EVALUATING COSTS WITH UNMEASURED CONFOUNDING: A SENSITIVITY ANALYSIS FOR THE TREATMENT EFFECT

BY ELIZABETH A. HANDORF[1,*], JUSTIN E. BEKELMAN[2,†],
DANIEL F. HEITJAN[3,†] AND NANDITA MITRA[3,†]

Fox Chase Cancer Center * *and University of Pennsylvania*[†]

Estimates of the effects of treatment on cost from observational studies are subject to bias if there are unmeasured confounders. It is therefore advisable in practice to assess the potential magnitude of such biases. We derive a general adjustment formula for loglinear models of mean cost and explore special cases under plausible assumptions about the distribution of the unmeasured confounder. We assess the performance of the adjustment by simulation, in particular, examining robustness to a key assumption of conditional independence between the unmeasured and measured covariates given the treatment indicator. We apply our method to SEER-Medicare cost data for a stage II/III muscle-invasive bladder cancer cohort. We evaluate the costs for radical cystectomy *vs.* combined radiation/chemotherapy, and find that the significance of the treatment effect is sensitive to plausible unmeasured Bernoulli, Poisson and Gamma confounders.

1. Introduction. Payers and health care providers, such as Medicare, private insurers and hospitals, routinely collect data on medical expenditures [Hornberger and Wrone (1997)]. Some of these data sources are readily available and have high external validity [Black (1996)] and therefore are widely used to compare treatment costs. Like all observational data, however, they are subject to confounding. That is, if all relevant confounders are measured and present in the data set, we can find a consistent estimate of the treatment effect by using a correctly specified statistical model for cost. If some confounders are unmeasured or unavailable, the model-based estimate is potentially biased. Here we develop a method to assess the sensitivity to potential unmeasured confounders when evaluating the effect of a medical intervention on mean cost.

Cost data are nonnegative and often highly skewed, features that can be readily described statistically in the generalized linear model framework. For example, a Gamma regression with a log link is often a suitable model for cost data [Dodd et al. (2006)]. The situation is more complex when some costs are censored.

Costs accrue over time, and they may accrue at different rates among individuals. Therefore, cost at the time of censoring is generally correlated with cost at the time of ultimate failure, making censoring informative [Etzioni et al. (1999), Lin et al. (1997)]. One popular method which accounts for this is Inverse Probability Weighting (IPW), where observed costs are weighted by the inverse probability of censoring [Lin (2000, 2003), Bang and Tsiatis (2000, 2002), Tian and Huang (2007), Willan et al. (2002)]. In the context of observational cost data, the method of Lin (2003) is particularly useful. He applies IPW to generalized linear models, allowing adjustment for confounders in a Gamma regression. Other approaches which handle informative censoring are based on Bayesian methods or use multi-part models [Basu and Manning (2010), Heitjan, Kim and Li (2004), Liu (2009), Liu, Huang and O'Quigley (2008)].

The methods cited above do not address the potential bias due to unmeasured confounding. A natural approach to evaluate this bias is a sensitivity analysis, in which one posits models for the distribution of an unmeasured confounder and its effects on cost, then evaluates their effects on estimates of the treatment effect. Researchers have developed various methods implementing this idea [Axelson and Steenland (1988), Flanders and Khoury (1990), Gail, Wacholder and Lubin (1988), Hosman, Hansen and Holland (2010), McCandless, Gustafson and Levy (2007), Rosenbaum and Rubin (1983a), Schlesselman (1978), Yanagawa (1984)]. Further developments include the method of Lin, Psaty and Kronmal (1998), in which the authors present a formula to determine the magnitude of the bias due to an unmeasured Bernoulli or Normal covariate. Mitra and Heitjan (2007) extended this idea to assess the sensitivity of survival outcomes using a Weibull model.

As we have indicated, the method of Lin, Psaty and Kronmal (1998) assesses the sensitivity of the treatment effect to unknown binomial or normal confounders. Wang and Krieger (2006) have shown that in a matched pairs study, conclusions are most sensitive to a Bernoulli confounder and, thus, such an analysis is sufficient to describe maximal sensitivity. It is unclear whether an analogous result holds in cohort studies, however, or indeed whether the exploration of maximal sensitivity is most desirable in every context. Therefore, in this paper we generalize the Lin method to derive corrections for a broader range of distributions of unmeasured confounders. For example, applying our formula for a Gamma confounder, one can adjust for an unmeasured skewed continuous variable such as personal income. Or using our formula for a Poisson confounder, one can adjust for an unmeasured count variable such as pack-years smoked.

A major drawback of the Lin method and related sensitivity analyses for the effects of an unmeasured confounder is the assumption of independence between measured and unmeasured confounders given treatment status. Hernán and Robins (1999) demonstrated that such an assumption is implausible, as conditioning on treatment induces a correlation even between unrelated covariates. VanderWeele (2008) showed how one of Lin's formulas can apply in certain circumstances with

relaxed assumptions, and VanderWeele and Arah (2011) developed bias adjust-ment formulas that are valid more generally without requiring conditional inde-pendence. The implementation of these approaches can be complex, however, and a simpler alternative such as the Lin method would be invaluable in applications. In this paper, therefore, we both extend the basic Lin method to a range of loglin-ear models and use simulations to explore the limits of its validity under departures from conditional independence.

We demonstrate our method with a comparison of mean medical costs for two treatments for stage II/III muscle-invasive bladder cancer. As this is the ninth costliest cancer [National Cancer Institute (2011)], it is of substantial interest to determine cost-effective treatments for it. Here, we compare the lifetime costs for radical cystectomy, an aggressive surgical procedure that is the current standard of care, to combined radiotherapy and chemotherapy, an alternative organ-sparing curative treatment. The data are censored because many registry participants were still living when the database was closed. We base our cost comparisons on ob-servational data from the linked SEER-Medicare registry, assessing the sensitivity of the treatment effect to potential Bernoulli, Poisson and Gamma unmeasured confounders.

The paper is organized as follows: in Section 2 we develop a general formula to quantify the magnitude of the bias of treatment effect due to an unknown con-founder, and derive the correction for Poisson and Gamma confounders. We fur-ther show how these formulas and those of Lin et al. [Lin, Psaty and Kronmal (1998)] can be used to assess sensitivity with cost data, including censored costs. We specifically address criticisms that have been raised of the assumption of con-ditional independence of observed and unobserved covariates given treatment. In Section 3 we assess the performance of our method using Monte Carlo simulations, evaluating robustness to departures from conditional independence. In Section 4 we apply our method to the SEER bladder cancer data.

