions are local in scope and short in duration. Accordingly, reducing the frequency, scope, and duration of these power interruptions is a principal focus of utility reliability planning activities. The ICE Calculator is specifically designed to estimate the benefits of reliability-enhancing activities that reduce costs of these power interruptions.

Second, survey methods are best suited to collecting information about costs that customers experience, themselves. They are not well-suited to collecting information about costs that power interruptions might create for other customers that have power, but whose activities are affected by those without power. For some power interruptions, these "costs" may even be "benefits." For example, when a retail customer loses business because they have experienced a power interruption, a competitor who has not lost power may gain the business that has been lost by this customer.

While this example does not change the fact that the customer without power has been impacted, it can temper or place a larger context around which power interruption costs should be considered in value-based reliability planning activities. In particular, while we believe that these sorts of offsetting effects are limited to only certain sectors (e.g., retail, dining, and lodging) for short-duration, localized power interruptions, they are likely to be of greater significance for widespread, long-duration power interruptions.

Estimating the power interruption costs associated with widespread, long-duration power interruptions requires methods that extend beyond the collection of survey information. First, the respondent's ability to estimate the costs they would experience via a survey may be unreliable because they may have never experienced such events, and because generalizing from details of the scenarios required to describe them meaningfully (e.g., the impacts on other critical infrastructures, such as road conditions, communications, etc.) is not straightforward. Second, the impacts of these power interruptions on customers who have not lost power are likely to be much more significant and so must also be considered.

Methods to estimate the cost of widespread, long-duration power interruptions are an active focus of current research.[28] It is our hope that these methods will soon contribute to a more comprehensive understanding of the cost of power interruptions: whether localized and short-duration, or widespread and long-duration.

While this research progresses, the ICE Calculator remains the most comprehensive, up-to-date, fully vetted, and publicly available tool available to estimate the cost of power interruptions.

[28] See Larsen et al., 2024, Larsen et al., 2019, and Sanstad, 2016 for reviews of approaches and a recent example.

References

Box, G., & Cox, D. (1964). An Analysis of Transformations. *Journal of the Royal Statistical Society. Series B (Methodological), 26*(2), 211-252.

Su, W., Bogdan, M., & Cand`es, E. (2017). False Discoveries Occur Early on the Lasso Path. *The Annals of Statistics, 45*(5), 2133-2150.

Belloni, A., & Chernozhukov, V. (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli, 19*(2), 521-547.

Cooper, J. C., Hanemann, M., & Signorello, G. (2002). One-and-one-half-bound dichotomous choice contingent valuation. *Review of Economics and Statistics, 84*(4), 742-750.

Deb, P., & Norton, E. C. (2018). Modeling health care expenditures and use. *Annual review of public health, 39*(1), 489-505.

Desboulets, L. (2018). A review on variable selection in regression analysis. *Econometrics, 6*(4), 45.

Doornik, J. A. (2009). Econometric Model Selection with More Variables Than Observations. *Oxford: Economics Department, University of Oxford. Unpublished Work.*

Farewell, V., Long, D., Tom, B., Yiu, S., & Su, L. (2017). Two-part and related regression models for longitudinal data. *Annual review of statistics and its application, 4*(1), 283-315.

Freijeiro-González, L., Febrero-Bande, M., & González-Manteiga, W. (2022). A critical review of LASSO and its derivatives for variable selection under dependence among covariates. *International Statistical Review, 90*(1), 118-145.

Hanemann, M., Loomis, J., & Kanninen, B. (1991). Statistical efficiency of double-bounded dichotomous choice contingent valuation. *American journal of agricultural economics, 73*(4), 1255-1263.

Homrighausen, D., & McDonald, D. (2018). A study on tuning parameter selection for the high-dimensional. *Journal of Statistical Computation and Simulation, 88*(15), 2865-2892.

Homrighausen, D., & McDonald, D. J. (2017). Risk consistency of cross-validation with lasso-type procedures. *Statistica Sinica, 27*(3), 1017-1036.

Hurvich, C. M., & Chih-Ling Tsai. (1990). The impact of model selection on inference in linear regression. *The American Statistician, 44*, 214–17.

Kanninen, B. J. (1993). Optimal experimental design for double-bounded dichotomous choice contingent valuation. *Land Economics*, 138-146.

Khuri, A. I., Mukherjee, B., Sinha, B. K., & Ghosh, M. (2006). Design issues for generalized linear models: A review. *Statistical Science, 21*(3), 376-399.

Larsen, P. H., J. P. Carvallo, A. H. Sanstad, S. Baik, I. Sue-Wing, D. Wei, A. Rose, J. Smith, C. Ramee, and R. Peterson. (2024). Power Outage Economics Tool: A Prototype for the Commonwealth Edison Service Territory. Retrieved from https://eta-publications.lbl.gov/sites/default/files/poet_final_report_6may2024.pdf

Larsen. P. H., A. H. Sanstad, K Hamachi-LaCommare, J. H. Eto. (2019). Frontiers in the Economics of Widespread, Long-Duration Power Interruptions: Proceedings from an Expert Workshop. Ret rieved from: https://eta-publications.lbl.gov/sites/default/files/long_duration_interruptions_workshop_proceedings.pdf

Lawrence Berkeley National Laboratory. (2023). 2023 Progress Report for ICE Calculator 2.0. https://icecalculator.com/recent-updates

Manning, W. G., & Mullahy, J. (2001). Estimating log models: to transform or not to transform? *Journal of health economics, 20*(4), 461-494.

Park, R. E. (1966). Estimation with heteroscedastic error terms. *Econometrica, 34*(4), 888.

Resource Innovations. (2022). *Memorandum.* Icecalculator.

https://icecalculator.com/assets/documents/ICE-2.0-Survey-Design-Memo.pdf
Sanstad, A. H. (2016). Regional Economic Modeling of Electricity Supply Disruptions: A Review and Recommendations for Research. Retrieved from https://eta-publications.lbl.gov/sites/default/files/lbnl-1004426.pdf

Sheather, S. (2009). *A Modern Approach to Regression with R.* Springer Science & Business Media.

Steyerberg, E. W., Marinus J. C. Eijkemans, & Habbema., J. D. (1999). Stepwise selection in small datasets: A simulation study of bias in logistic regression analysis. *Journal of Clinical Epidemiology, 52*, 935–42.

Sullivan, M. J., & Keane, D. M. (1995). *Outage Cost Estimation Guidebook. Palo Alto CA USA: Electric Power Research Institute Report No EPRI-TR-106082.* San Francisco CA USA: Freeman, Sullivan and Co.

Sullivan, M. J., Collins, M. T., Schellenberg, J., & Larsen, P. H. (2018). Estimating power system interruption costs: A guidebook for electric utilities. *Report #: LBNL-2001164*. Retrieved from https://escholarship.org/uc/item/0xn4r200

Sullivan, M. J., Collins, M. T., Schellenberg, J., Larsen, P. H., Hanemann, M., Morgan, G., . . . Vaishnav, P. (2019). Elicitation of Electric Utility Customer Power Interruption Costs: A Roadmap for Conducting a National Survey. Retrieved from https://icecalculator.com/assets/documents/National-Study-Roadmap.pdf

Sullivan, M. J., Mercurio, M., & Schellenberg, J. (2009). Estimated value of service reliability for electric utility customers in the United States. *Lawrence Berkeley National Laboratory Report No. LBNL-2132E.*

Sullivan, M. J., Schellenberg, J., & Blundell, M. (2015). Updated Estimated Value of Service Reliability for Electric Utility Customers in the United States. *Lawrence Berkeley National Laboratory Report No. LBNL-6941E.*

Taylor, J., & Tibshirani, R. J. (2015). Statistical learning and selective inference. *Proceedings of the National Academy of Sciences, 112*(25), 7629-7634.

Vasisht, A. (2007). Logit and Probit Analysis. *New Delhi: Technical report: I.A.S.R.I.*

Watanabe, M. (2010). Nonparametric estimation of mean willingness to pay from discrete response valuation data. *American Journal of Agricultural Economics, 92*(4), 1114-1135.

Whittingham, M. J., Stephens, P. A., Bradbury, R. B., & Freckleton., R. P. (2006). Why do we still use stepwise modelling in ecology and behaviour? *Journal of Animal Ecology, 75*, 1182–89.

Wooldridge, J. M. (1999). Quasi-likelihood methods for count data. *Handbook of applied econometrics volume 2: Microeconomics*, 321--368.

Appendix A: Residential Valuation Methodology

The residential survey employed a willingness-to-pay (WTP) valuation method called the "one-and-one-half bound discrete choice" (OHDC) survey methodology. The OHDC methodology involves presenting residential customers with a price that represents what they may or may not be willing to pay to avoid a power interruption. This Appendix describes how the prices used in the OHDC were developed.

Following the common procedure used in contingent valuation studies, the prices for ICE 2.0 were developed via an iterative testing process utilizing a pre-test study (Hanemann et al., 1991). Initially, a small number of preliminary prices were chosen based on the results of previous interruption cost studies. These prices were presented to respondents in the pre-test for the first utility enrolled in ICE 2.0 in the fall of 2022.

After administration of the pre-test, the distribution of "Yes" and "No" responses to these prices were reviewed and adjusted. Prices that were near-universally accepted (i.e., where the vast majority of respondents said "Yes") were increased, and prices that were near-universally rejected were lowered. Existing literature on OHDC bid design recommends, in the absence of knowing the true distribution of costs, selecting bids that are centered on the estimated median cost but do not drift too far into the upper or lower tails (Kanninen, 1993). Following this recommendation, we selected a series of lower bid values that on average yield between a 60/40 to 70/30 Yes/No split, and a series of upper bid values that on average yield between a 30/70 to 40/60 No/Yes split.

Table A1 presents the current price values presented to respondents in the ICE 2.0 study. For ICE 2.0, five "sets" of high and low prices were developed for each duration scenario. Each respondent was randomly assigned one of these sets, and therefore received the corresponding price values for each duration. These sets were designed to be internally consistent within a given set, where the high price for a longer duration would always be higher than the high price for a shorter duration and the low price corresponding to a longer duration was higher than the low price corresponding to a shorter duration within the same set. This ensured that each respondent received a plausible set of prices where price increased with duration.

Table A1. Residential Price Values

Set	Duration							
	Momentary		2 Hours		8 Hours		24 Hours	
	Low Price	High Price	Low Price	High Price	Low Price	High Price	Low Price	High Price
1	$0.10	$1.00	$0.50	$4.00	$3.00	$10.00	$5.00	$20.00
2	$0.15	$2.00	$1.00	$10.00	$5.00	$25.00	$10.00	$40.00
3	$0.20	$4.00	$2.00	$12.00	$7.00	$40.00	$15.00	$60.00
4	$0.25	$6.00	$3.00	$14.00	$9.00	$55.00	$20.00	$80.00
5	$0.30	$8.00	$4.00	$16.00	$11.00	$70.00	$30.00	$120.00

Appendix B: Sampling Strategy

Section 2 described the consumption-based strata used to develop the initial samples of customers that were selected. This Appendix describes the procedures used to recruit customers from these samples.

