Chapter 3

THE COST OF POWER OUTAGES TO HETEROGENEOUS HOUSEHOLDS

An Application of the Mixed Gamma-Lognormal Distribution

David F. Layton
Daniel J. Evans School of Public Affairs
University of Washington
dflayton@u.washington.edu

Klaus Moeltner
Department of Resource Economics
University of Nevada, Reno
moeltner@cabnr.unr.edu

Abstract We use a repeated dichotomous choice contingent valuation survey to elicit households' willingness to pay to avoid unannounced interruptions in electricity service. The data pose multiple econometric challenges including: correlated responses for a given household, heteroskedastic errors, and a willingness to pay distribution with large mass near zero. We address these issues by combining a gamma distribution for outage costs with a lognormally distributed scale parameter defined as a function of household characteristics, outage attributes, outage history, and random coefficients. The model is estimated through simulated maximum likelihood. We demonstrate that cost estimates are sensitive to the interaction of attributes of previously experienced and hypothetical interruptions.

Keywords: Power outage costs, non-market valuation, gamma distribution, random coefficients, maximum simulated likelihood.

1. Introduction

Sudden power outages can greatly disrupt social and economic activities of residents and firms in affected areas. Historically, the vast majority of interruptions in electricity supply in the U.S. occurred due to damages to the distribution network of a utility, usually during inclement weather conditions. Outages related to generation or transmission failures have been less common, as regulators required vertically integrated utilities to maintain a generating margin in electricity production and reliability reserves in their transmission system. However, with the advent of deregulation of the electric power sector in many states during the last decade, service interruptions caused by disparities between supply and demand for electricity have become more frequent. As became evident during the 2000/2001 power crisis in California, newly implemented deregulation schemes may lack the right incentive structure for generators to maintain or expand capacity reserves. This shortcoming, in combination with congested transmission grids and rigidities in retail pricing can result in market failures, supply shortages, and—ultimately—widespread blackouts (Faruqui *et al.*, 2001; Joskow, 2001; Borenstein, 2001). Since most existing deregulation frameworks are still in their infancy and thus may be susceptible to similar design problems, the risk of generation and transmission type outages is likely to remain higher in deregulated states compared to regulated markets in the near future.

Regardless of the specific cause of a service interruption and the regulation status of the utilities involved, the development of efficient policies to reduce the risk of blackouts requires knowledge of the economic costs they cause to customers. In a traditional power market regulators can induce utilities to step up reliability efforts (for example by replacing distribution lines with underground connections) by allowing them to recoup their increased cost of service through higher electricity rates. Such rate changes will only be acceptable to end-users if they are proportionate to their value of improved service reliability. In restructured markets, power distributors (retailers) are separate entities and generally remain under some degree of regulation. Many states are considering incentive frameworks that link retailers' rate of return to actual service performance (Energy Information Administration, 1998). Outage prevention and service restoration times are commonly used yardsticks in such performance-based contracts (Warwick, 2000). Naturally, an understanding of costs incurred by customers during service interruptions will be vital to set performance criteria and design economically efficient incentive structures in this context. Markets for generation and transmission in process of deregulation generally rely on an independent system operator (ISO) to ensure acceptable reliability levels. These control units usually apply a mix of economic and administrative tools to maintain adequate generation and transmission levels, such as transmission tariffs, defaulting fines, and regulations governing capacity markets. Again, economic efficiency dictates that the costs of system failure, i.e. service interruptions, ought to enter into the design of these instruments.

The focus of this study is on outage costs to residential customers. This segment comprises over 85% of retail customers in the U.S. and contributes more to retail sales and revenues than any other user group (35% and 43%, respectively; Energy Information Administration, 1998). Also, households are subjected to the highest risk of interruptions, as they rely on a more extensive infrastructure of distribution lines, substations, and transformers than larger commercial and industrial users (Warwick, 2000). We propose an innovative survey and estimation framework to elicit costs to households associated with specific power outages. Our model allows for the inclusion of both household characteristics and outage features, while capturing unobserved heterogeneity in household preferences for service reliability. In addition, we extend the existing outage cost literature by explicitly analyzing the joint effect of experienced and hypothetical interruptions on welfare losses. Due to the presence of high-dimensional integrals, the estimation of this model requires the application of simulated maximum likelihood techniques.

In the next section we provide a brief discussion of previously used approaches to estimating residential outage costs and motivate our new approach based on repeated dichotomous choice valuation questions. In section we develop an econometric model designed for the repeated dichotomous choice data. In section we discusses the data, estimation results, and policy implications. Section concludes the paper.

2. Approaches to Estimating Residential Outage Costs

Despite the importance of this topic in a time of continued growth in electricity reliance (Energy Information Administration, 2001), there exist only a few studies on outage costs to residential customers in the published literature. Three general methodologies to derive cost estimates have been proposed in existing work. In the first approach households are asked directly their willingness to pay (*WTP*) to avoid a specific outage type (Woo *et al.*, 1991; Beenstock *et al.*, 1998). The second method, as applied in Wacker *et al.* (1985) and Doane *et al.* (1988a) is based on households' direct estimates of itemized costs associated with a given menu of mitigating actions during a service interruption. The third methodology is anchored in a discrete choice framework where households are asked to select one or rank all of several outage scenario/payment options. Examples are Goett *et al.* (1988), Doane *et al.* (1988b) and Beenstock *et al.* (1998).