2. Bias with an unmeasured confounder. We first consider uncensored costs. Assuming a generalized linear model (GLM) with a log link, we define the true model as

$$E(Y|X, Z, U) = \exp(\alpha + \beta_0 X + \phi_0 U + \psi_1 XU + \theta'Z),$$

where Y is the total cost, X is the treatment of interest ($X = 0$ for control and $X = 1$ for treated), U is the unmeasured confounder and Z is a vector of measured covariates. The true regression parameters are α, β_0, θ, ϕ_0 and ψ_1. We allow the unmeasured confounder U to have different effects on cost in the treatment groups, so β_0 is the effect of X when $U = 0$, and $\beta_0 + \psi_1$ is the effect of X when U is nonzero. Likewise, ϕ_0 is the effect of U in the control group, and $\phi_0 + \psi_1$ is the effect of U in the treated group. For convenience, we express the model as

$$(1) \qquad E(Y|X, Z, U) = \exp(\alpha + \beta_0 X + \gamma_X U + \theta'Z).$$

This is equivalent to the model specified above, where γ_0 is the effect of U in the control group and γ_1 is the effect of U in the treated group. For simplicity, we henceforth suppress the subscript in β_0.

In practice, we cannot fit model (1) because we do not know the value of U. We define the reduced model, which we can fit, as

$$(2) \qquad \mathrm{E}(Y|X, Z) = \exp(\alpha^* + \beta^* X + \theta^{*\prime} Z).$$

Fitting model (2) yields an estimate of the apparent treatment effect β^*.

By conditional expectation,

$$\mathrm{E}(Y|X, Z) = \int_{-\infty}^{\infty} \mathrm{E}(Y|X, Z, U)\, dF(U|X, Z)$$

$$= \exp(\alpha + \beta X + \theta' Z) \int_{-\infty}^{\infty} \exp(\gamma_X U)\, dF(U|X, Z)$$

$$= \exp(\alpha + \beta X + \theta' Z) \mathrm{M}_{U|X,Z}(\gamma_X),$$

where $\mathrm{M}_{U|X,Z}(\gamma_X)$ is the moment generating function (m.g.f.) of $U|X, Z$. If we assume that U is conditionally independent of Z given X so that $\mathrm{M}_{U|X,Z}(\gamma_X) = \mathrm{M}_{U|X}(\gamma_X)$, the following simplification is possible:

$$(3) \qquad \begin{aligned} \mathrm{E}(Y|X, Z) = \exp\big(&\alpha + \ln(\mathrm{M}_{U|0}(\gamma_0)) \\ &+ X[\beta + \ln(\mathrm{M}_{U|1}(\gamma_1)) - \ln(\mathrm{M}_{U|0}(\gamma_0))] + \theta' Z\big). \end{aligned}$$

Setting (2) equal to (3) gives

$$(4) \qquad \beta = \beta^* - \ln(\mathrm{M}_{U|X=1}(\gamma_1)) + \ln(\mathrm{M}_{U|X=0}(\gamma_0)).$$

Thus, we can conduct a sensitivity analysis by positing distributions for $U|X$ and values for γ_X in the treated and control groups. Equation (4) holds for any distribution of U that one can characterize by a moment-generating function, a large class of distributions. In the following sections we show the solutions for β when U is distributed as Poisson or Gamma. We also present the formulas developed by Lin, Psaty and Kronmal (1998) for Bernoulli and Normal unmeasured confounders, and show how they are applicable in the cost setting.

2.1. *Poisson unmeasured confounder.* For $U \sim \mathrm{Poisson}(\lambda_{X,Z})$, we have

$$\mathrm{M}_{U|X,Z}(\gamma_X) = \exp(\lambda_{X,Z}(e^{\gamma_X} - 1)),$$

and under conditional independence ($\lambda_{X,Z} = \lambda_X$), (4) becomes

$$\beta = \beta^* - (\lambda_1 e^{\gamma_1} - \lambda_0 e^{\gamma_0} + \lambda_0 - \lambda_1),$$

where λ_0 is the mean of U in the control group and λ_1 is the mean in the treated group, β is the true treatment effect, β^* is the observed treatment effect, and γ_0 and γ_1 are the effects of U in the control and treated groups, respectively.

2.2. *Gamma unmeasured confounder.* An unobserved covariate

$$U \sim \text{Gamma}(\theta_{X,Z}, \kappa_{X,Z})$$

has m.g.f.

$$\text{M}_{U|X,Z}(\gamma_X) = (1 - \theta_{X,Z}\gamma_X)^{-\kappa_{X,Z}}.$$

Under conditional independence $(\theta_{X,Z} = \theta_X, \kappa_{X,Z} = \kappa_X)$, (4) implies

$$\beta = \beta^* - \ln \frac{(1 - \theta_0\gamma_0)^{\kappa_0}}{(1 - \theta_1\gamma_1)^{\kappa_1}},$$

where the mean and variance of U are $\kappa_0\theta_0$ and $\kappa_0\theta_0^2$ for the control group and $\kappa_1\theta_1$ and $\kappa_1\theta_1^2$ for the treated group.

2.3. *Bernoulli and Normal unmeasured confounders.* Lin, Psaty and Kronmal (1998) derived the relationship between the true and apparent treatment effects for a binary outcome with a log link when unmeasured confounders are Bernoulli or Normal. Because the derivation also applies to a continuous outcome with a log link, their results are special cases of (4). Specifically, if $U \sim \text{Bernoulli}(\pi_{X,Z})$, and U is conditionally independent of Z $(\pi_{X,Z} = \pi_X)$, then

$$\beta = \beta^* - \ln \frac{e^{\gamma_1}\pi_1 + (1 - \pi_1)}{e^{\gamma_0}\pi_0 + (1 - \pi_0)},$$

where π_0 is the prevalence of U in the control group, π_1 is the prevalence of U in the treated group, and the other parameters are as above. Similarly, if $U \sim \text{Normal}(\mu_{X,Z}, 1)$ and again is conditionally independent of Z $(\mu_{X,Z} = \mu_X)$, then

$$\beta = \beta^* + (\mu_0\gamma_0 - \mu_1\gamma_1 + (\gamma_0^2 - \gamma_1^2)/2),$$

where μ_j is the mean of U in treatment group j, $j = 0, 1$, and β, β^*, γ_0 and γ_1 are as above.