To ensure that the survey samples were representative of the customers served by the Phase 1 utility sponsors, two sequential procedures were used to select and recruit customers to take the surveys.

Residential Customers

Data from national sources were used to ensure that respondents in the survey were representative of the demographics of each Phase 1 utility customer population. The information is used to pre-screen potential respondents during the survey recruitment process.

The U.S. Census Bureau makes available estimates of the number of household members in each age group, income group, and age/income combination at the ZCTA (Zip Code Tabulation Area) level. ZCTAs are generalized area representations of ZIP Code service areas identified by the most common ZIP code in each census tract. Based on this relationship, estimates of age, income, and age/income groups at the ZCTA level were used to roughly estimate the demographics of the population living in a given ZIP code.

To estimate the population of customers, we used the average demographics of all the ZCTAs that correspond to ZIP codes in which the utilities had active customers in the customer population data, weighing the results by the number of customers in each ZIP code. By weighting these demographics by the number of customers located in each ZIP code, we estimated the residential customer population. Figure B1 provides an example of the process for estimating demographics for a hypothetical utility with 30 customers living in two ZIP codes with two corresponding ZCTAs.

Figure B1. Utilizing ZCTA Demographics

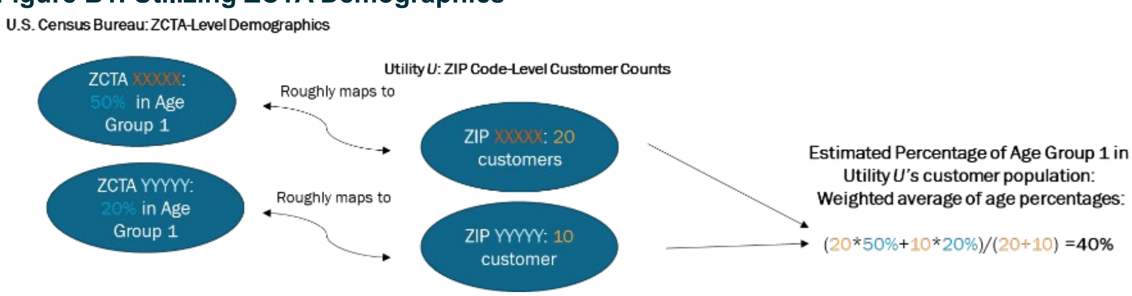

These estimated demographics were used in survey pre-screening to ensure that the respondents who completed the survey were representative of the customer population at large. The key demographics used for pre-screening were usage, age, and income level.

Survey bias can be a concern, especially if respondents of some usage classes, age groups, or income levels are more likely to complete the survey than others. To address this, caps were defined identifying the maximum number of respondents in each usage class, age group, and income level that would be allowed to take the survey. As shown in Table B1, the response caps for each age and income group were informed by the estimated demographics of each utility's service area.

Table B1. Example of Residential Age and Income Response Caps

Income	Age			
	18–30	30–60	60+	Total Income Cap
0–$50,000	25	88	50	**150**
$50,000–$150,000	25	100	50	**150**
$150,000+	25	75	25	**50**
Total Age Cap	**75**	**200**	**88**	**250**

In addition to setting response caps based on respondent age and income, we also set caps for residential customers in each demand stratum. While we initially sampled customers based on these demand strata, the caps for each stratum were set to ensure that the customers who ultimately completed the survey were similar in their distribution of usage to the initial sample of customers who were solicited. These caps were set by establishing a threshold for each stratum slightly above the percentage of that stratum's total population usage. For each stratum, this threshold represented the maximum number of respondents from that stratum, expressed as a percentage of the total number of survey respondents, which would be allowed to complete the survey. Once this threshold was reached for a given stratum, the survey was closed to any remaining respondents whose average demand fell within that stratum. Table B2 shows the caps by usage strata.

Table B2. Example of Residential Stratum Response Caps

Strata	Average Demand	Total Usage Cap
1	0–0.5 kW	25
2	0.5–1.0 kW	50
3	1.0–2.0 kW	150
4	2.0–5.0 kW	150
5	5.0–10.0 kW	25
Total Cap	**All**	**250**

Finally, once the minimum target of total responses (250) was received, all caps based on income, age, and usage were removed. At this point, all respondents were allowed to take the survey to increase the sample size past the minimum target.

Non-residential Customers

Small and medium non-residential customers were defined as customers with an average hourly demand of less than 200 kW, while LNR customers were defined as having a demand greater than 200 kW.[29] These categories include both facilities operated by commercial customers and those operated by government/nonprofit organizations. We sought to survey 250 SMNR and 67 LNR customers. Like those in the residential segment, these customers were also stratified by average demand and sampled by the relative percentage of total usage in the preliminary sample.

The utilities provided customer-level data for non-residential customers, including customer name, address, and NAICS industry code. This information allowed us to filter the preliminary sample based on firmographic characteristics. When sampling non-residential customers, it was critical to include respondents from a variety of companies and industries. To avoid gathering redundant information, we sought to avoid retaining multiple facilities of the same facility type from the same firm in the final sample.

We utilized two criteria to identify whether facilities from the same firm could be considered the same facility type, industry, and usage. First, some firms have facilities that are identified as operating in different industries based on their NAICS codes (e.g., a firm may have both wholesale distribution centers and retail stores). However, even within NAICS codes, one could likely lump facilities into subclasses based on size. For example, in the retail sector, there are small, medium, and large retail stores; each of these subclasses likely has different implications in terms of interruption costs.

To identify subclasses within facility designations, we utilized the utilities' population data and employed k-means clustering within each industry, as defined by two-digit NAICS code. "K-means clustering" is a method of classification that, given a specified number of clusters *k*, assigns observations to clusters based on similarity in one or more specified variables (in this case, electricity demand). For simplicity, Resource Innovations specified the number of clusters, k=3 for each NAICS code. The objective was to assign facilities to a "small," "medium," or "large" cluster, depending on their demand relative to the demand of other facilities with the same NAICS code. Figure B2 presents an example of k-means clustering for the retail trade industry (NAICS codes 44–45) in the small and medium non-residential segment.

[29] For PSE, a split between the SMNR and LNR customer classes of 25 kW was used.

Figure B2. K-means Clustering Example – Retail Trade

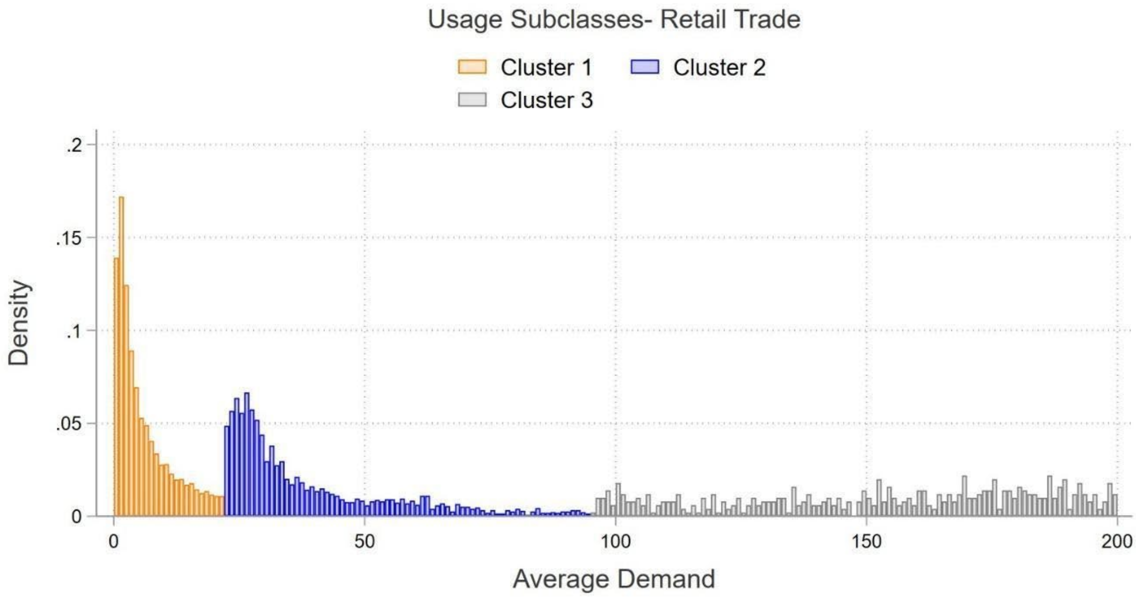

Next, following the identification of clusters and the preliminary sampling of customers based on demand, we identified customers within the preliminary sample that were duplicates in terms of (1) Firm/Organization; (2) NAICS code; and (3) sub-industry demand cluster. Take, for example, firm *i*, which operates three different types of facilities: small retail stores, large retail stores, and wholesale distributors. If several of this firm's facilities in the same cluster within the same industry were included in this preliminary sample (i.e., two small retail facilities), only one of the facilities was retained. If several facilities from the same firm were sampled but were identified as being in different industries based on NAICS code or were in different identified demand clusters within the same industry, all of these facilities were retained. It is important to note that these clusters were based on demand trends within industries among the utilities' customer population, not based on firm- or industry-specific definitions of different facility classes.

Following the preliminary sampling, we also implemented response caps to ensure that respondents completing the survey were similar to the non-residential customer population. For both non-residential segments, these caps were based on industry data (as reported by NAICS code) and usage strata. For each possible usage stratum/industry combination, initial thresholds were set as the percentage of total population usage in each usage stratum and industry, based on the industry information provided in the utilities' population data.

For each industry/stratum combination, the survey would remain open until one of four criteria is reached:

1. The total sample size for the entire population is achieved
2. The sample quota for the total industry is reached
3. The sample quota for the total demand stratum is reached
4. The sample quota for the industry/stratum combination is reached.

Table B3 and Table B4 display the response caps for the SMNR and LNR surveys, respectively. Once the minimum target of total responses was received (i.e., 250 for SMNR and 67 for LNR), all caps based on industry and usage were removed. At this point, all respondents were allowed to take the survey to increase the sample size past the minimum target.