The first two methods lead to regression models with continuous dependent variables (cost or *WTP* in dollars) that are either estimated through simple OLS (Wacker *et al.*, 1985, Woo *et al.*, 1991), Tobit (Beenstock *et al.*, 1998), or a 2-Stage Heckman model (Doane *et al.*, 1988b). The latter two estimation techniques take account of the fact that some residents report zero costs or *WTP* for a given power interruption. Studies following the third elicitation approach apply variants of the conditional logit model (McFadden, 1974) to generate cost estimates.

Each approach has its benefits and drawbacks. Models based on open-ended *WTP* reports, while computationally convenient, are susceptible to strategic response and non-response bias (Arrow *et al.*, 1993; McFadden, 1994; Beenstock *et al.*, 1998). Asking respondents to itemize costs may mitigate these problems to some extent. However, as such cost menus only capture outlays for a limited number of actual market transactions (purchase of candles and batteries, dining out, etc.), there is a risk of missing non-market welfare losses associated with blackouts, such as health and safety concerns, disruption of work or study, and interference with social events and past-time. Thus, cost estimates from such lists can only be interpreted as lower bounds for actual welfare losses, assuming truthful responses to each line item.

Discrete choice elicitation methods based on conditional logit analysis, in turn, have the theoretical ability to capture both market and non-market values associated with specific outage types. In the residential outage costs literature there is some evidence that such multi-choice models may trigger 'status quo' bias (i.e. a household's inherent resistance to any changes in service provision) and asymmetry effects (i.e. the value of service deteriorations categorically exceeds the value of service improvements), as shown in Doane *et al.* (1988b), Hartman *et al.* (1991), and Beenstock *et al.* (1998).[1]

In this study, we promote the use of repeated dichotomous choice questions to elicit the cost of power outages to residential customers. The dichotomous choice, or referendum-style, format has been found to provide a more familiar decision making context to respondents, and to largely avoid creating strategic response incentives (Arrow *et al.*, 1993; Hanemann, 1994). In our application, each respondent is presented with a series of hypothetical outage scenarios, differing in length and time of occurrence. For each scenario, households have to decide if they would be willing to pay a given amount to avoid a specific interruption, or tolerate the outage with no change in electricity costs. This format collects a large amount of information from each respondent and allows

[1]Furthermore, unlike in most of the studies that have used multi-choice conjoint experiment-designs for their valuation questions, the policy objective for this study was NOT an examination of the potential for price-differentiated service packages. Thus there was no need to present more than two relevant "choices". In addition, as the power interruptions we model in this study are by definition unannounced random events, it would appear counterintuitive to ask a given respondent to "choose" from a set of different outages.

for the valuation of attributes that describe outages. As the econometric model provides household-level estimates of willingness-to-pay-to-avoid outages, or, equivalently, outage cost, as a function of outage and household characteristics, it allows for forecasts of costs to residential customers of new types of outages not explicitly included in the experimental design.

There are a number of econometric challenges associated with the repeated dichotomous choice approach. In the next section we develop an econometric model appropriate for using repeated dichotomous choice data for estimating the costs of residential power outages.

3. The Econometric Model

The econometric model is designed to handle four important features in the data. First, many households may have a near-zero cost attributable to a power outage, especially outages of short or momentary length. Second, it can reasonably be assumed that no household obtains positive value from a power outage. Considering these two points together suggests a distribution of *WTP* that is non-negative but can allow for substantial mass near zero.[2] Third, the survey data consists of up to four responses per respondent, and so we anticipate that intra-respondent *WTP* responses will exhibit correlation due to common unobservables for a given household. Fourth, the interaction of household unobservables with outage attributes will likely lead to heteroskedastic errors. As is well known, ignoring such heteroskedasticity in a discrete choice context will lead to inconsistent parameter estimates, especially if elements of the error variance are correlated with regressors (e.g. Hanemann and Kanninen, 1999).

The model we develop incorporates these four features. We follow Cameron (1988) and specify directly a probability density function (pdf) for latent individual *WTP*. Specifically, we choose a gamma kernel for the distribution of *WTP*. The gamma distribution for dichotomous choice data has been previously considered by McFadden (1994) and Werner (1999). It constrains *WTP* to be non-negative, but is also flexible enough to allow for exponential or normal like behavior, with much of the mass near or far away from zero as implied by the data.[3] As in McFadden (1994) and Werner (1999) we express the scale

[2]See Haab and McConnell (1998) for a detailed discussion of willingness to pay distributions.

[3]McFadden (1994) and Werner (1999) estimate models with a discrete-continuous distribution for *WTP*. The population is modeled as having two components, those with a zero *WTP*, and those with a positive *WTP*. There is a discrete probability of a respondent having a zero *WTP*. The positive component of the *WTP* distribution is modeled using a continuous distribution. This model has come to be called a "spike" model, given the discrete spike at zero in the *WTP* distribution modeled with an extra parameter. See Hanemann and Kanninen (1999) for a more detailed discussion. This approach has not been applied to repeated dichotomous choice valuation with multiple scenarios. The spike model is difficult to generalize to multiple response valuation for a number of reasons. First, the spikes for each scenario are unlikely to be equal requiring additional parameters, or need to be modelled as functions of the attributes of the scenario, along with the continuous portion of the *WTP* distribution which may make identification in

parameter of the gamma distribution as an exponential function of explanatory variables. We then extend this specification by modeling some of the coefficients associated with the regressors as random parameters. In addition to introducing the desired intra-household correlation across choice occasions and heteroskedasticity, this specification allows for an explicit analysis of the interactive effect of various outage attributes on *WTP*. We describe below the econometric model and its estimation via simulated maximum likelihood.