2.4. *Assumption of conditional independence.* The conditional independence assumption (that U is conditionally independent of Z given X) cannot be tested in practice since U is not observed. However, it has been well-established that the assumption cannot actually hold in practice. Conditional independence would hold if the marginal correlation perfectly cancels the conditional correlation or if the covariates Z are not truly confounders, but either condition is highly improbable. Hernán and Robins (1999) observed that conditional independence cannot hold if Z and U are marginally independent and are both confounders. VanderWeele (2008) showed that closed-form solutions are available for normal unmeasured confounders and for a binary confounder with a normal outcome under a relaxed additivity assumption. VanderWeele and Arah (2011) developed general formulas for calculating an adjusted estimate of the treatment effect without assumptions

about the relationship between the measured and unmeasured covariates. In the most general case one must specify $E(Y)$ at each level of X, U and Z. When U or Z is continuous, this would be very difficult to implement in practice. For the case where the relationship between Y and U is constant across levels of Z, as would occur in a linear model, VanderWeele and Arah (2011) developed simplified formulas that require fewer sensitivity parameters. In our loglinear regression setting, however, an additional bootstrapping approach would be required to use these formulas.

Although we do not expect our corrections to hold exactly in practice due to weak tenability of the conditional independence assumption, if conditional dependence is modest, they may still be useful. In Section 3.2 we present simulations to investigate the effect of dependence within treatment strata.

2.5. *Analysis with censored costs.* We derived our adjustment formulas assuming uncensored data; nevertheless, because they are statements about parameters, they are valid also in the censored case provided that one estimates the parameters consistently. Lin (2003) developed a generalized regression method, appropriate for censored costs, that combines the generalized estimating equation approach with inverse probability of censoring weighting. This approach produces consistent estimates but forfeits efficiency because it uses only the uncensored data. Bang and Tsiatis (2002) developed a more efficient estimator for linear regression that uses all of the data, but an extension to the GLM has yet to be developed. We therefore recommend using the Lin (2003) approach, but note that our formulas are applicable with any method that gives consistent estimates for a loglinear mean model.

2.6. *Sensitivity analysis for cost regression.* The sensitivity analysis proceeds in four steps: first, estimate β^* in (2) by any consistent method. Second, specify the type of the unmeasured confounder: for example, Poisson for a count confounder; Gamma for skewed, positive continuous; Normal for symmetric continuous; and Bernoulli for binary. Third, hypothesize possible parameter values η for the distributions of the unmeasured confounder in the treated and control groups, and also the possible effect γ of the unmeasured confounder on cost. Finally, apply the correction via (4) to obtain $\hat{\beta}$, the adjusted treatment effect.

The analysis assumes that the sensitivity parameters η are known. Because the adjustment is in every case additive, it follows that

$$\mathrm{Var}(\hat{\beta}|X, Z, \eta) = \mathrm{Var}(\hat{\beta}^*|X, Z).$$

We use this relationship to calculate confidence intervals and determine the significance of the adjusted treatment effect, as was done by Lin, Psaty and Kronmal (1998), Mitra and Heitjan (2007) and VanderWeele and Arah (2011). Note that this does not imply that $\mathrm{Var}(\hat{\beta}|X, Z, U) = \mathrm{Var}(\hat{\beta}^*|X, Z)$. If we were able to measure U, the variance of our estimate of β would be smaller, however, our correction

does not reduce the uncertainty of our estimate, only the bias. Therefore, the appropriate variance estimate for the corrected $\hat{\beta}$ is $\text{Var}(\hat{\beta}^* | X, Z)$. We confirm this *via* simulation in Section 3.1 where we show that coverage probabilities are close to the nominal level, indicating that $\text{Var}(\hat{\beta}^* | X, Z)$ is a reasonable estimate for the variance of the corrected treatment effect.

The investigator should repeat this procedure for a range of combinations of parameter values and effect sizes, assessing the plausibility of any combination of inputs that materially changes the conclusions. If the significance of the treatment effect changes due to unmeasured confounders with a small effect on cost and similar distributions in the treated and control groups, one can conclude that inferences are vulnerable to the effects of unmeasured confounding. We demonstrate the procedure in Section 4.

3. Simulations. The general plan of the simulations was to generate costs knowing the values of both measured and unmeasured covariates, then fit the reduced regression omitting the unmeasured covariate and compute the adjusted treatment effect. We evaluated performance by the bias of the adjusted estimate and the coverage probability of the resulting 95% confidence interval.

3.1. *Method performance.* Initial simulations evaluated the performance of the method under independence of the unmeasured confounder U and measured confounder Z given treatment status X. We constructed 1000 replications in each simulation, with 100, 200 and 500 subjects per treatment stratum. We drew U independently from its distribution $f_U(U | X, \eta_X)$ using the parameter values η_X shown in Table 1. For example, a Poisson unmeasured confounder was Poisson($\lambda = 1$) in the control group and Poisson($\lambda = 1.58$) in the treated group. We next drew a scalar measured confounder Z, which can represent either a single confounder or a linear combination of a vector Z. Each Z was an independent draw from a Normal distribution with variance 1 and mean 0 in the control group and mean 1 in the treated group. We computed the mean cost according to (1), where $\alpha = 5$, $\beta = 1$, $\theta = 1$, and $\gamma_0 = \gamma_1 = \gamma$. We drew cost outcomes from $f_Y(Y | X, Z, U)$, where f_Y was the Gamma with variance equal to the mean. We drew censoring indicators using a Bernoulli with varying probabilities as indicated in Table 1. We drew failure times from the Exponential with mean 5 and censoring times uniformly over $[0, 10]$, independently of both treatment group and costs. It was only necessary to identify the subjects with censored costs, not generate their actual censored cost values, because the IPW method uses only noncensored observations [Lin (2003)]. In each replication, we calculated $\hat{\beta}^*$ from the apparent model and $\hat{\beta}$ from the adjustment formula. For uncensored data, we used the Gamma GLM, and for censored data the IPW estimator.