Table B3. Example of SMNR Industry and Stratum Response Caps

Survey Industry	Strata 1	Strata 2	Strata 3	Strata 4	Strata 5	Industry Cap Total
Accommodation and Food Services	13	13	50	25	13	**75**
Administrative and Support Services	13	13	13	13	13	**25**
Agriculture, Forestry, Fishing and Hunting	13	13	13	13	13	**13**
Arts, Entertainment, and Recreation	13	13	13	13	13	**25**
Construction	13	13	13	13	13	**25**
Educational Services	13	13	13	13	13	**25**
Finance and Insurance	13	13	13	13	13	**13**
Health Care and Social Assistance	13	13	13	13	13	**25**
Information, Data, and Telecommunications	13	13	13	13	13	**25**
Management of Companies	13	13	13	13	13	**13**
Manufacturing	13	13	13	13	13	**25**
Mining	13	13	13	13	13	**13**
Professional, Scientific, and Technical Services	13	13	13	13	13	**25**
Public Administration/Government	13	13	13	13	13	**25**
Real Estate	13	13	13	13	13	**13**
Rental and Leasing Services	13	13	13	13	13	**13**
Retail Trade	13	25	50	25	25	**100**
Transportation	13	13	13	13	13	**13**
Utilities	13	13	13	13	13	**13**
Warehousing and Storage	13	13	13	13	13	**13**
Waste Management and Remediation Services	13	13	13	13	13	**13**
Wholesale Trade	13	13	13	13	13	**13**
Other	13	25	25	13	13	**50**
Stratum Cap Total	**25**	**75**	**125**	**75**	**100**	**250**

Table B4. Example of LNR Industry and Stratum Response Caps

Survey Industry	Strata 1	Strata 2	Strata 3	Strata 4	Strata 5	Industry Cap Total
Accommodation and Food Services	3	3	3	3	3	**7**
Administrative and Support Services	3	3	3	3	3	**7**
Agriculture, Forestry, Fishing and Hunting	3	3	3	3	3	**7**
Arts, Entertainment, and Recreation	3	3	3	3	3	**7**
Construction	3	3	3	3	3	**7**
Educational Services	3	3	3	3	30	**34**
Finance and Insurance	3	3	3	3	3	**7**
Health Care and Social Assistance	3	3	3	3	17	**20**
Information, Data, and Telecommunications	3	3	3	3	3	**7**
Management of Companies	3	3	3	3	3	**7**
Manufacturing	3	3	3	3	10	**13**
Mining	3	3	3	3	3	**7**
Professional, Scientific, and Technical Services	3	3	3	3	3	**7**
Public Administration/Government	3	3	3	3	3	**7**
Real Estate	3	3	3	3	3	**7**
Rental and Leasing Services	3	3	3	3	3	**7**
Retail Trade	3	3	3	3	3	**7**
Transportation	3	3	3	3	3	**7**
Utilities	3	3	3	3	3	**7**
Warehousing and Storage	3	3	3	3	3	**7**
Waste Management and Remediation Services	3	3	3	3	3	**7**
Wholesale Trade	3	3	3	3	3	**7**
Other	7	7	7	7	7	**13**
Stratum Cap Total	**13**	**13**	**13**	**20**	**50**	**67**

Appendix C: Review of Survey Responses

The results from all individual Phase 1 surveys were combined into a single residential and a single non-residential modeling dataset for use in updating the customer damage functions (CDFs) in the ICE Calculator. Prior to this, each survey response was reviewed to confirm its validity and logical consistency and to remove outliers. Separate procedures were developed for each survey. This Appendix first describes these criteria and then presents how they were used to identify and remove responses from the final datasets.

Invalid responses refer to customer responses that indicate they are answering a different question than the one being asked. For example, customers sometimes react to questions about interruption costs by redefining the question so that it relates to their ability to pay, their satisfaction with service, or whether they think they are being fairly charged for the service they receive. Such responses do not accurately reflect the cost of an interruption for a customer. *Protest* responses are invalid responses where respondents overstate their interruption cost or refuse to give a WTP because they believe that their utility should be responsible for paying for the interruption.

Illogical responses refer to cases where a respondent expresses decreasing costs as the interruption duration increases. For instance, a respondent might indicate that a shorter interruption would cost more than a longer interruption. Because the survey respondents received multiple survey scenarios, it is possible to identify respondents that gave illogical responses by comparing their response to a shorter-duration interruption to their response to a longer-duration interruption.

Finally, *outlier* responses are responses that greatly exceed the typical range of interruption costs for a given duration. Outliers are dropped because they have a disproportionate impact on the mean interruption cost. Respondents may erroneously provide unrealistically high estimates because of human error, misjudging their interruption costs, or misunderstanding the question. Due to the structure of the residential surveys, outliers were only a concern for the non-residential segment.

Residential Customers

Eight to nine percent of residential respondents were not included in the final residential dataset because they provided invalid or illogical responses. For the residential survey, the review focused on removing responses that were not based on economic considerations (e.g., unwilling to pay any price to avoid an interruption for reasons other than economic ones) or were internally inconsistent (e.g., willing to pay a higher price to avoid short- vs. long-duration power interruptions).

If a respondent indicated they were unwilling to pay any of the prices presented to them, they were given an open-ended follow-up question asking why they rejected the prices. Responses were examined individually. If the respondent verified they were unwilling to pay because the prices presented exceeded the cost and inconvenience they experienced, their response was confirmed as valid and included in the cost estimate calculations. However, if the respondent reported an unwillingness to pay related to a protest response, then their response was deemed invalid and not included in the cost estimate calculations. For example, some respondents indicated they were unwilling to pay the presented prices because “[their] utility should pay instead,” or “I should not have to pay for this.” These respondents were likely rejecting the price to make a statement about their utility rather than demonstrating a true unwillingness to pay. Because including these responses would bias the results, they were dropped.

Responses were also dropped if respondents expressed a higher WTP for a shorter-duration scenario than a longer-duration scenario, given the same interruption circumstances. For instance, if a respondent accepted a bid value for a momentary duration scenario but rejected an equal or lower bid value for a 2-hour duration scenario, all their responses were dropped.

Due to the nature of the OHDC contingent valuation methodology, the results received from the residential segment consist only of a series of responses to the price bids presented, instead of a distribution of actual costs as given in an open-ended cost survey. For this reason, the survey responses in the residential segment did not have “outliers” in the sense that open-ended surveys do, so conducting an outlier analysis similar to that in the non-residential segment was not necessary.

Table C1 summarizes the prevalence of invalid and illogical responses by interruption duration in the residential survey.

Table C1. Summary of Invalid Responses – Residential

Interruption Duration	Total Respondents	Invalid or Illogical Respondents		Valid Respondents
		N	%	
Momentary	2,514	202	8.0%	2,312
2 Hours	2,461	229	9.3%	2,232
8 Hours	2,492	222	8.9%	2,270
24 Hours	2,481	217	8.7%	2,264

Non-residential Customers

Thirteen to fourteen percent of non-residential responses were not included in the final non-residential dataset for any one of four reasons: first, if they were a residential or rental property; second, if their comments suggest a protest response; third, if they reported illogical responses; and fourth, if their response was deemed an outlier.

Customer information is imperfect, and occasionally the non-residential survey was sent to residential customers. Thus, the analysis screened for instances where residential customers responded to the non-residential survey. This was generally achieved by reviewing the response comments and flagging surveys in which the respondent reported they were a residential customer.

While reviewing the survey results, it also became clear that some non-residential customers used the survey to vent frustrations with their utility. Therefore, the analysis screened for protest responses that suggested the customer did not earnestly answer the survey questions. To this end, the surveys were screened for hostile language in the comments. Screening for these comments is a particularly important step, as some customers may genuinely experience $0 costs while others may report a zero value out of protest.

Next, survey responses were also screened to remove illogical responses. Illogical responses reported a higher cost for a shorter-duration interruption than for a longer one during the same day type, season, and onset time.

Finally, interruption costs were screened separately for each duration to identify (but not remove) large outliers. Outliers were identified by first normalizing interruption cost by average demand (kW) and then calculating the log of each normalized interruption cost per kW. Normalizing the costs by usage avoids excluding high-usage customers with large interruption costs. Taking the log of each normalized costs yields a normal distribution. The log normalized interruption larger than 1.5 times the interquartile range plus the 75th percentile was marked as an outlier. Outliers were included in the final dataset but were not included in the calculation of the mean interruption costs from the surveys.

Table C2 summarizes the prevalence of invalid responses by interruption duration in the non-residential surveys. The percentage of responses deemed invalid varied from 13.1% for a momentary interruption to 14.5% for a 24-hour interruption. The majority of invalid responses were illogical responses.

Table C2. Summary of Invalid, Illogical, and Outlier Responses – Non-residential

Interruption Duration	Total Respondents	Invalid or Illogical Respondents		Valid Respondents	Outlier Respondents
		N	%		
Momentary	3,421	447	13.1%	2,974	66
2 Hours	3,383	471	13.9%	2,912	66
8 Hours	3,342	473	14.2%	2,869	57
24 Hours	3,354	488	14.5%	2,866	63

Appendix D: Sample Weighting

This Appendix describes methods used to aggregate survey responses to ensure the calculations incorporating them yield results that are statistically representative of the populations from which they were drawn.

As discussed in Section 2.3, we intentionally oversampled customers with higher electricity usage because past research has shown that customers with high usage report the highest and greatest spread in interruption costs. For example, if customers in the highest usage stratum, taken together, accounted for 20% of the total energy usage (in kWh) of the entire customer population, we sought to recruit 20% of the total number of surveys from customers in this stratum. Oversampling, in this instance, refers to the fact that the customers in this stratum account for far less than 20% of the utility's total number of customers. Therefore, in order to estimate the mean interruption costs for the entire population, the responses from these customers must be weighted lower than those from customers in the other, lower-usage strata.

Table D1 presents the percentages of the residential population and sample in each stratum.

Table D1. Residential Strata Distribution in Population and Sample

Stratum	Usage Category (Average kW)	Percent of Population in Strata	Percent of Survey Responses in Strata
1	0–0.5	18.5%	4.0%
2	0.5–1	29.0%	19.1%
3	1–2	37.1%	41.1%
4	2–5	15.0%	33.9%
5	5–10	0.4%	1.8%
Total		**100%**	**100%**

Table D2 presents the percentages of the non-residential population and sample in each stratum.

Table D2. Non-residential Strata Distribution in Population and Sample

Stratum	Usage Category (Average kW)	Percent of Population in Strata	Percent of Survey Responses in Strata
1	0.25–2	48.4%	6.6%
2	2–10	32.7%	23.2%
3	10–50	13.7%	33.3%
4	50–100	2.5%	13.0%
5	100–200	1.3%	7.1%
6	200–400	0.7%	8.3%
7	400–1,000	0.4%	4.8%
8	1,000–2,000	0.1%	1.6%
9	2,000–5,000	0.07%	1.4%
10	>5,000	0.05%	0.6%
Total		**100%**	**100%**

The strata weights are described in Equation D1. Let S_n represent the number of customers in the n^{th} strata at the population level. Similarly, let s_n represent the number of survey respondents that fall within the n^{th} strata. The data weights, denoted as sw_n, are calculated by dividing the proportion of the population that fit within the strata by the proportion of the survey respondents.