First we consider the *WTP* model without unobserved preference heterogeneity. Respondents are presented with a question that asks whether they would be willing to pay $\$B$ (B stands for bid) to prevent a power outage. Each power outage is described by a set of characteristics, as is each respondent. Denoting the vector of power outage and household characteristics by *x*, we model each respondent i's *WTP* for a given outage, j , as a function $WTP_{ij}(x_{ij}, \theta)$ where θ is a set of parameters. Bids can vary across outage types and respondents, and respondents will answer that they will pay B_{ij} if B_{ij} is less than their WTP_{ij}. Formally, we follow standard practice, (see Cameron and James 1987) and denote a "yes" response by $Y_{ij} = 1$, and a "no" response by $Y_{ij} = 0$. Then:

$$Y_{ij} = \begin{cases} 1, & \text{if } B_{ij} < WTP_{ij}(x_{ij}, \theta) \\ 0, & \text{if } B_{ij} > WTP_{ij}(x_{ij}, \theta) \end{cases} \tag{3.1}$$

We take a fully parametric approach. To ensure a positive *WTP*, one can either formulate the model in terms of a transformed Bid variable, usually by taking logs, and assuming that an additive error term is from a normal or logistic distribution (for example, Cameron and James, 1987). The alternative approach we follow is similar to the suggestions outlined in Haab and McConnell (1998). Specifically, we consider equation (3.1) as written with the bid, B_{ij}, in levels, but utilize a distribution for WTP_{ij} that takes only non-negative values. We begin by assuming that the WTP_{ij} are independently distributed across households and outages as $Gamma(b_{ij}, c)$, so that the density function $f(WTP_{ij})$ is:

practice difficult. Second, one would expect substantial correlation across scenarios as those who have a zero *WTP* for one scenario are far more likely to have a zero *WTP* in another scenario. This would necessitate the use of a multivariate distribution over the "spikes" which is difficult to implement. On an intuitive level, it is not unreasonable in our context to assume that every household experiences at least incremental costs or disutility from even the shortest outage. In both studies mentioned above, the item to be valued are wilderness areas or wildlife sanctuaries. In those applications, it seems more likely that some respondents' *WTP* is truly nonpositive, and that this subpopulation is distinctly different from other stakeholders.

$$f(WTP_{ij}) = \frac{\left(WTP_{ij}b_{ij}^{-1}\right)^{c-1}\left[\exp\left(-WTP_{ij}b_{ij}^{-1}\right)\right]}{b_{ij}\times\Gamma(c)} \tag{3.2}$$

for $0 \leq WTP_{ij} < \infty$, with $b_{ij}, c > 0$, and $\Gamma(c)$ is the gamma function evaluated at c (e.g. Evans *et al.*, 2000). Following McFadden (1994) and Werner (1999) we model the scale parameter, b_{ij}, as an exponential function of a linear combination of explanatory variables x_{ij} and associated coefficient vector θ_i, i.e. $b_{ij} = \exp(x_{ij}'\theta_i)$. This ensures that $b_{ij} > 0$, as required. As indicated above, to capture household heterogeneity and to introduce correlation across intra-household responses we model these coefficients as stochastic terms. Specifically, we let θ_i follow a multivariate normal distribution with mean vector μ and variance-covariance matrix Ω, i.e. $\theta_i \sim mvn(\mu, \Omega)$. This specification allows the elements of θ_i to – a priori – have unrestricted sign and magnitude. As illustrated in Moeltner and Layton (2002) the estimated covariance terms of Ω can provide additional information on joint effects of different regressors on the dependent variable.

The assumption of multivariate normality for θ_i implies a lognormal-gamma mixture distribution for WTP_{ij}. Note that in contrast to b_{ij} we treat the shape parameter c as common to all respondent-scenario combinations. The mean for this distribution can be conveniently expressed as cb_{ij} (Evans *et al.*, 2000). Note that expressing the scale parameter, b_{ij}, as an exponential of a linear function of covariates and outage attributes will make expected outage costs for a given respondent a non-linear function of all of the covariates and outage attributes interacted together. We will graphically illustrate the rich results this feature of the model yields when we discuss our application.

Since each household responded to up to four outage scenario/bid combinations, computation of the joint probability of observing a specific choice sequence for a given respondent requires computing a multi-dimensional integral. We approximate these probabilities using the random parameter simulator as described in Revelt and Train (1998), McFadden and Train (2000), and Layton and Brown (2000). To simulate the probability of each respondent's set of responses, we first compute the gamma probability conditional on the θ_i, then we simulate the unconditional probability using draws from θ_i's multivariate normal distribution. Considering (3.1), conditional on θ_i, the probability that a respondent says "yes" to a particular valuation question is $1 - F(B_{ij})$, and the probability of a "no" is $F(B_{ij})$ where:

$$F(B_{ij}) = \int_0^{B_{ij}} f(WTP_{ij})dWTP_{ij},$$

and $f(.)$ is the gamma pdf shown in (3.2). $F(.)$ is not closed form, but is readily computed in standard statistical packages. Denote the appropriate "yes" or "no" conditional probability for a particular valuation scenario by $P_{ij}|\theta_i$. Under our assumptions, the $P_{ij}|\theta_i$ are statistically independent for person i, across all j. Thus the probability, $P_i|\theta_i$, of a series of responses conditional on θ_i is