Results of the simulation with 100 subjects per treatment stratum appear in Table 1. We observe that the correction worked well for Bernoulli and Normal confounders with modest censoring. As the proportion censored increased, coverage

TABLE 1
Monte Carlo simulations of the adjusted treatment effect, n = 100 per treatment stratum

Distribution	γ	Mean of $\hat{\beta}^{\dagger}$ (coverage probability)			
		Uncensored	25% censoring	50% censoring	75% censoring
Bernoulli$^{\diamond}$	0.25	0.999 (0.963)	1.001 (0.925)	1.001 (0.933)	0.997 (0.817)
	0.5	1.001 (0.958)	0.998 (0.936)	1.003 (0.885)	1.004 (0.790)
	0.75	1.000 (0.954)	1.004 (0.924)	0.999 (0.918)	1.005 (0.780)
	1	1.003 (0.946)	1.007 (0.935)	1.008 (0.909)	1.016 (0.744)
Normal*	0.25	1.001 (0.931)	0.996 (0.950)	1.000 (0.911)	1.000 (0.856)
	0.5	1.002 (0.949)	1.000 (0.950)	1.006 (0.910)	0.995 (0.841)
	0.75	1.010 (0.946)	0.994 (0.928)	1.003 (0.924)	0.983 (0.841)
	1	1.006 (0.941)	0.995 (0.923)	0.982 (0.892)	1.028 (0.815)
Poisson‡	0.25	1.002 (0.959)	1.003 (0.943)	1.001 (0.935)	0.998 (0.841)
	0.5	0.994 (0.939)	0.995 (0.936)	0.986 (0.899)	0.994 (0.804)
	0.75	0.981 (0.924)	0.979 (0.892)	0.965 (0.873)	0.895 (0.768)
	1	0.925 (0.894)	0.922 (0.850)	0.877 (0.828)	0.811 (0.758)
Gamma$^{\triangle}$	0.25	1.001 (0.957)	0.999 (0.948)	1.002 (0.924)	0.996 (0.855)
	0.5	0.999 (0.947)	0.997 (0.919)	0.996 (0.910)	0.997 (0.851)
	0.75	0.998 (0.951)	0.989 (0.913)	0.970 (0.893)	0.957 (0.840)
	1	0.931 (0.881)	0.908 (0.853)	0.915 (0.847)	0.882 (0.780)

† For all models, the true value of β is 1.

$^{\diamond}$ Unmeasured Bernoulli covariate with $\pi_0 = 0.3, \pi_1 = 0.866$.

* Unmeasured Normal covariate with $\mu_0 = 0, \mu_1 = 1, \sigma_0 = \sigma_1 = 1$.

‡ Unmeasured Poisson covariate with $\lambda_0 = 1, \lambda_1 = 1.58$.

$^{\triangle}$ Unmeasured Gamma covariate with $\kappa_0 = \kappa_1 = 0.75, \theta_0 = 0.5, \theta_1 = 0.868$.

probability declined from the nominal level, although the bias remained small. For example, the treatment effect with a Bernoulli unmeasured confounder of effect size $\gamma = 0.5$ had a coverage probability of 0.96 with 0% censoring and 0.79 with 75% censoring, yet the mean of $\hat{\beta}$ remained close to its true value of 1 regardless of censoring. The method worked well for Poisson and Gamma confounders with moderate effect sizes ($\gamma \leq 0.5$), but as the effect of U drew close to the effect of X, the corrected $\hat{\beta}$ incurred a bias and coverage probabilities fell below nominal levels. For example, with Gamma U and uncensored costs, if the effect of U was small ($\gamma = 0.25$), the bias of $\hat{\beta}$ was 0% with a coverage probability of 0.96, but if the effect was large ($\gamma = 1$), the bias was -7% with a coverage probability of 0.88. Increased proportions of censoring further degraded the performance of the estimator. In the uncensored case, increasing the sample size to 2000 per treatment stratum resulted in negligible bias even for large effects (not shown). We conclude that with substantial confounding the adjusted estimates may incur bias in samples of moderate size.

With Gamma and Poisson confounders, 1%–4% of the reduced regression models failed to converge. This occurred only when U had a large effect on cost ($\gamma \geq 0.75$). Analysis of the simulation data revealed that all true models converged, and reduced models failed only when there were single large outliers. Therefore, divergence of the reduced model may itself serve as an indicator of misspecification due to an unmeasured confounder.

3.2. *Sensitivity to departures from conditional independence.* A second set of simulations assessed robustness to violations of conditional independence. Following Hernán and Robins (1999) and VanderWeele (2008), we generated the treatment X conditional on the confounders. We first drew $Z \sim \text{Normal}(1, 1)$, then $U \sim f_U(U|Z = z)$ defined as follows: $\text{Bernoulli}(\pi = \text{expit}(0.5 + 0.2z))$; $\text{Normal}(\mu = 1 + 0.1z, \sigma = 1)$; $\text{Poisson}(\lambda = 0.9 + 0.1z)$; or $\text{Gamma}(\kappa = 0.5, \theta = 0.65 + 0.2|z|)$. With these parameters the unconditional correlation between Z and U is approximately 0.1. Next, we generated $X \sim \text{Bernoulli}(\Pr(X = 1) = \text{expit}(\phi_1 + \phi_2 z + \phi_3 u))$, varying ϕ_1, ϕ_2, ϕ_3 to obtain different patterns of correlation. As before, we drew costs from a Gamma with mean function (1).

Figure 1 plots the bias in the corrected estimate vs. the maximum absolute within-stratum correlation (the larger of the correlations in the two treatment strata) for $n = 500$ under 25% censoring. We see that with correlations above 0.15, biases exceeded 5% even for moderate effect sizes, and the bias increased with the strength of U. Figure 2 plots the bias of the adjusted estimate vs. that of the unadjusted estimate, revealing that the adjustment was substantially more accurate in all cases. Figure 3 plots the coverage probability vs. the maximum absolute within-stratum correlation. We note that effects on coverage probability can be large for modest partial correlations, even when the confounder effect size is also modest.
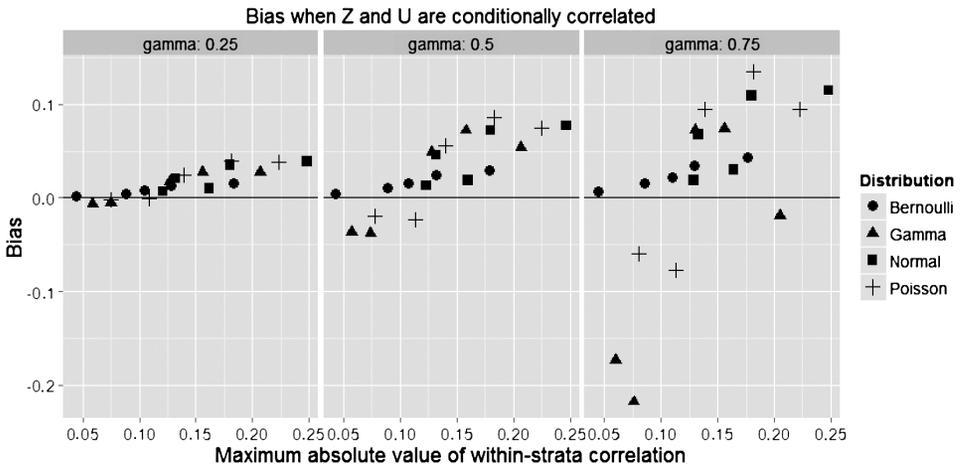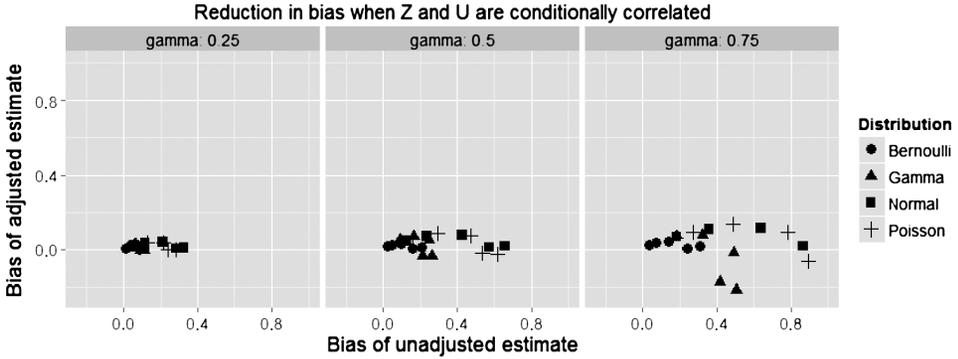