Equation D1. Strata Weights

$$P_n = \frac{S_n}{\sum_{i=1}^{10} S_i}$$

$$p_n = \frac{s_n}{\sum_{i=1}^{10} s_i}$$

$$SW_n = \frac{P_n}{p_n}$$

In addition to the above strata weights, interruption scenarios were also weighed to reflect the fact that interruptions are more likely to occur during some times of the year. To reflect this, interruption weights were also introduced. These weights are described by Equation D2. Let t_n represent the number of interruptions that occur in one of three considered seasonal interruption scenarios: a winter weekday interruption, a summer weekend interruption, and a summer weekday interruption. The interruption weights, IW_n, are calculated by finding the proportion that each scenario occurs.

Equation D2. Interruption Weights

$$IW_n = \frac{t_n}{\sum_{i=1}^{3} t_i}$$

Appendix E: Additional Survey Findings

This section presents selected additional findings from the surveys that either by design or through our analysis did not influence the final customer damage functions (CDFs), which are the focus of this report. We first review the qualitative responses customers provided describing how they would respond or react to the hypothetical three-day power interruption scenario described in Section 2.2. We then discuss both the residential and non-residential responses regarding back-up generation. Next, we summarize the information residential respondents reported on working from home. Finally, we present the firmographic information reported by non-residential respondents.

Responding to a Long-duration Power Interruption

In addition to the four short-duration interruption cost questions, respondents were also asked about how they would respond to a longer-duration interruption scenario. The scenario was described as an interruption affecting a 20-mile radius around the respondent's home or facility, and they were told they could expect the interruption to last up to three days (72 hours). This scenario was presented as occurring in either the summer or the winter, with the season of the scenario being randomly assigned.

Residential respondents were asked how they would respond to this long-duration interruption scenario, with possible choices including:

- Stay home and do activities that do not require electricity
- Stay home running your own backup generator or battery storage unit to power critical appliances and devices
- Temporarily move some or all family members to places that have electrical power (e.g., houses of family or friends outside the affected area, hotels, or emergency shelters outside the affected area)
- Other – Please explain: ____

Figure E1 presents the responses to this question. The majority of respondents indicated that they would temporarily move some or all family members to places that have electrical power. Other responses included "permanently relocat[ing] to an area that can supply power," "us[ing] neighbor's generator for refrigerators," "go[ing] camping," and doing a combination of the listed responses. Some respondents who selected "Other" indicated that an interruption would be highly challenging for their household due to medical conditions. For instance, some respondents indicated that they or family members use CPAP machines that require electricity. The responses were not significantly different by season.

Figure E1. Residential Responses to a Long-duration Interruption

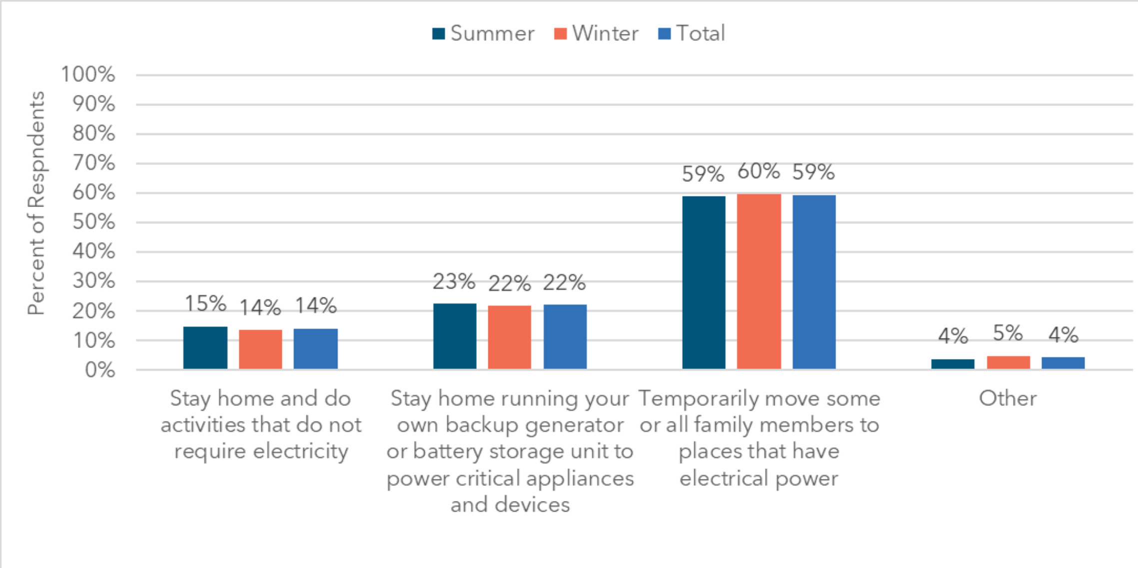

Similarly, non-residential customers were asked how they would respond this same interruption scenario (i.e., affecting a 20-mile radius around their facility and expected to last up to three days).

Non-residential respondents were asked how they would respond to this long-duration interruption scenario, with possible choices including:

- Shut down the facility throughout the power interruption
- Continue operation with backup generation
- Temporarily shut down the facility until backup generation is rented
- Transfer activities to another facility that has electricity

Figure E2 displays the non-residential responses. The majority of respondents indicated that they would shut down the facility for the duration of the interruption. The least common response was to transfer activities to another facility that has electricity. The responses were not significantly different by season.

Figure E2. Non-residential Responses to a Long-duration Interruption

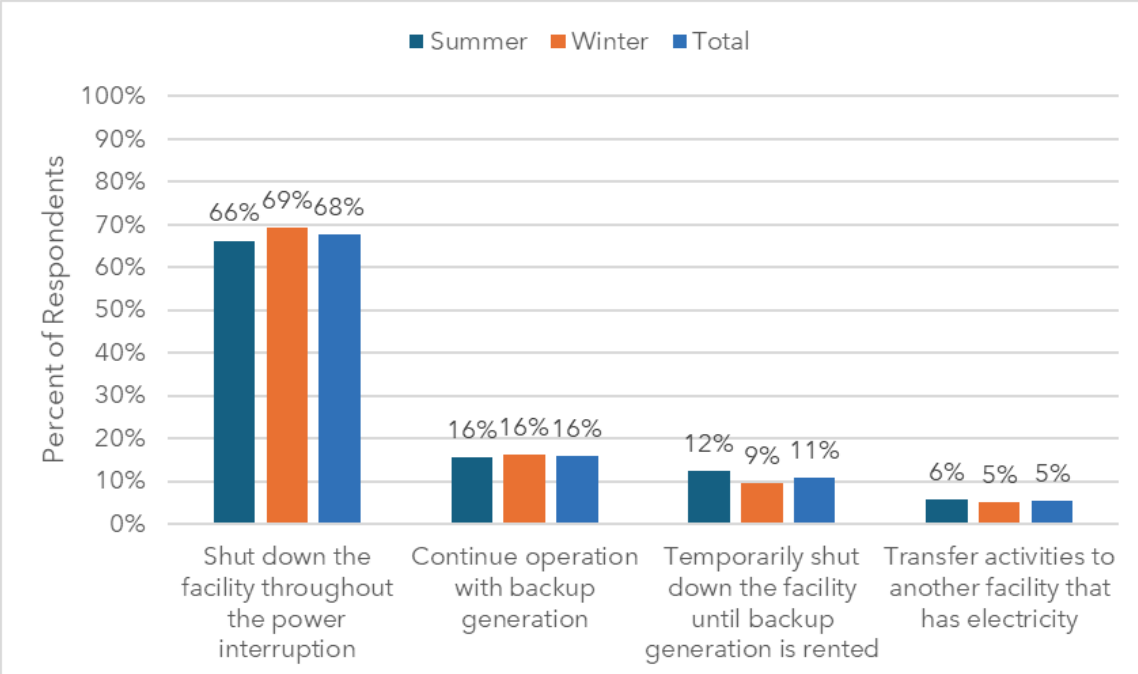

Ownership of Backup Generation

Survey respondents were also asked whether they had backup generation equipment that could be utilized during a power interruption. This question was intended to give the researchers a sense of how resilient customers would be to interruptions.

Table E1 presents residential ownership of backup generation by usage stratum. Generally, respondents in the higher usage strata had higher ownership rates of backup generation, with ownership in the lowest stratum (0 to 0.5 kW) being 11.3%, and ownership in the highest stratum (5 to 10 kW) being 40.0%. The population-weighted average prevalence of backup generation across the entire survey population was 19.0%.

Table E1. Residential Ownership of Backup Generation

Stratum	Usage Category (Average kW)	Respondents with Backup Generation	% of Population in Stratum
1	0–0.5	11.3%	18.5%
2	0.5–1	15.6%	29.0%
3	1–2	20.4%	37.1%
4	2–5	30.8%	15.0%
5	5–10	40.0%	0.4%
Total		**19.0%**	**100.0%**

Table E2 presents non-residential ownership of backup generation by usage stratum. The prevalence of backup generation increases with usage for strata 1 through 5 (usage ranging from 0.25 to 200 kW). Strata 6 through 9 (usage ranging from 200 to 5,000 kW) had backup generation at a frequency that fell between 48% and 60%. The largest stratum (5,000 kW and above) had backup generation at a rate of 71%. The population-weighted average prevalence of backup generation across the entire survey population was 15.5%.

Table E2. Non-residential Ownership of Backup Generation

Stratum	Usage Category (Average kW)	Respondents with Backup Generation	% of Population in Stratum
1	0.25–2	11.0%	5.6%
2	2–10	13.6%	21.1%
3	10–50	18.8%	32.7%
4	50–100	29.7%	11.5%
5	100–200	41.7%	7.1%
6	200–400	59.3%	10.7%
7	400–1,000	55.2%	6.1%
8	1,000–2,000	56.9%	2.3%
9	2,000–5,000	48.3%	2.0%
10	5,000+	70.7%	1.0%
Total		**15.5%**	**100.0%**

Work From Home

Residential respondents were asked whether they or any household members earned an income working from home. This question was intended to give a sense of how households with telecommuters would be impacted by interruptions.

Table E3 presents results from this question by usage stratum. Respondents in the higher usage strata were more likely to report themselves or a household member working from home than those in the lower strata. The weighted population average prevalence of working from home was 37.0%.

It should be noted that survey results were collected during the post-Covid period from December 2022 to April 2024, when many individuals continued to work from home.

Table E3. Residential Work From Home Prevalence

Stratum	Usage Category (Average kW)	Respondents Who Work From Home	% of Population in Stratum
1	0–0.5	30.1%	18.5%
2	0.5–1	37.0%	29.0%
3	1–2	38.6%	37.1%
4	2–5	41.2%	15.0%
5	5–10	50.0%	0.4%
Total		**37.0%**	**100.0%**

Those respondents who said they or a household member worked from home were also asked to note whether they worked for an hourly wage, a salary, were self-employed, or other. Figure E3 presents the prevalence of each of these categories. The majority of respondents who said they or a household member worked from home mentioned that they or their household members were salaried employees (53.2%), with a smaller percentage being self-employed (26.6%) or working for an hourly wage (18.1%). Respondents who selected "Other" in some cases indicated they performed both salaried and self-employed work or worked for a commission.