$$P_i|\theta_i = \prod_{j=1}^{j=m} P_{ij}|\theta_i, \tag{3.3}$$

where m indexes the number of WTP questions. The unconditional probability for person i, is:

$$P_i = \int_{-\infty}^{\infty} \left(\prod_{j=1}^{j=m} P_{ij}|\theta_i \right) f(\theta_i) d\theta_i, \tag{3.4}$$

where the dimension of the integration is equal to the number of random parameters in θ_i. Simulation of P_i is straightforward following Brownstone and Train (1999) and McFadden and Train (2000). At each iteration of the maximum likelihood routine we draw R sets of θ_i as $MVN(\mu, \Omega)$, and compute the simulated P_i, $\tilde{P}_i$, as the average over the R draws:

$$\tilde{P}_i = \frac{1}{R} \sum_{r=1}^{r=R} \left(\prod_{j=1}^{j=m} P_{ij}|\theta_{ir} \right). \tag{3.5}$$

The elements of μ and Ω are updated throughout the optimization process.

4. Empirical Analysis

4.1 Data

The data are from a fall 1998 survey of residential customers implemented by a U.S. utility. The main objective of the survey was to identify priority neighborhoods for reliability improvements in power distribution based on the *WTP* to avoid an outage. Each household was presented with four outage scenarios. For each scenario, households could avoid the outage by use of a pre-installed backup generator for which they would pay a specific fee every time the generator was activated by a power interruption. The selection of scenarios was based on Sullivan *et al.* (1996) and was subjected to further pre-testing using focus groups. Each scenario differed in terms of season (summer versus winter), outage timing and duration, and corresponding bid amounts. The timing and duration were chosen in consultation with the utility. Given the duration of the outages, bid amounts were based on the open-ended *WTP* data from Sullivan *et al.* (1996) in conjunction with results of the focus groups.

The following additional considerations guided the experimental design. First, the bids are such that for a given respondent a one hour outage never costs less than a momentary outage, a four hour outage never costs less than a one hour outage, and so on. Given that the open-ended data previously collected by Sullivan *et al.* (1996) for another utility revealed a fairly long right tail in the distribution of *WTP*, we allowed for a number of fairly high bid levels to be able to adequately model the skewed distribution. A wide range of bids was used to account for the fact that the location (mean, median) of the *WTP* distribution for the consumers in question might be significantly higher or lower than in the previously available data from another utility. Finally, the survey versions were carefully designed to avoid any implicit ordering of the bids through the four scenarios.

The mail survey yielded a 63% response rate. After elimination of protest responses and observations with missing household characteristics, 4,528 observations from 1,421 households were retained for this analysis. Seven of the eight administered scenarios were for winter time outages which is our focus here. Table 3.1 summarizes the seven scenarios we utilize, bid ranges, and sample counts. Household characteristics were collected as part of the survey and are supplemented with information available from customer accounts. Table describes a set of variables that relate to the types of electricity needs a household may have, which we utilize in our model estimation.

Table 3.1. Scenario and bid design

Scenario	*Duration (hrs)*	*Time*	*Bid Levels*		*No. of Obs.*
			lowest	highest	
1	1	7 pm	0.5	30	652
2	4	7 pm	1.0	50	642
3	1	8 am	0.5	40	656
4	Moment (1-2 sec.)	7 pm	0.5	30	654
5	1	midnight	0.5	40	665
6	12	7 pm	15.0	100	623
7	1	3 pm	0.5	40	636
				Total:	4,528

Note: All outages occur on a winter weekday and are unannounced.

4.2 Model Estimation

Our model specification includes a dummy variable for evening outages (*evening*), the log of the duration of a given outage scenario in minutes (

ln_dur), and the household attributes listed in Table 3.2. In various combinations, most of these variables have been considered as determinants of residential outage costs in existing studies (e.g. Doane *et al.*, 1988a; Beenstock *et al.*, 1998). We add to this traditional set of regressors a dummy variable for mobile home residences and a dummy variable taking the value of one if a given household has access to non-electric power sources for heating. The last two variables in Table 3.2, the number and log of total duration of outages during the preceding 12 months, are included in the model to measure the effect of outage history on *WTP* (or cost) estimates. While other studies have captured the impact of past outage occurrences on households' *WTP* to avoid future interruptions (Doane *et al.*, 1988b; Hartman *et al.*, 1991; Beenstock *et al.*, 1998) the separate inclusion of historic outage counts and combined duration appears to be novel. As we will show, these two indicators have significant and offsetting effects on cost estimates.