FIG. 1. *Bias of treatment effect.*

FIG. 2. *Reduction in bias from unadjusted estimate.*

For example, consider a Bernoulli U where $\gamma = 0.75$ with $\phi_1 = -1$, $\phi_2 = 1$, and $\phi_3 = 2$. For these parameters, the partial correlation of U and Z is -0.086 in the $X = 1$ stratum and -0.046 in the $X = 0$ stratum. Whereas the unadjusted estimate of β has a bias of 30.8% with 0% coverage, the adjusted treatment effect has a bias of only 1.6% and a coverage of 98.6%.

In further simulations, we explored the relative importance of correlations between U and individual elements of Z vs. the aggregate correlation of U with a summary of Z. We generated a four-variate normal vector (U, Z) where U was a scalar unmeasured confounder and Z was a trivariate measured confounder, assuming two different correlation models: in model 1, the correlations between U and elements of Z were all set to 0.1; in model 2, the correlations were 0.3, -0.4 and 0. We then generated the treatment indicators and computed the propensity scores [estimated $e = \Pr(X = 1|Z)$] [Rosenbaum and Rubin (1983b)] to sum-
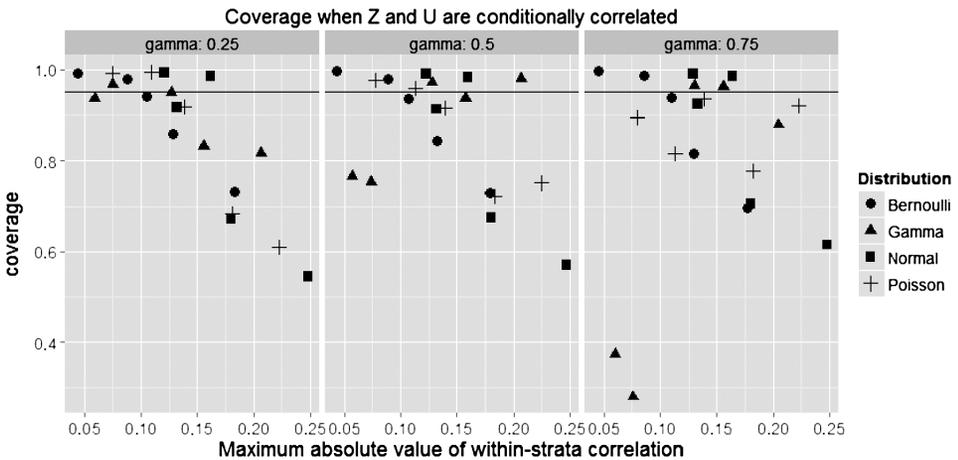

FIG. 3. *Coverage of treatment effect.*

marize the effects of Z. Although individual partial correlations were larger in model 2 than model 1, within-stratum correlations between U and the propensity score were moderate under both models (e.g., -0.02 in the $X = 1$ stratum and -0.10 in the $X = 0$ stratum under model 2). Biases were 0% for model 1 and 2% for model 2. Thus, the simulation, although not exhaustive, suggested that the total correlation of U with the propensity score is more important than individual correlations between U and elements of Z. We also considered the effect of the unconditional correlation between U and Z on bias and coverage. Briefly, larger unconditional correlations of 0.2 and 0.3 led to larger within-stratum correlations, but the bias was largely driven by the within-stratum correlation.

Our simulations suggest that our correction is useful for moderately strong unmeasured confounders when the total within-stratum correlation between the unmeasured and measured confounders is less than 0.15. While the adjusted treatment effect may have some residual bias under conditional dependence, it can still give the investigator a sense of how strongly an unmeasured confounder might influence the treatment effect.

4. Bladder cancer study. We have applied our method to data from a cohort of stage II/III bladder cancer patients diagnosed between 1995 and 2005 and appearing in the linked SEER-Medicare data set [Schrag et al. (2005), Yabroff et al. (2008)]. We restrict analysis to the two main treatments: the standard therapy of radical cystectomy, and a bladder-sparing regimen consisting of radiotherapy and cisplatin-based chemotherapy (rad/chemo). We excluded patients who were diagnosed before age 65, for whom bladder cancer was not the first primary malignancy, who did not have continuous Medicare part A/B coverage or who had coverage from an HMO during the treatment period, or for whom key data such as histology were missing. Radiotherapy and chemotherapy could be given up to 3 months apart, so in order to avoid survivorship bias, we excluded patients who died before 90 days. The final cohort included 1860 patients, of whom 77.4% were treated with cystectomy.

We extracted payment data for Medicare parts A and B from the Carrier Claims file, the Outpatient file, and the Medicare Provider Analysis and Review Record. These files contained payment information for physicians, outpatient services from institutions such as hospitals and health clinics, and inpatient services from hospitals and long-term care facilities [National Cancer Institute (2010)]. All payments were adjusted to year 2000 dollars using the Medicare Economic Index [Centers for Medicare and Medicaid Service (2010), Yabroff et al. (2008)]. This method allowed us to compare total costs of care rather than costs due to treatment only. Two patients had zero costs, suggesting some incompleteness in the payment data; we corrected for this by adding half of the lowest nonzero cost to all observations [Mosteller and Tukey (1977)].