Figure E3. Residential Work From Home Employment Types

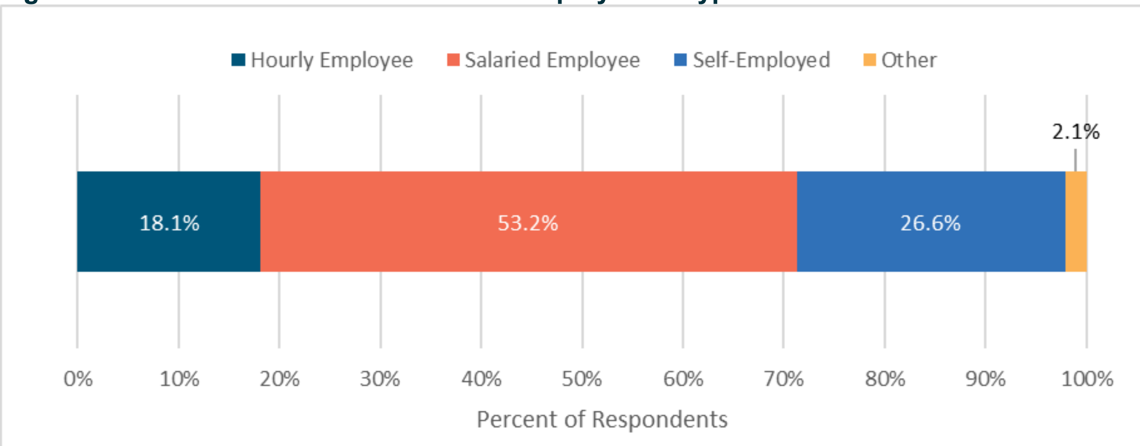

Industries Represented by Non-residential Respondents

Non-residential respondents were asked to identify what industry type best describes their establishment. The industry types are based on the two-digit NAICS codes. Respondents also had the opportunity to write in the industry type that best describes their establishment. For such cases, they were assigned a NAICS code that best fit their description. These results are summarized in Table E4.

Table E4. Non-residential Industry Type Prevalence

Industry	NAICS Code	Respondents
Agriculture, Forestry, Fishing and Hunting	11	1.8%
Mining, Quarrying, and Oil and Gas Extraction	21	0.8%
Utilities	22	1.5%
Construction	23	2.3%
Manufacturing	31-33	17.3%
Wholesale Trade	42	3.7%
Retail Trade	44-45	12.9%
Transportation and Warehousing	48-49	2.5%
Information	51	1.5%
Finance and Insurance	52	1.5%
Real Estate and Rental and Leasing	53	3.5%
Professional, Scientific, and Technical Services	54	4.7%
Management of Companies and Enterprises	55	0.3%
Administrative and Support and Waste Management and Remediation Services	56	1.5%
Educational Services	61	5.6%
Health Care and Social Assistance	62	8.6%
Arts, Entertainment, and Recreation	71	3.9%
Accommodation and Food Services	72	12.3%
Other Services (except Public Administration)	81	11.9%
Public Administration	92	2.0%
Total	**-**	**100%**

Appendix F: Selection of a Functional Form to Represent Costs of Varying Interruption Event Durations

This Appendix describes the development of the functional form used to represent costs as a function of the interruption duration. The development of an appropriate functional form was guided by two considerations. First, the selected functional form had to be able to produce estimates (or predictions) of interruption costs that were similar to the interruption costs reported in the surveys. As presented in Section 2.7, direct analysis of the survey results produced population-weighted average interruption costs for four interruption durations: momentary, 2-hour, 8-hour, and 24-hour interruptions. Second, the selected functional form also had to produce internally consistent costs for all durations between 0 and 24 hours. For example, the selected functional form had to produce an estimate (or prediction) of the cost of a 4-hour interruption that exceeded the population-weighted average survey cost for both a momentary and 2-hour interruption, and that did not exceed the cost for an 8-hour interruption. Accordingly, we restricted the possible functional forms for duration to those that were monotonically increasing in cost with respect to duration over the region from 0 to 24 hours. With these constraints in mind, we tested a series of functional forms to meet these requirements. Separate functional forms were developed for the residential and non-residential CDF.

Residential Functional Form Selection

For the residential CDF, the first functional form examined involved only a single, linear term for duration. Since a log link function of the Poisson regression is used, including duration as a linear input does *not* imply a linear relationship between duration and cost. Figure F1 shows the average predictions yielded by this functional form for durations between 0 and 24 hours, as compared to the mean costs for the four included durations. While this functional form satisfies the constraint of a monotonically increasing function with respect to duration, the predicted costs for the four surveyed durations are not particularly accurate. For instance, for an 8-hour interruption, the mean *observed* survey cost is $29.97, whereas the *predicted* cost for the same duration is $16.34. This indicates that this functional form is not appropriate.

Figure F1. Residential Model Test – Average Cost Predictions with a Single Linear Term for Duration

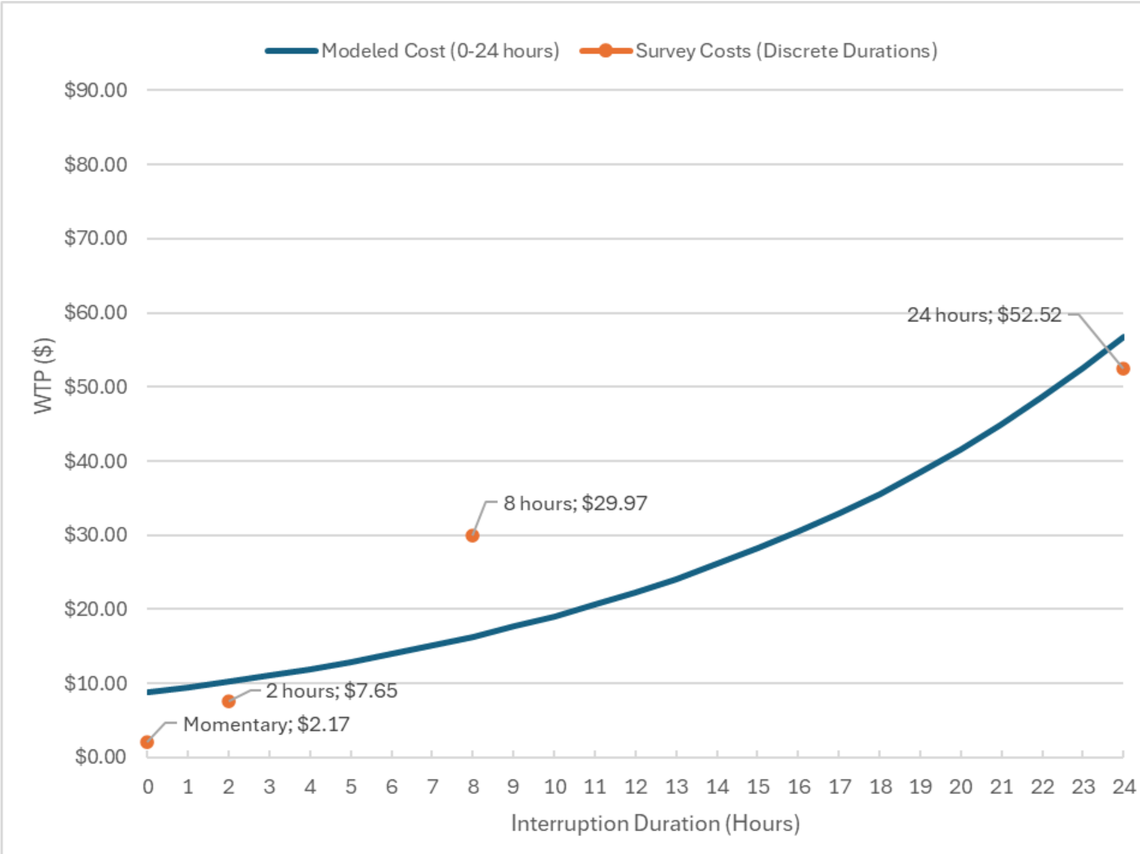

Next, several additional functional forms were tested utilizing a linear term for duration as well as squared and cubed terms (duration + duration^2 and duration + duration^2 + duration^3). While these functional forms produced more accurate predictions for the included durations, they failed to satisfy the constraint of monotonically increasing costs. Figure F2 and Figure F3 show cost predictions for a linear and squared term, and a linear, squared, and cubed term for duration, respectively. The predictions for each of the four included durations were very accurate when compared to the observed surveyed duration means, but both functional forms yielded decreasing cost with respect to duration in some ranges between 0 and 24 hours.

Figure F2. Residential Model Test – Average Cost Predictions with Duration and Duration²

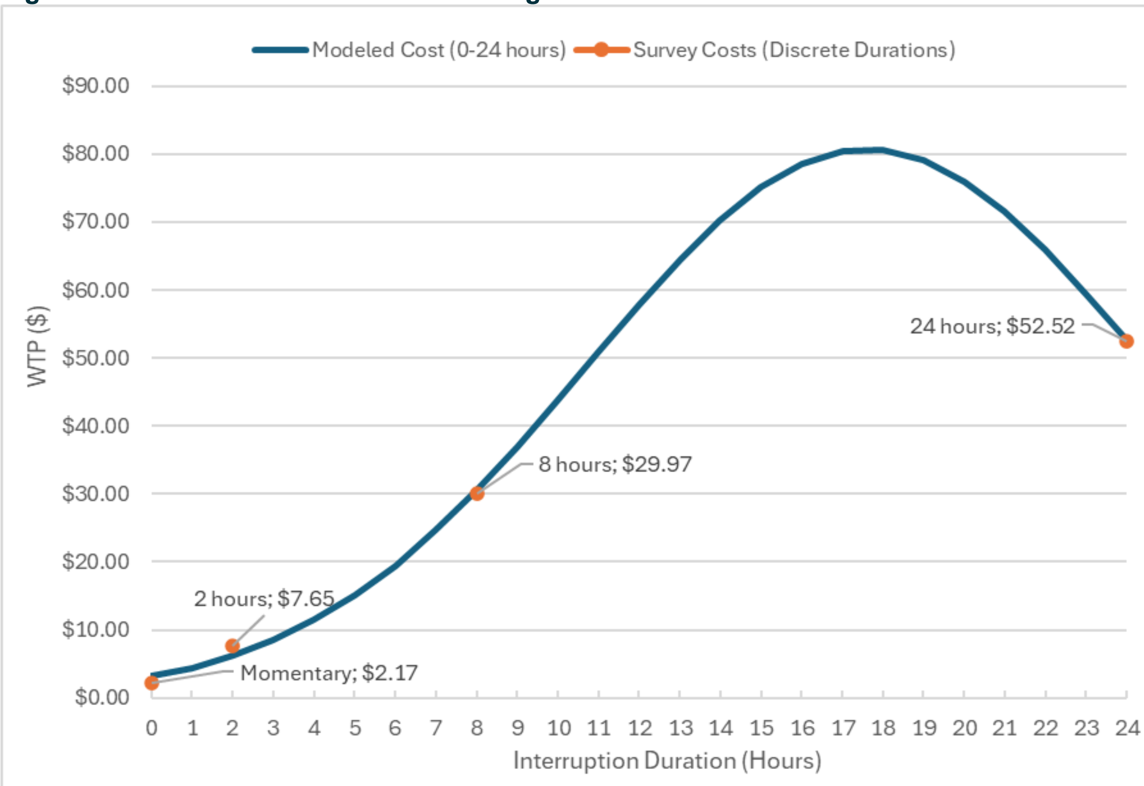

Figure F3. Residential Model Test – Average Cost Predictions with Duration, Duration^2, and Duration^3