Table 3.2. Household Characteristics

Variable	*Description*	*Mean*	*Std. Dev.*
generate	1 = home has generator	0.17	0.37
business	1 = business at home	0.13	0.34
medical	1 = medical need at home	0.03	0.17
home	1 = someone at home most of the time	0.60	0.49
hh_ size	household size (persons)	2.60	1.45
over64	number of persons over 64	0.36	0.70
inc000	annual income, $1000	53.02	27.83
mobile	1 = mobile home	0.10	0.29
other_ heat	1 = secondary heating source available	0.49	0.50
ln_ cons	log of avg. monthly electricity consumption in kwh	6.75	0.71
num_ out	number of outages in past 12 months	5.58	8.36
out_ past	log of total duration of outages in past 12 months (hours)	1.17	1.87

We model the outages in the dichotomous choice scenarios as consisting of two components: A short momentary component of less than a minute, and then any additional duration beyond one minute. Momentary outages that last for less than a minute have a particular suite of impacts on some households but not on others. Sensitive electrical equipment such as medical devices and home office systems may fail, but most other home electricity uses will not be greatly affected. Longer outages share this initial effect and as duration increases other costs begin to mount. The literature suggests that the impact of outage duration increases, but at a decreasing rate, so we model the effect of duration beyond the first minute in log form. Following Moeltner and Layton (2002) we include an intercept term in our model while setting the value of log duration for a momentary (1 minute or less) outage scenario to zero. Thus the intercept term is the effect of a momentary outage on *WTP* (*moment*), and the coefficient on log duration (*ln_dur*) measures the impact of a duration length longer than a minute.

Specifying all k =15 elements of θ_i as correlated random coefficients would require estimation of k elements of μ plus $k(k+1)/2$ elements of Ω for a total of 135 parameters. Such a large number of parameters are not likely to be identified without a prohibitively large data set. Further, estimation is not computationally feasible given the need to simulate the response probability at each function evaluation (Keane, 1997). We thus restrict randomness to variables of primary interest with likely heterogeneity in preferences. These are past outage duration (*out_past*) and occurrence (*num_out*), as well as the two main attributes of the hypothetical outage scenarios, *ln_dur* and *moment*.[4] Adding the resulting ten variance-covariance terms in Ω and the gamma shape parameter c to the 15 elements of θ_i yields a total number of 26 model parameters. We estimate this model through simulated maximum likelihood using R =1,000 repetitions for the simulated probabilities described in (3.5).

Table 3.3 summarizes the estimation results from the mixed Gamma-Lognormal model. Generally, the model exhibits a reasonably good fit with the underlying data with a pseudo-R^2 of 0.25. The majority of the coefficient estimates are significant at the 5% level or higher. Specifically, the gamma shape parameter, c, is estimated with high precision. A value of c less than one indicates a high probability mass near zero (Evans *et al.*, 2000). This result is compatible with similar findings by Doane *et al.*, (1988a) and Beenstock *et al.* (1998), who report a preponderance of zeros in their open-ended *WTP* elicitations even after purging their data of potential strategic and protest responses. Evidently, a large share of residential customers in our sample does not consider the bulk of the power outages described as especially bothersome. We discuss in turn the results for the outage attributes, household characteristics, and past outage history before elaborating on the covariance estimates. We conclude this section with a comparison of our results to those available in the literature.

4.3 Outage Attributes

The effect of evening outages emerges as insignificant compared to a combined baseline of afternoon and morning interruptions. A possible explanation for this finding may be that the period of daylight during winter is relatively short in the survey region. Accordingly, electricity needs for lighting are re-

[4] The evening dummy is specified as a fixed coefficient in part on the basis of preliminary work, which suggested it had little impact in our data set.

Table 3.3. Estimation Results

Parameters	*Coeff.*	*Stand. err.*	
c	0.232	(0.023)	***
evening	–0.015	(0.132)	
generate	–0.494	(0.174)	***
business	0.338	(0.200)	*
medical	0.304	(0.407)	
home	0.658	(0.148)	***
hh_ size	–0.124	(0.048)	**
over64	–0.215	(0.109)	**
inc000	0.026	(0.003)	***
mobile	0.457	(0.222)	**
other_ heat	–0.467	(0.134)	***
ln_ cons	0.133	(0.101)	
ln_ dur	0.455	(0.056)	***
num_ out	–0.033	(0.011)	***
out_ past	0.260	(0.066)	***
moment	–0.933	(0.808)	
Variance and Covariance Terms			
ln_ dur	0.077	(0.062)	
ln_ dur / num_ out	0.001	(0.006)	
num_ out	0.000	(0.000)	a
ln_ dur / out_ past	0.138	(0.069)	**
num_ out / out_ past	0.002	(0.010)	
out_ past	0.257	(0.154)	*
ln_ dur / moment	–0.899	(0.514)	*
num_ out / moment	–0.015	(0.065)	
out_ past / moment	–1.647	(0.651)	**
moment	10.648	(4.075)	***
Log-likelihood	2,356.100		
Pseudo-R^2 = 1-[–2,356.1/ln(0.5)]	0.250		

Note: Standard Errors in parentheses. a = rounded to zero;
*significant at 10% level; ** significant at 5% level;
*** significant at 1% level.

quired for much of the day. In addition, many businesses in the particular metropolitan area that generated this sample offer staggered work shifts, which distributes electricity needs more evenly over a 24 hour time period. An alternative explanation is that the limited number of outage scenarios that could be valued in the survey did not permit sufficient contrast between duration and time of day. For instance a 12 hour evening outage would cover evening, late night, and morning, thus mitigating much of the time of day effect. This suggests that when time of day is a variable of important policy interest in conjunction with outages of long duration, many survey versions will be required – perhaps prohibitively many.