We modeled costs using a GLM with a log link and a gamma variance function. Because the observed total costs were subject to censoring, we estimated the model

by the IPW method [Lin (2003)]. Our analysis considered all health care costs, for which the censoring rate was a moderate 39.0%.

Using a Gamma GLM accounting for censoring, we estimated the effect of rad/chemo on cost, adjusting for the available measured covariates tumor grade, sex, ethnicity, marital status, age, median income of the census tract, size of the metropolitan area, number of comorbidities, year of diagnosis and SEER site. We found that rad/chemo lowers the cost of treatment by a factor of 0.873 (95% CI: 0.793, 0.960).

4.1. *Potential magnitude of conditional correlation.* To determine the plausible extent of violations of the assumption $Z \perp U | X$, we calculated partial correlations between each of the measured confounders, Z_j, and the combined effect of the other covariates estimated by the propensity score excluding Z_j. Some individual pairwise correlations are substantial, for example, the correlation between sex and marital status among the treated is -0.405. When considering the aggregate correlation with all other confounders, however, we found that correlations were typically smaller, often less than 0.1 (see Table 2). The highest combined within-stratum correlations were with race, with the largest observed correlation being 0.148. Although we can never be sure of the magnitude of correlations between U and Z, if aggregate correlations with the unmeasured confounder are no more extreme than the highest observed value, our procedure will yield valid results.

Because our unadjusted estimate was statistically significant, we were primarily interested in assessing potential unmeasured confounders that would displace the effect of treatment toward the null. To this end, we posited three potential confounders: access to health care, manifested as a Bernoulli variable; smoking history, manifested as Poisson; and personal income, manifested as Gamma.

4.2. *Unmeasured Bernoulli confounder.* Although this cohort had Medicare insurance that guaranteed coverage of their care, other factors that are unavailable in SEER-Medicare such as lack of transportation or low physician availability may have limited access [Penchansky and Thomas (1981)]. For example, if rad/chemo patients had poorer access, their overall medical costs may also have been lower. We performed a sensitivity analysis to determine how strong this unmeasured confounder would need to be to explain the treatment effect on costs. As access to health care is a difficult concept to quantify, we categorized it as binary, with $U = 1$ representing good access and $U = 0$ poor access.

Our sensitivity analysis (Table 3) varied both the effect of U on expected cost and the distributions of U in the treatment strata. We see that moderate imbalances in U produced tangibly different inferences from the original analysis, even when the effect of the unmeasured confounder was small. The adjusted treatment effect was nonsignificant when the prevalence of good access to care was 80% in the cystectomy group and 40% in the rad/chemo group and the confounder had an effect of 1.1 on the cost ratio. Alternatively, if the effect of access on cost was 1.25, prevalences of 70% in the cystectomy group and 50% in the rad/chemo group pro-

TABLE 2
Observed correlations in bladder cancer study

	Corr with propensity score			Largest individual corr	
	Unconditional	**Rad/chemo**	**Cystectomy**	**Rad/chemo**	**Cystectomy**
Grade					
1	0.004	−0.004	0.000	0.150	0.082
3	0.006	−0.085	0.024	0.153	0.124
4	−0.007	0.076	−0.013	0.183	0.136
5	−0.006	0.049	−0.051	0.153	0.061
Sex	−0.021	−0.035	−0.059	−0.405	0.373
Race					
White	0.140	0.111	0.148	−0.189	−0.235
Black	−0.075	−0.070	−0.084	0.235	0.177
Other	−0.125	−0.099	−0.125	0.167	0.381
Hispanic	−0.074	−0.001	−0.071	0.264	0.202
Marital status					
Married	−0.003	−0.001	0.012	0.384	0.373
Unmarried	0.014	0.007	0.003	0.405	0.351
Unknown	−0.032	−0.017	−0.044	0.153	0.322
Age	0.134	0.115	0.061	0.199	0.162
Urban Category					
1	−0.053	−0.066	−0.043	−0.408	−0.311
2	0.013	0.009	0.005	0.328	0.297
3	0.048	0.067	0.045	0.490	0.359
Comorbidities					
0	−0.043	−0.044	−0.019	0.097	−0.097
1	−0.071	−0.075	−0.051	0.128	0.066
2 or more	0.091	0.093	0.058	0.140	0.108

duced nonsignificant confidence intervals. We conclude that although the apparent treatment effect is significant, this could be easily explained by an unmeasured Bernoulli confounder.

4.3. *Unmeasured Poisson confounder.* Smoking is a major risk factor for bladder cancer [Baris et al. (2009)] and for co-morbid conditions such as lung and artery disease [Freund et al. (1993)]. Treating these conditions is likely to increase costs. It is common to quantify smoking history in pack-years, defined as the number of packs smoked per day times the length of smoking in years; because this

TABLE 3
Sensitivity of the estimated bladder cancer treatment cost ratio to an
unmeasured Bernoulli confounder

π_0	π_1	Effect of confounder	Cost ratio	95% CI
0.7	0.5	1.1	0.89	(0.81, 0.98)
0.8	0.4	1.1	0.91	**(0.82, 1.00)**
0.8	0.3	1.1	0.92	**(0.83, 1.01)**
0.7	0.5	1.25	0.91	**(0.83, 1.00)**
0.8	0.4	1.25	0.95	**(0.87, 1.05)**
0.8	0.3	1.25	0.97	**(0.89, 1.07)**
0.7	0.5	1.5	0.94	**(0.86, 1.04)**
0.8	0.4	1.5	1.02	**(0.93, 1.12)**
0.8	0.3	1.5	1.06	**(0.97, 1.17)**

Unadjusted cost ratio = 0.87 (95% CI 0.79–0.96).
Boldface confidence intervals denote changes from the unadjusted inference.
π_0 = prevalence of unmeasured confounder in cystectomy group.
π_1 = prevalence of unmeasured confounder in rad/chemo group.

is a discrete count variable, it is natural to model it as Poisson [see Wang and Heitjan (2008)]. Smoking intensity is related to medical costs; a decrease of 20 pack-years has been associated with 10% lower costs [Leigh, Hubert and Romano (2005)]. This means that we could expect costs to increase by a factor of approximately 1.005 per additional pack-year smoked. We estimate the mean number of pack-years in our bladder cancer population to be between 15 and 30 [Baris et al. (2009)]. If exposure differs between treatment strata, the treatment effect may be subject to unmeasured confounding.