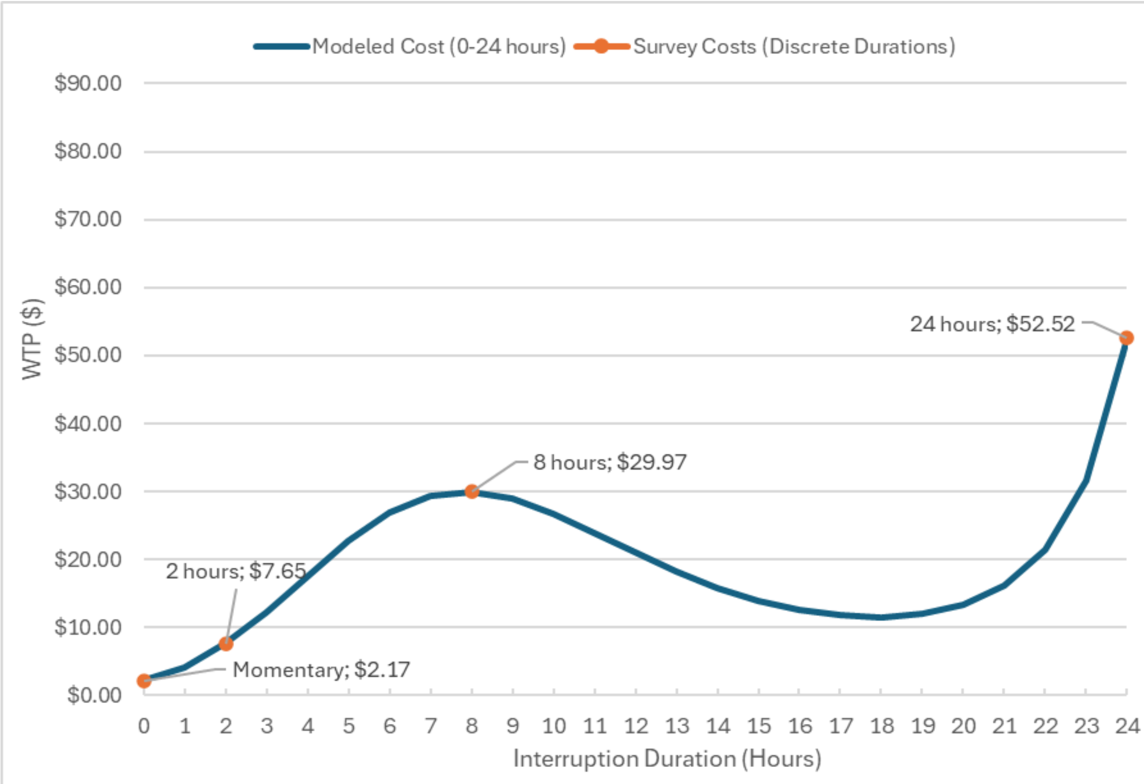

The functional forms shown in Figures F1 through Figure F3 used untransformed duration values (in minutes). The next series of functional forms tested transformations of duration. A common transformation for explanatory variables is a natural log transformation, where the outcome variable y is predicted as a function of ln(x). Natural log transformations are common in econometrics because they are straightforward to interpret: the relationship between y and ln(x) is approximately equal to the relationship between y and a percentage change in x. Several natural log functional forms were tested.

Figure F4 shows the average predicted interruption costs of the final functional form selected, which utilized the natural log of duration and duration squared (ln(duration) and ln(duration)^2). This functional form satisfies both constraints: costs are reasonably accurate with respect to the survey results, and they increase monotonically with respect to duration. This functional form also predicts that costs will increase at a decreasing rate as the duration of an interruption increases, which again is consistent with the survey findings.

Figure F4. Residential Model Test – Average Cost Predictions with Ln(Duration) and $Ln(Duration)^2$

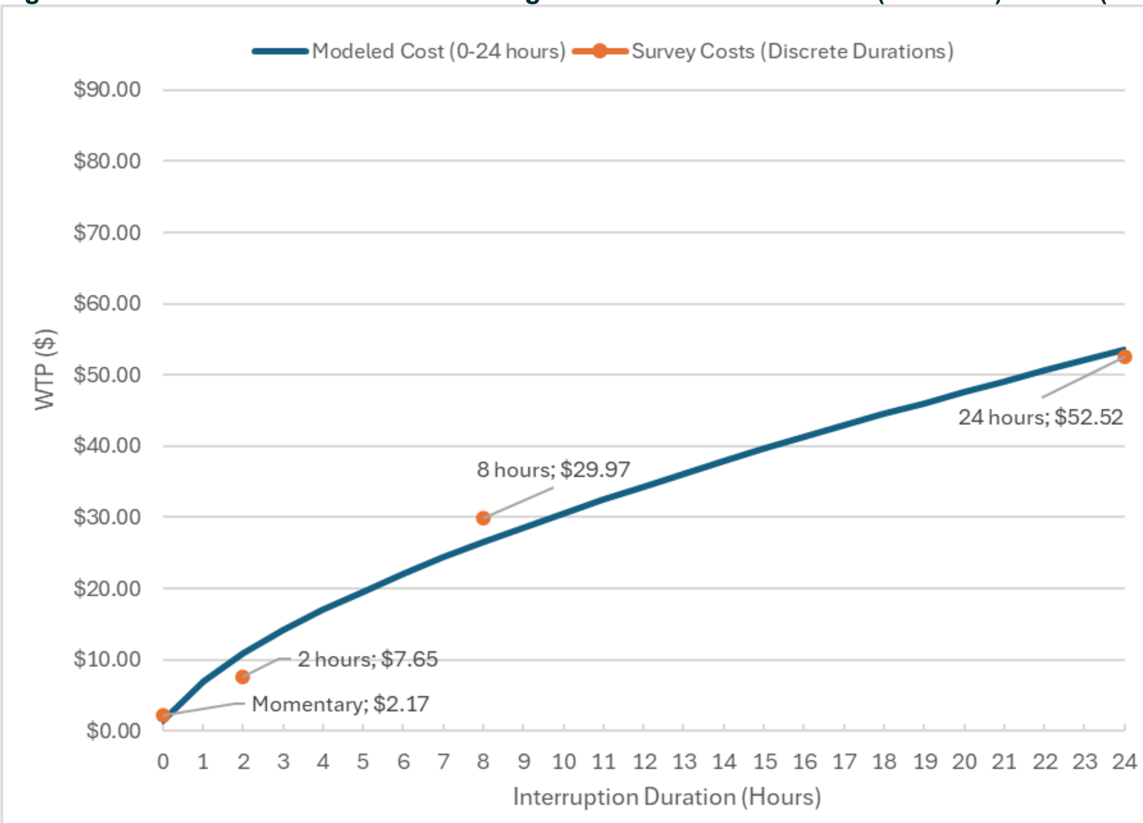

Non-residential Functional Form Selection

For the non-residential CDF, we considered various functional forms by examining different transformations to the duration variable following roughly the same sequence of testing used for the residential CDF. The first functional form considered for both the probit and GLM model was a single, linear term for duration (in minutes). As noted earlier, since the GLM model is implemented with a log link function, including duration as a linear input does not imply a linear relationship between duration and cost. Figure F5 compares the resulting model's average relationship between cost and duration with the costs directly calculated from the surveys. This functional form satisfies the constraint of a monotonically increasing function. However, it does not accurately predict the survey costs at the four interruption durations used in these surveys.

Figure F5. Non-residential Model Test – Average Cost Predictions with Duration

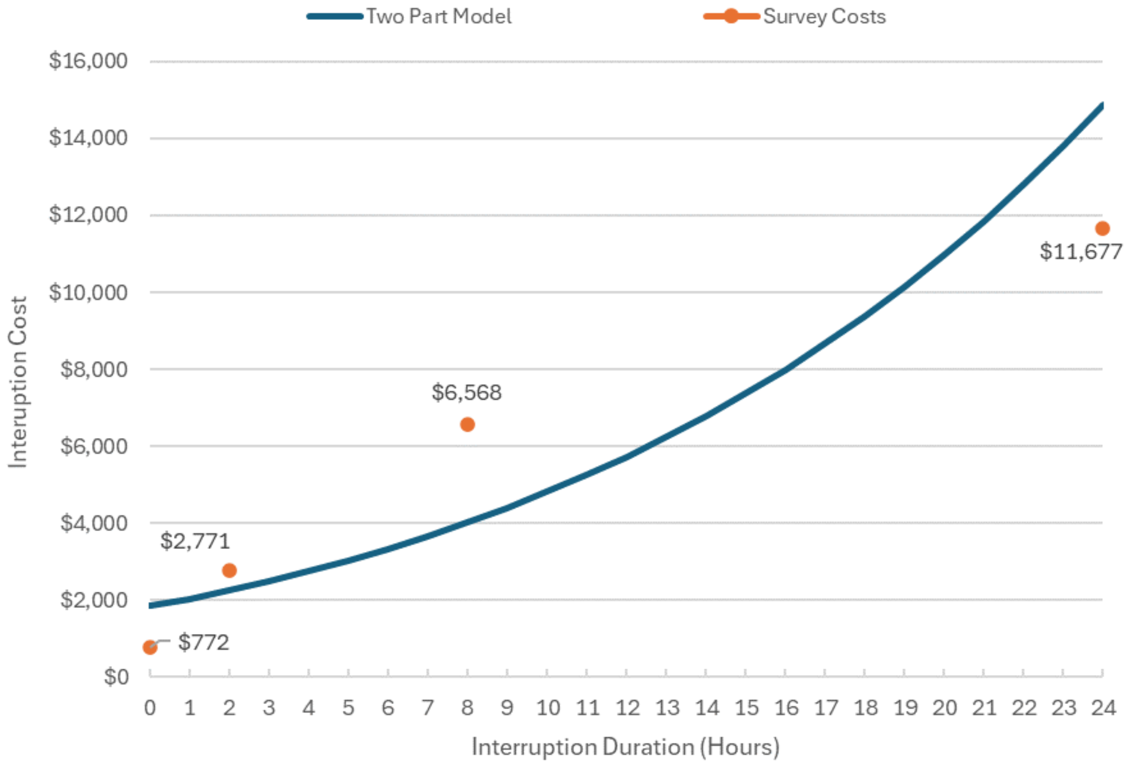

Next, we considered the functional form with duration transformed by the natural log to the first and third power. During the variable selection process, the probit model only included log(duration). However, the LASSO regression suggested GLM models that also included both log(duration) and $\log(\text{duration})^3$. To identify which of these functional forms were more appropriate, we considered two-part models that include log(duration) in the probit but different functional forms in the GLM.