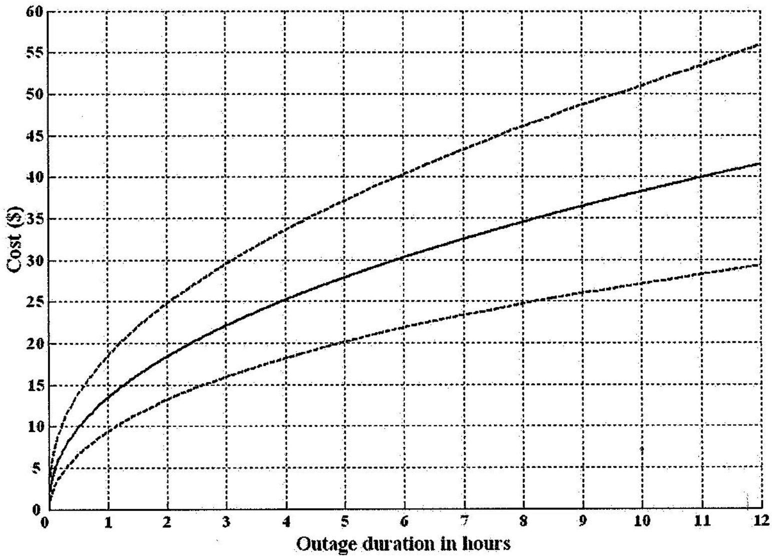

Figure 3.1. Outage Costs Versus Duration (mean and 95% confidence intervals averaged over all households). Outage costs increase at a decreasing rate with outage duration.

As reflected by the insignificant coefficient for moment, a purely instantaneous interruption does not cause any sizeable costs to the average household. This is consistent with findings reported in Caves *et al.* (1990). Outage costs and associated *WTP* values do, however, increase with the duration of an interruption as indicated by the positive sign and high level of significance for *ln_dur*. Figure 3.1 depicts the resulting duration-cost function for a prototypical household and an interruption starting at 7pm. The 95% confidence intervals are based on the empirical distribution of household-specific estimated *WTP* averaged over all respondents. Given the distributional assumptions in our model, the expectation of WTP_{ij} is itself a random variable following a lognormal distribution with mean $E(b_{ij}c) = c \times \exp\left(x'_{ij}\mu + 0.5x'_{ij}\Omega x_{ij}\right)$. Due to some outliers, this expression generates excessively high values for some households. We therefore use the median, $c \times \exp(x'_{ij}\mu)$, as the basis for our point and interval estimates of WTP_{ij}.

Figure 3.1 depicts how median costs change over outage duration. Consistent with results reported by Doane *et al.* (1988a), outage costs increase at a decreasing rate with increasing duration. This is intuitively sound as longer outages give households more time to take countervailing measures. At the same time, the variability of outage damages to individual households increases with duration as indicated by the widening spread of the confidence interval. For example, for a one-hour evening interruption our point estimate is \$13 with a 95% confidence interval of \$9 to \$18. At a duration of 12 hours the point estimate is \$42, with a 95% confidence interval of \$29 to \$56.

4.4 Household Characteristics

Turning to the effect of household characteristics, we note from Table 3.3 that the presence of medical needs and annual electricity consumption do not significantly affect *WTP* values. As expected, the presence of business activities run from home, the presence of residents at home during most of the day, and income have a positive and significant effect on cost estimates. Similarly, households residing in mobile homes have a significantly higher sensitivity to power interruptions. This is an anticipated result given the reduced insulation of such dwellings and the corresponding higher reliance of their occupants on uninterrupted heating. As expected, the availability of non-electric heating significantly reduces the *WTP* to avoid a specified interruption. The negative and significant coefficients for household size and the number of persons over age 64 are probably indicative of reduced disposable income for such families.

4.5 Past Outage History

One of the key insights provided by this analysis flows from our specification of a rich structure for past outage history. By using two components of past outage history, log duration and number of occurrences, we show that past outage history is not a uni-dimensional concept, but instead illustrate that different components of outage history have contrasting effects on *WTP*. These contrasting effects derive from the significant and opposite signs for the number of past outages during the preceding 12 months (*num_out*) and the log of combined duration of such outages in hours (*out_past*). This implies that an increase in historic outage frequency, ceteris paribus, decreases a household's *WTP* to avoid further interruptions. This could be indicative of a learning-to-cope, or preparedness effect induced by frequent outage events. In stark contrast, however, estimated *WTP* increases with the combined duration of recent interruptions. Evidently, one or more longer blackouts in the recent past stir up decidedly unpleasant memories in affected respondents and seem to induce substantially higher costs than they generate learning gains.

These results may explain the contradicting findings in the existing outage cost literature on the role of outage history on cost estimates for a specified future interruption. Specifically, Doane, Hartman *et al.* (1988b) find that households that traditionally experience a larger number of outages have a decreased *WTP* to avoid additional blackouts, while Doane *et al.* (1988a) and Beenstock *et al.* (1998) reach the exact opposite conclusion. This apparent discrepancy could be a result of different average length of past outages in each of these cases. None of these studies incorporate measures of historic outage duration in their estimation models.