Table 4 presents the sensitivity analysis for smoking history. A single pack-year of smoking will not affect costs substantially, so moderate differences in smoking history between the groups are required to change the significance of the treatment effect. If the effect of a single pack-year on cost is 1.005, an average of 34 pack-years in the cystectomy group and 26 pack-years in the rad/chemo group would yield a nonsignificant confidence interval. Because a difference this large is plausible, we deem that the treatment effect is sensitive to smoking history manifested as a Poisson confounder.

4.4. *Unmeasured Gamma confounder.* Income is positively associated with health care spending at the population level [Centers for Medicare and Medicaid Service (2003), Di Matteo (2003)]. Median income for the census tract was available in SEER-Medicare data, and our analysis showed no evidence of imbalance on aggregate, but there was still a possibility of imbalance at the patient level. We therefore hypothesized that an effect of income on spending would be small, and a potential imbalance of income between the patient groups would also be small.

TABLE 4
Sensitivity of the estimated bladder cancer treatment cost ratio to an
unmeasured Poisson confounder

λ_0	λ_1	Effect of confounder	Cost ratio	95% CI
15	13	1.005	0.88	(0.80, 0.97)
15	11	1.005	0.89	(0.81, 0.98)
17	11	1.005	0.90	(0.82, 0.99)
19	11	1.005	0.91	**(0.83, 1.00)**
19	9	1.005	0.92	**(0.83, 1.01)**
15	13	1.01	0.89	(0.81, 0.98)
15	11	1.01	0.91	**(0.83, 1.00)**
17	11	1.01	0.93	**(0.84, 1.02)**
30	28	1.005	0.88	(0.80, 0.97)
30	26	1.005	0.89	(0.81, 0.98)
32	26	1.005	0.90	(0.82, 0.99)
34	26	1.005	0.91	**(0.83, 1.00)**
34	24	1.005	0.92	**(0.83, 1.01)**
30	28	1.01	0.89	(0.81, 0.98)
30	26	1.01	0.91	**(0.83, 1.00)**
32	26	1.01	0.93	**(0.84, 1.02)**

Unadjusted cost ratio = 0.87 (95% CI 0.79–0.96).
Boldface confidence intervals denote changes from the unadjusted inference.
λ_0 = mean of unmeasured confounder in cystectomy group.
λ_1 = mean of unmeasured confounder in rad/chemo group.

Income distributions are heavy-tailed. According to the U.S. census [US Census (2008)], the ratio of the variance to the mean (θ in the Gamma distribution) is approximately 3 for households with members 65 and older. We took income to be measured as a ratio of individual income to mean income. Therefore, if the true effect of income was 1.25, an increase from the mean to twice the mean would cause a 25% increase in medical spending. We further assumed equal θ across treatment strata. Our sensitivity analysis (Table 5) revealed that the treatment effect on cost became nonsignificant when income was 25% higher in the cystectomy group than the rad/chemo group if the effect of income on spending was only 1.05 and $\theta = 3$. As such small differences and effect sizes are plausible, we conclude that the treatment effect is also sensitive to this Gamma confounder.

5. Conclusion. We have developed a method to correct estimates of the treatment effect on mean cost for a range of potential unmeasured confounders. The method is simple to apply and readily accommodates potential binary, count and continuous confounders. Although it assumes independence of the measured and unmeasured confounders given treatment status, simulations reveal that the adjustments are robust to modest departures from this condition. Moreover, in all

TABLE 5
Sensitivity of the estimated bladder cancer treatment cost ratio to an unmeasured Gamma confounder

$Mean_0/Mean_1$	$Var/Mean_p$	Effect of confounder	Cost ratio	95% CI
1.1	2	1.05	0.88	(0.80, 0.97)
1.25	2	1.05	0.89	(0.81, 0.98)
1.5	2	1.05	0.91	**(0.83, 1.01)**
1.1	3	1.05	0.89	(0.81, 0.98)
1.25	3	1.05	0.91	**(0.82, 1.00)**
1.5	3	1.05	0.94	**(0.85, 1.03)**
1.1	2	1.1	0.89	(0.81, 0.98)
1.25	2	1.1	0.92	**(0.83, 1.01)**
1.5	2	1.1	0.96	**(0.87, 1.06)**
1.1	3	1.1	0.90	(0.82, 0.99)
1.25	3	1.1	0.95	**(0.86, 1.04)**
1.5	3	1.1	1.02	**(0.92, 1.12)**

Unadjusted cost ratio $= 0.87$ (95% CI 0.79–0.96).
Boldface confidence intervals denote changes from the unadjusted inference.
$Mean_0 =$ mean of unmeasured confounder in cystectomy group.
$Mean_1 =$ mean of unmeasured confounder in rad/chemo group.
$Mean_p =$ mean pooled across both groups.

conditions we examined the adjustment reduced bias compared to an unadjusted analysis.

Conditional independence is a generally untestable assumption that is unlikely to hold exactly in applications. Investigators should therefore evaluate patterns of correlation within their data before applying the method. We recommend exploring the possible size of partial correlations by successively omitting individual observed confounders and estimating their partial correlations with the remaining measured confounders. Simulations show that the adjustment is robust to within-stratum correlations less than 0.15, a condition that appears to be satisfied in our bladder cancer example.

In some types of studies one expects to find substantial correlations among measured and unmeasured confounders. For example, in a study to examine cost of an intervention to improve student performance on a standardized exam, measured confounders related to socioeconomic status could be highly correlated with unmeasured confounders such as parental attitudes toward education. If most measured confounders are tightly interrelated with the unmeasured variable, this would render our adjustments inaccurate. Here the investigator should use the method of VanderWeele and Arah (2011), perhaps with simplifying assumptions about the confounder to render the procedure more tractable.

A second limitation of our method is that adjustments for strong Poisson and Gamma confounders can be biased in moderate sample sizes. The correction may also leave bias when there is a large proportion of censored data, as would occur in studies with short follow-up. This is likely due to the decreased efficiency of the Lin estimator when the proportion of censoring is high, a problem that could be ameliorated by splitting the time period into intervals [Lin (2003)] or the creation of a more efficient GLM for censored costs.

With our simple additive corrections, assuming fixed values of the sensitivity parameters leads us to use the unadjusted $\text{Var}(\hat{\beta}^*)$ to develop confidence intervals. A fully Bayesian approach would quantify uncertainty about the parameters using prior distributions and propagate the uncertainty using Bayes's theorem. One could also conduct a probabilistic sensitivity analysis using Monte Carlo simulations as demonstrated in Arah, Chiba and Greenland (2008). Although it is appealing to incorporate uncertainty about the sensitivity parameters, both methods require sufficient knowledge about unmeasured parameters to fully characterize their distributions.