Figure F6 shows the average cost predictions of the two-part model with a GLM that only includes log(duration). This model performs well for the 24-hour duration but less well for the shorter interruptions. More specifically, the model predictions for a momentary interruption cost had a relative error of 32%, a 2-hour interruption cost had a relative error of 22%, an 8-hour interruption cost had a relative error of 4%, and a 24-hour interruption cost had a relative error of 1%.

Figure F6. Non-residential Model Test – Average Cost Predictions with log(Duration)

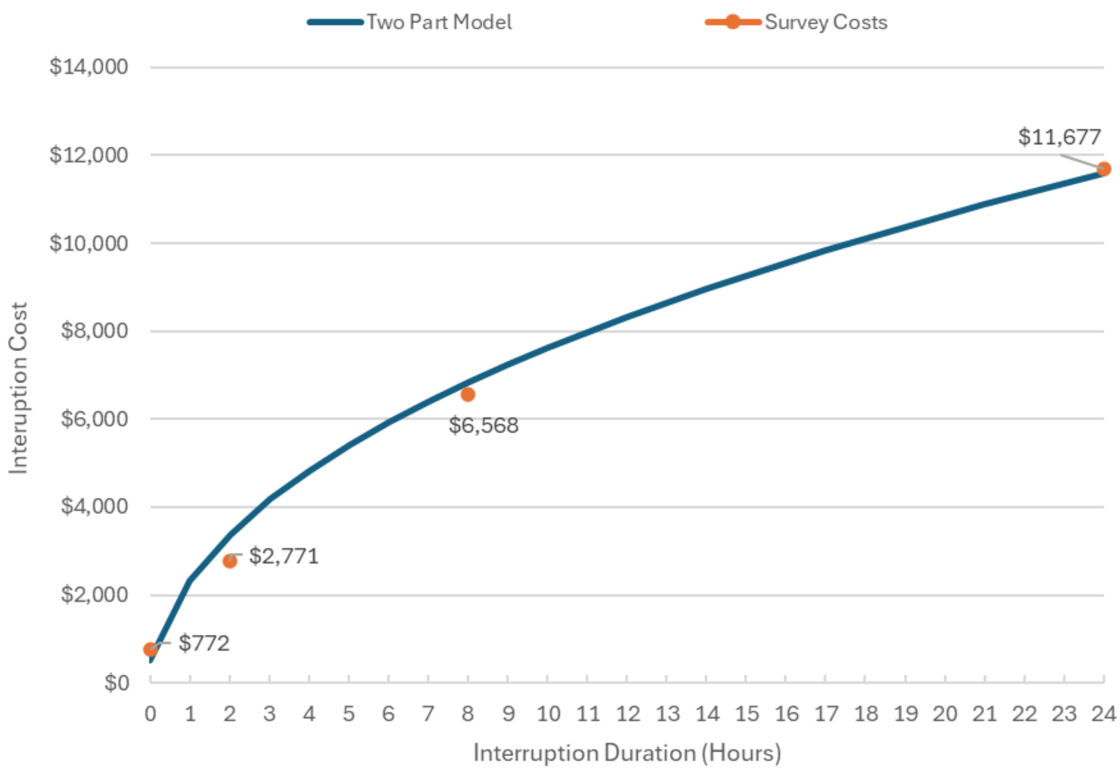

Figure F7 shows the average cost predictions of the two-part model with a GLM that includes both log(duration) and log(duration)^3. This model tended to be fairly accurate across all interruption durations. The model predictions for a momentary interruption cost had a relative error of 21%, a 2-hour interruption cost had a relative error of 3%, an 8-hour interruption cost had a relative error of 6%, and a 24-hour interruption cost had a relative error of 8%.

Figure F7. Non-residential Model Test – Average Cost Predictions with log(Duration) and log(Duration)^3

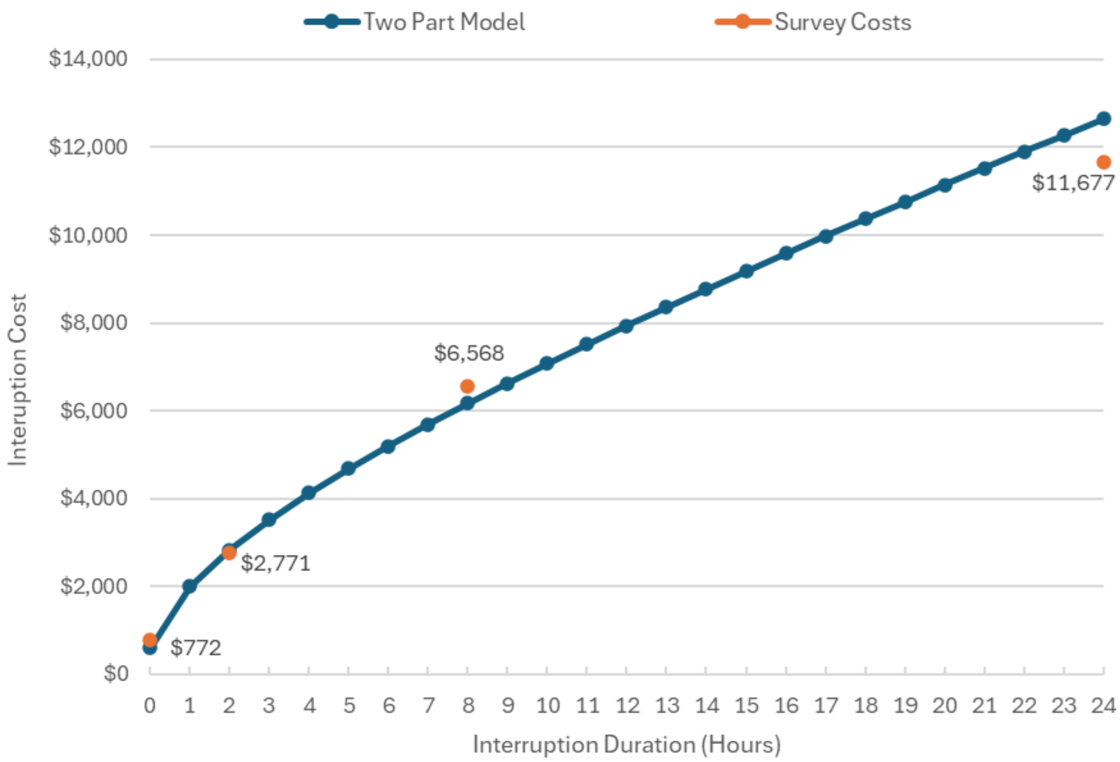

Based on the relative errors, the model that includes both log(duration) and log(duration)^3 is more accurate than the model that included only log(duration). The model that included both terms had a maximum relative error of 21%, while the model that only included log(duration) had a maximum relative error of 32%. Furthermore, the log(duration) model had an average relative error of 14.75%, while the model that included both terms had an average relative error of 9.5%. Thus, we utilized the functional form that included both log(duration) and log(duration)^3 for the non-residential CDF.

Appendix G: LASSO Regression

The initial selection of explanatory variables for inclusion in both the residential and non-residential CDFs was guided by an automated selection process called Least Absolute Shrinkage and Selection Operator (LASSO). LASSO is a regression method commonly used for variable selection that is designed to avoid overfitting, and results in accurate but parsimonious models (Desboulets, 2018).

The LASSO regression method implements an important improvement that makes it superior to the historically more common method, called “forward stepwise selection.” As its name suggests, the forward stepwise selection method operates by adding variables one at a time to build up a final regression model. Typically, variables are added in order of the greatest additional explanatory power contributed by the added variable. However, the well-documented disadvantage of this approach is that once a variable has been chosen, it is never retested after additional variables have been added (Doornik, 2009). Several studies have found that not retesting (and sometimes rejecting variables that have already been selected) leads to biased estimates and inconsistencies in the final variables selected (Hurvich and Chih-Ling, 1990; Steyerberg, et al., 1999; Whittingham et al., 2006).

The LASSO regression avoids these limitations. Instead of selecting variables sequentially for inclusion, LASSO considers all potential explanatory variables simultaneously. LASSO then assesses the individual contribution of each variable to the overall explanatory power of a model by systematically testing the influence of all variables in a consistent manner. By establishing different testing thresholds (using a penalty term), LASSO identifies which variables – when all are taken together – contribute more (and less) than the other variables to the overall explanatory power of a model.

LASSO regression models have the general form:

$$\min_{\boldsymbol{w} \in \mathbb{R}^p} \left\{ RSS(\boldsymbol{w}) + \lambda \sum_{i=1}^{p} |w_i| \right\}$$

where RSS represents the Residual Sum of Squares, λ is the tuning parameter, and w is the coefficient vector.

The tuning parameter is a non-negative value that adjusts the strength of the penalty term. When the tuning parameter is set to zero, the penalty term is dropped, and the model will produce the same coefficients as a least squares regression involving all the explanatory variables. When the tuning parameter approaches infinity, all coefficients become zero. When the tuning parameter is set between these two extremes, the coefficient values will be between zero and that which would emerge from a least squares regression. The rate at which individual coefficients shrink toward zero (as the penalty term is increased) depends on the influence (or importance) of the corresponding covariate. Hence, potential explanatory variables (or covariates) that do not have a large impact on (i.e., do not contribute significantly to explaining) the dependent variable get dropped (i.e., their coefficients become zero) for relatively small values of λ. Variable selection is accomplished by using the remaining variables with non-zero coefficients for a given λ.

LASSO regression has several practical limitations that affect its application as a variable selection tool (Freijeiro-González et al., 2022). Notably, the penalty term introduces bias that increases with the tuning parameter. To overcome this bias, we implemented a two-step process. First, LASSO regression is used to select covariates. Then, in a second step, the selected covariates are used to build a separate Generalized Linear Model that does not include a penalty term. Similar two-step processes have been previously studied (Belloni and Chernozhukov, 2013; Taylor and Tibshirani, 2015).

In addition, LASSO regression has trouble distinguishing between strongly correlated covariates and will often randomly choose one among them and drop the rest. This phenomenon does not affect the current model, as none of the available covariates are highly correlated with one another. The largest correlation between covariates was 0.33 for residential and 0.35 for non-residential.

Finally, it has been well established that LASSO regression always introduces noise to the model and may therefore mistakenly include noisy covariates (Su et al., 2017). This potential issue was mitigated by using cross-validation to choose the final model.