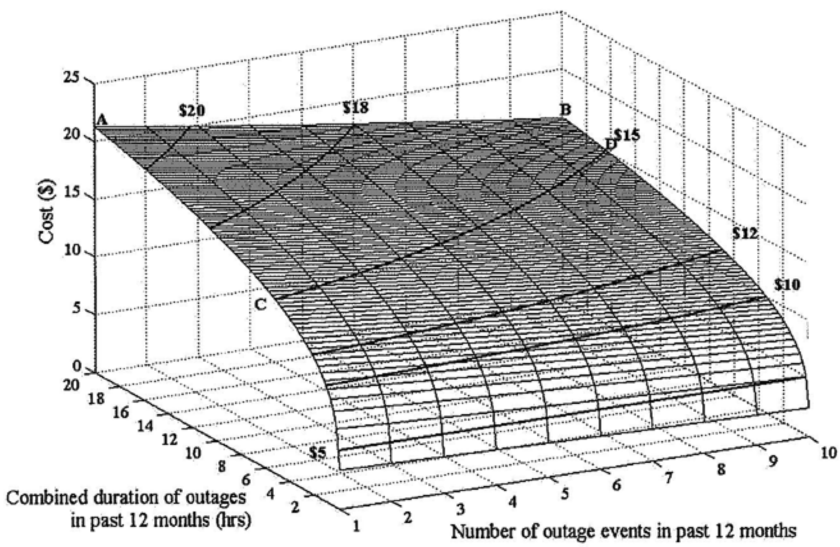

Figure 3.2. *WTP* for a One-Hour Outage: Past Duration vs. Past Events. The iso-*WTP* lines show that *WTP* to avoid this particular outage is the same for households that have experienced one long outage or several shorter interruptions. Alternatively, *WTP* is lower if a given past duration is distributed over several outage occurrences. This may be indicative of two countervailing forces: a cost-awareness factor versus a learning-to-cope effect.

Figure 3.2 illustrates these offsetting effects. The surface plane of figure 3.2 connects simulated cost estimates for a hypothesized evening outage of one-hour duration at different combinations of number and duration of past interruptions. In each case, cost estimates were first generated for each respondent and then averaged over households. For example, a prototypical household that experienced one 20-hour outage in the past would be willing to pay approximately $22 to avoid the stipulated interruption (point A). If the combined duration of 20 hours is distributed over, say, 10 individual events, *WTP* decreases to approximately \$16 (point B). The darker lines crossing the cost surface represent iso-*WTP* curves for several dollar amounts. For instance, a *WTP* value of \$15 could be reached with one 6-hour interruption (point C), ten outages with a combined duration of 16 hours (point D) or any of the frequency / duration pairs along the line C-D.

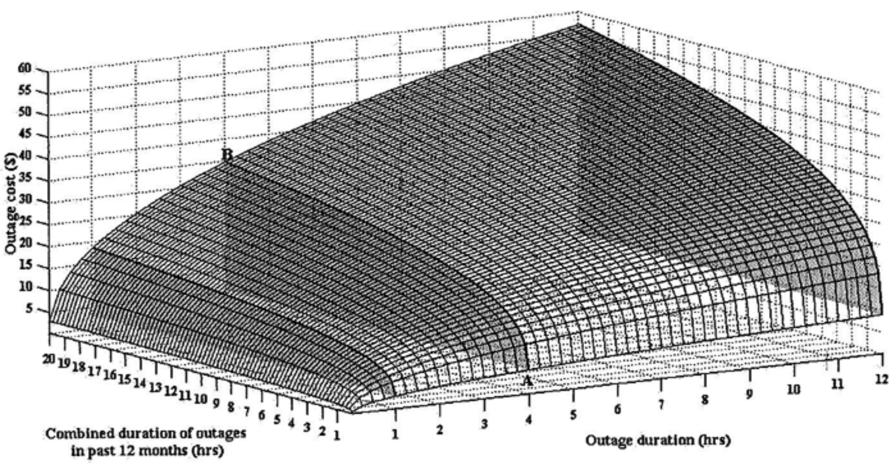

Figure 3.3. Outage Costs by Scenario Duration and Past Duration. The WTP to avoid future outages of a given duration increases with the combined duration of past outages. This effect is relatively stronger for longer hypothetical future outages.

Figure 3.3 shows the importance of capturing historic duration effects when analyzing the value of electric reliability to residential customers from a different perspective. The figure depicts cost estimates associated with a new unannounced future outage as a function of combined duration of past interruptions. For ease of interpretation, cost estimates for one, four, and 12-hour duration are highlighted through cross-section planes. For example, a four-hour evening interruption causes costs of approximately \$7 to a household that has not experienced any blackouts over the preceding 12 months (point A). In contrast, *WTP* to avoid a four-hour outage is about five times higher for a household with a combined past duration of 20 hours (point B). Comparing the vertical height of the cross-sections at any historic duration value, one can also note that combined past duration affects *WTP* estimates relatively more for longer proposed interruptions.

The recognition of these interactive effects of frequency and duration of past blackouts may offer additional guidance to utilities in identifying residential neighborhoods with relatively high sensitivity to power interruptions. Clearly, a sole focus on the number of past outage events in this context may lead to sub-optimal allocation of reliability efforts.

4.6 Covariance Parameters

The bottom half of Table 3.3 shows estimation results for the random elements of Ω. Based on a likelihood ratio test, we strongly reject the null hypothesis that all elements of Ω are zero, i.e. that outage features and outage history have an equal effect on *WTP* for all customers. The variance term for moment is relatively large and highly significant. This indicates strong heterogeneity in costs from an instantaneous interruption across individual households. Thus, even though the costs caused by momentary outages are negligible for the average household, there are some families that experience considerable damage even from a very short blackout. This mirrors the results for costs of momentary outages to commercial/industrial firms in Moeltner and Layton (2002). The remaining elements of Ω that emerge as significant at the 5% level or higher are the covariance terms for *out_past* with *ln_dur* and *moment*, respectively. The first term is positive and lends itself to an intuitively sound interpretation: households that are relatively more affected by combined past duration are also more sensitive to duration as specified in the hypothetical scenarios. The negative sign on the second term suggests that households whose *WTP* values depend more strongly on the combined duration of experienced outages are relatively less sensitive to instantaneous interruptions. Conversely, households whose *WTP* is affected less strongly by the duration of past interruptions are relatively more concerned about instantaneous blackouts. Presumably, such residents experience the bulk of outage costs during the initial period of a blackout and are relatively more capable to cope with prolonged duration. The negative sign for the covariance term between *moment* and scenario duration (*ln_dur*) supports this hypothesis even though this term is only significant at the 10% level.

4.7 Comparison with Previous Estimates

Our results are compared to those stated in other studies in Table 3.4. As recommended by Caves *et al.* (1990) and Woo and Pupp (1992) we report outage costs in both absolute dollars and in terms of dollars per kwh unserved, using additional information provided by the Utility on energy consumption by a prototypical household for the season and time period of consideration.

Table 3.4. Cross-study Comparison of Cost Estimates

	This Study	Doane *et al.* (1988)	Doane *et al.* (1988b)	Woo *et al.* (1991)
Data Year:	1998	1986	1986	1989
Timing:	winter evening	winter evening / morning	winter eve./ mor.	winter
Method:		2-stage Heckman	Self-stated	OLS
Duration		Cost (1998 $)		
1 hr	13.45	16.33	13.66	9.83
4 hrs	25.17	29.16	26.79	13.10
8 hrs	34.49	N/A	N/A	19.65
12 hrs	41.51	49.39	58.11	30.13
		Cost ($/kwh unserved)		
1 hr	5.34	14.61	N/A	12.71
4 hrs	2.66	5.29	N/A	7.34
8 hrs	2.29	N/A	N/A	4.98
12 hrs	2.06	3.38	N/A	3.28

As can be seen from the table, our cost estimates in absolute dollars are reasonably close to those found in Doane *et al.* (1988a), and the results based on self-stated costs in Doane *et al.* (1988b). The estimates by Woo *et al.* (1991) are clearly lower than those produced by this analysis and the other two comparison studies for all listed outage durations. To some extent, this may be related to the fact that Woo *et al.* (1991) use OLS regression to generate these cost estimates. This is likely to place the mean of the resulting underlying cost distribution closer to zero than would be the case in models that impose non-negativity constraints on outage costs, as applied in the other three studies. At the same time, these lower estimates may simply indicate a relatively lower reliance on electric power of the particular population captured in that analysis. When expressed in terms of dollars per kwh unserved Woo *et al.* (1991)'s cost estimates are close to the ones reported in Doane *et al.* (1988a), which would lend support to the latter hypothesis. While following the same general decrease-with-duration pattern, our cost estimates in dollars per kwh unserved are about 50% smaller in magnitude than those generated by the two comparison sources. This suggests that cumulative energy consumption by a representative household from our population during the listed outage periods is about twice as high as underlying consumption for the sample considered in Doane *et al.* (1988a), and approximately three to four times higher than for households in Woo *et al.* (1991). This relatively pronounced difference may be attributable to the different geographic location, differing sample characteristics, and changes in electricity consumption during the intervening time periods that separate the three studies.

5. Discussion and Conclusion

We have developed an econometric model that captures the essential features of repeated dichotomous choice non-market valuation data. By using a gamma distribution for the kernel of *WTP*, the model allows for the distribution of *WTP* to have large amount of mass near zero while still constraining WTP to be non-negative. This is crucial for estimating *WTP* for goods that may not be worth much to many households such as preventing a momentary electricity outage. Our model allows for heterogeneity in *WTP* by specifying the scale parameter of the gamma distribution to be lognormally distributed in the population. The lognormal distribution for the scale parameter captures both heteroskedasticity and within-subject correlation in responses to the multiple dichotomous choice valuation questions. This models important features of the data. For example, as shown by the small mean but large variance for a momentary outage, a momentary outage imposes little cost on average but it imposes large costs on *some* households.

It appears that whether deregulation of the retail electricity market continues or not, with higher electricity usage and insufficient construction of new generating capacity in some areas of the country, rational management of the risks of power outages will become more, not less important in the coming years. Rational management requires an understanding of whether the benefits of reliability improving actions outweigh the costs. Given the complexities of managing a transmission and distribution network, it is crucial that utilities or regulators be able to disaggregate their costs and benefits as much as possible so that they can effectively target projects to those that would most benefit. This is critical as in real world applications available budgets are likely to be exhausted before all beneficial projects have been implemented. Understanding how marginal *WTP* behaves as a function of the attributes of an unannounced outage and past outage history are crucial determinants of the relative benefits of mitigating outages.

Our model provides a rich analysis by using a number of covariates that are typically observable by utilities, such as the availability of non-electric heating, or whether the customer lives in a mobile home. The specification of the covariates in exponential form allows for non-linear surfaces describing the *WTP* to avoid different types of outages. Overall, the results conform to expectations based on the outage costs literature and experience. These factors can be used by a utility or a regulatory authority to better target reliability improvements to neighborhoods (or even at the circuit level) that would most benefit. For example, by conditioning *WTP* on both the number of previous outages the household has experienced and the duration of previous outages we can show that they have differential effects on welfare losses. As illustrated in figure 3.2, ceteris paribus, households that experience one long outage are willing to pay more to prevent *any* kind of future outage than households that have experienced a number of shorter outages. This illustrates another margin on which we can compare the relative benefits of different reliability improvements.