As previously noted, the covariates included by a LASSO regression depend on the value of the tuning parameter, λ. Many procedures have been purposed to identify an optimal value for λ. These procedures can be widely classified into three general categories: generalized information criteria (e.g., AIC or Bayesian Information Criteria (BIC)), resampling procedures such as cross-validation, and reformulations of the LASSO optimization problem (Homrighausen and McDonald, 2018). Cross-validation is widely used and has been shown to be an effective method of selecting λ (Homrighausen and McDonald, 2017). We employed a cross-validation approach to select the final tuning parameter value.

Appendix H: Confidence Intervals

For both the residential and non-residential CDFs, bootstrapping was used to calculate confidence intervals. Although confidence intervals for the CDFs are not presented in discussing the results from the application of the CDFs in Section 3.3 and 4.3, they are incorporated into the online ICE Calculator. This Appendix describes how they were developed.

Bootstrapping is a statistical process that begins by first calculating an average interruption cost using the full population of survey respondents. Next, a sample of survey respondents equal in size to the original population is randomly sampled (with replacement)[30] and the mean interruption cost is calculated using this sample of respondents. Then, another sample is taken and the mean interruption cost is calculated using this second sample. This process is repeated until mean interruption costs from 100 distinct samples drawn (with replacement) from the full population have been calculated. These 100 means form a *distribution* of values that is centered on the mean interruption cost of the full population. Accordingly, the distribution of sample means can then be used to estimate standard errors and corresponding confidence intervals of the population mean interruption cost. Figure H1 illustrates the process of utilizing bootstrapping to estimate confidence intervals around the mean interruption cost value for a given set of ICE user inputs.

[30] Sampling with replacement allows respondents to be sampled more than once within a given sample.

Figure H1. Process for Estimating Confidence Intervals with Bootstrapping

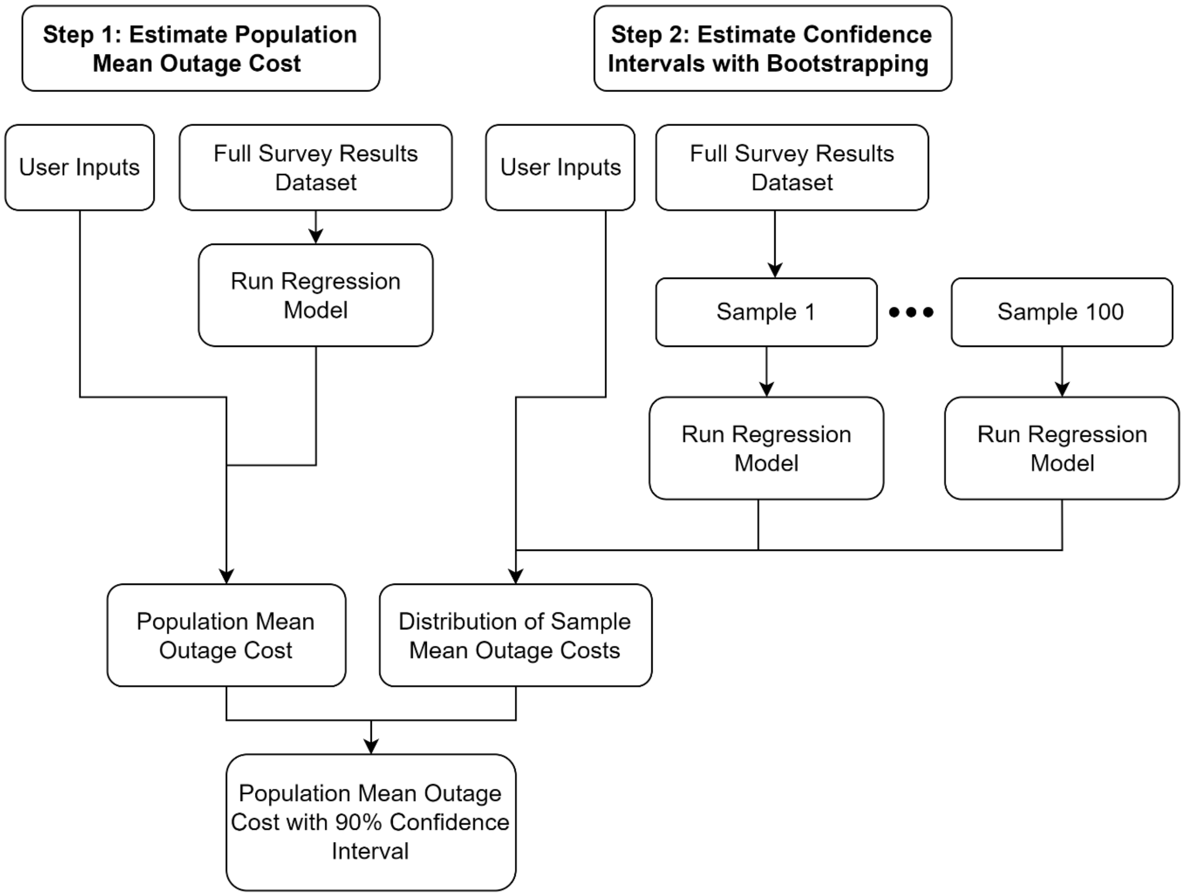

Bootstrapping allows for the estimation of precision around the mean interruption cost for any set of ICE user inputs, or any combination of customer- and interruption-level characteristics. Bootstrapping is performed at the cost-per-interruption-event level.

The mean interruption cost, combined with the upper and lower bounds of the confidence interval, allow ICE Calculator users to make inferences about both the magnitude and level of certainty surrounding the displayed average interruption cost.

Figure H2 presents an example of the residential interruption costs with the 90% confidence interval calculated by bootstrapping.

Figure H2. Example Residential Interruption Costs with Bootstrapped 90% Confidence Intervals

Figure H3 presents an example of the non-residential interruption costs with the 90% confidence interval calculated by bootstrapping.

Figure H3. Example Non-residential Interruption Costs with Bootstrapped 90% Confidence Intervals

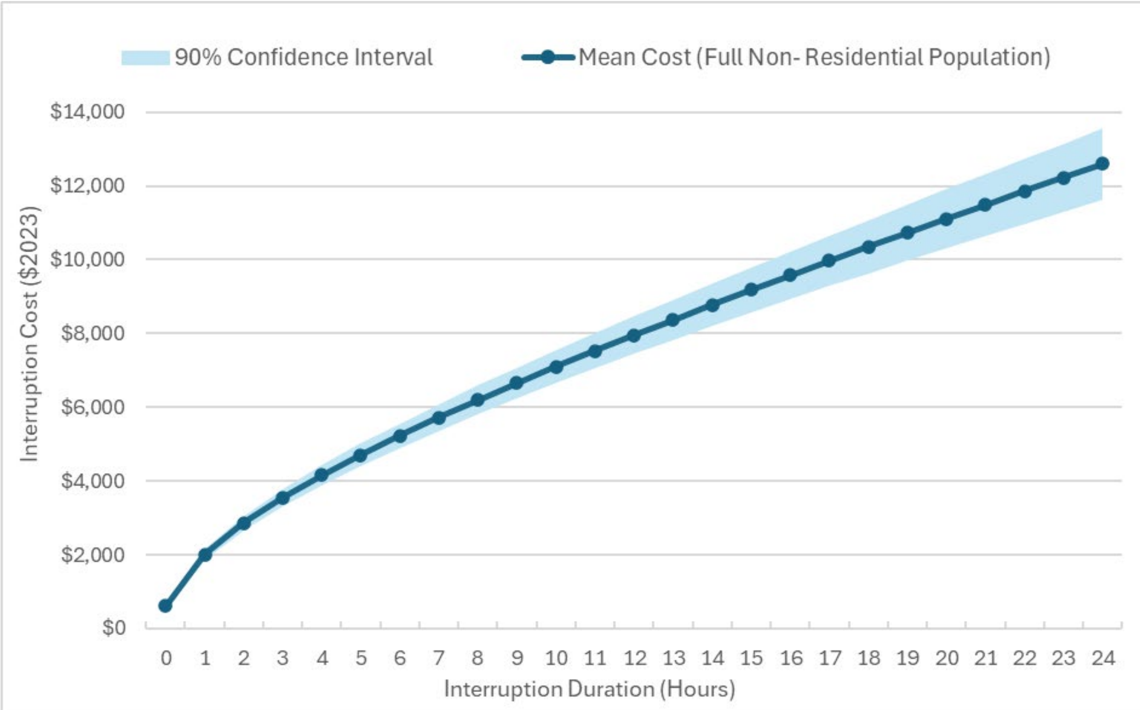

Appendix I: Evaluation of Separate Non-residential CDFs Segmented by Customer Size

The original ICE Calculator relied on two separate sets of CDFs: one for small non-residential customers (<50 MWh annual usage) and a second for medium and large non-residential customers (>50 MWh annual usage). This Appendix describes the development of CDFs for two different alternative groupings of non-residential customers and evaluates their explanatory power compared to that of the final non-residential CDF presented in Section 4.

The first alternative involved estimating distinct two-part CDFs using responses from each of three groupings of customers according to annual usage: small (0 to 50 MWh), medium (50 to 1,000 MWh), and large (over 1,000 MWh). The three CDFs were then combined piecewise to form a single, segmented CDF. The performance of the segmented CDF and the continuous non-residential CDF presented in Section 4 is evaluated by calculating the RMSE for the full dataset as well as that for customers with small, medium, and large annual usage.

Table I1 shows that the segmented CDF offers limited improvements over the continuous non-residential CDF presented in Section 4. The continuous CDF has a lower RMSE for both the full dataset and for the large customer group. The segmented CDF has a lower RMSE for the small and medium customer group. We concluded that these improvements were not sufficient to warrant the additional model complexity, discontinuities, and implementation challenges that would result from adopting a segmented CDF.

Table I1. Non-residential Continuous vs Three-segmented CDF Comparison

Model	Full Data RMSE	Small RMSE (0–50 MWh)	Medium RMSE (50–1,000 MWh)	Large RMSE (1,000 MWh and above)
Continuous CDF	302,726	5,920	34,518	575,174
Segmented CDF	310,215	5,782	34,506	589,981

The second alternative involved estimating distinct two-part models using responses from two groupings of customers according to annual usage: small (0 to 50 MWh) and medium to large (50 MWh and above). This alternative sought to account for the small sample size of the large non-residential respondents and provide a more direct comparison to the non-residential CDFs developed for the original ICE Calculator.

Table I2 shows that the segmented CDF offers limited improvements over the continuous CDF. The continuous CDF has a smaller error for the full dataset as well as for the medium-to-large customer group. The segmented CDF had a smaller error for only the small customer group. We concluded that the continuous CDF was more appropriate for implementation in the ICE Calculator update than the two-part segmented CDF.

Table I2. Non-residential Continuous vs Two-segmented CDF Comparison

Model	Full Data RMSE	Small RMSE (0–50 MWh)	Medium to Large RMSE (50 MWh and above)
Continuous CDF	301,649	5,918	339,521
Segmented CDF	303,028	5,817	341,073