Despite its limitations, our method provides a practical approach to evaluate the sensitivity of costs to unmeasured confounding. The bladder cancer study provides an example where our method can be easily implemented to evaluate the effect of a wide range of confounders. Further research should develop extensions to allow sensitivity analysis for combined cost-effectiveness outcomes.

REFERENCES

ARAH, O. A., CHIBA, Y. and GREENLAND, S. (2008). Bias formulas for external adjustment and sensitivity analysis of unmeasured confounders. *Ann. Epidemiol.* **18** 637–646.

AXELSON, O. and STEENLAND, K. (1988). Indirect methods of assessing the effects of tobacco use in occupational studies. *Am. J. Ind. Med.* **13** 105–118.

BANG, H. and TSIATIS, A. A. (2000). Estimating medical costs with censored data. *Biometrika* **87** 329–343. MR1782482

BANG, H. and TSIATIS, A. A. (2002). Median regression with censored cost data. *Biometrics* **58** 643–649. MR1926117

BARIS, D., KARAGAS, M. R., VERRILL, C., JOHNSON, A., ANDREW, A. S., MARSIT, C. J., SCHWENN, M., COLT, J. S., CHERALA, S., SAMANIC, C., WADDELL, R., CANTOR, K. P., SCHNED, A., ROTHMAN, N., LUBIN, J., FRAUMENI, J. F., HOOVER, R. N., KELSEY, K. T. and SILVERMAN, D. T. (2009). A case-control study of smoking and bladder cancer risk: Emergent patterns over time. *J. Natl. Cancer Inst.* **101** 1553–1561.

BASU, A. and MANNING, W. G. (2010). Estimating lifetime or episode-of-illness costs under censoring. *Health Economics* **19** 1010–1028.

BLACK, N. (1996). Why we need observational studies to evaluate the effectiveness of health care. *BMJ* (*Clinical Research Ed.*) **312** 1215.

CENTERS FOR MEDICARE AND MEDICAID SERVICE (2003). Medicare current beneficiary survey, 2003. Available at http://www.cms.gov/Research-Statistics-Data-and-Systems/Research/MCBS/index.html. (Accessed Nov 2010.)

CENTERS FOR MEDICARE AND MEDICAID SERVICE (2010). Medicare economic index. Technical report, Centers for Medicare and Medicaid Services. Available at http://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/MedicareProgramRatesStats/Downloads/mktbskt-summary.pdf.

DI MATTEO, L. (2003). The income elasticity of health care spending. A comparison of parametric and nonparametric approaches. *Eur. J. Health Econ.* **4** 20–29.

DODD, S., BASSI, A., BODGER, K. and WILLIAMSON, P. (2006). A comparison of multivariable regression models to analyse cost data. *J. Eval. Clin. Pract.* **12** 76–86.

ETZIONI, R. D., FEUER, E. J., SULLIVAN, S. D., D., L., HU, C. and RAMSEY, S. D. (1999). On the use of survival analysis techniques to estimate medical care costs. *Journal of Health Economics* **18** 365–380.

FLANDERS, W. D. and KHOURY, M. J. (1990). Indirect assessment of confounding: Graphic description and limits on effect of adjusting for covariates. *Epidemiology* **1** 239–246.

FREUND, K. M., BELANGER, A. J., D'AGOSTINO, R. B. and KANNEL, W. B. (1993). The health risks of smoking. The Framingham Study: 34 years of follow-up. *Ann. Epidemiol.* **3** 417–424.

GAIL, M. H., WACHOLDER, S. and LUBIN, J. H. (1988). Indirect corrections for confounding under multiplicative and additive risk models. *Am. J. Ind. Med.* **13** 119–130.

HEITJAN, D. F., KIM, C. Y. and LI, H. (2004). Bayesian estimation of cost-effectiveness from censored data. *Stat. Med.* **23** 1297–1309.

HERNÁN, M. A. and ROBINS, J. M. (1999). Method for conducting sensitivity analysis. *Biometrics* **55** 1316–1317.

HORNBERGER, J. and WRONE, E. (1997). When to base clinical policies on observational versus randomized trial data. *Ann. Intern. Med.* **127** 697–703.

HOSMAN, C. A., HANSEN, B. B. and HOLLAND, P. W. (2010). The sensitivity of linear regression coefficients' confidence limits to the omission of a confounder. *Ann. Appl. Stat.* **4** 849–870. MR2758424

LEIGH, J. P., HUBERT, H. B. and ROMANO, P. S. (2005). Lifestyle risk factors predict healthcare costs in an aging cohort. *American Journal of Preventative Medicine* **29** 379–387.

LIN, D. Y. (2000). Linear regression analysis of censored medical costs. *Biostatistics* **1** 35–47.

LIN, D. Y. (2003). Regression analysis of incomplete medical cost data. *Stat. Med.* **22** 1181–1200.

LIN, D. Y., PSATY, B. M. and KRONMAL, R. A. (1998). Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics* **54** 948–963.

LIN, D. Y., FEUER, E. J., ETZIONI, R. and WAX, Y. (1997). Estimating medical costs from incomplete follow-up data. *Biometrics* **53** 419–434.

LIU, L. (2009). Joint modeling longitudinal semi-continuous data and survival, with application to longitudinal medical cost data. *Stat. Med.* **28** 972–986. MR2518360

LIU, L., HUANG, X. and O'QUIGLEY, J. (2008). Analysis of longitudinal data in the presence of informative observational times and a dependent terminal event, with application to medical cost data. *Biometrics* **64** 950–958. MR2526647

MCCANDLESS, L. C., GUSTAFSON, P. and LEVY, A. (2007). Bayesian sensitivity analysis for unmeasured confounding in observational studies. *Stat. Med.* **26** 2331–2347. MR2368419

MITRA, N. and HEITJAN, D. F. (2007). Sensitivity of the hazard ratio to nonignorable treatment assignment in an observational study. *Stat. Med.* **26** 1398–1414. MR2345728

MOSTELLER, F. and TUKEY, J. W. (1977). *Data Analysis and Regression*. Addison-Wesley, Reading.

NATIONAL CANCER INSTITUTE (2010). SEER-Medicare: Medicare claims files. Available at http://healthservices.cancer.gov/seermedicare/medicare/claims.html.

NATIONAL CANCER INSTITUTE (2011). The cost of cancer. Available at http://www.cancer.gov/aboutnci/servingpeople/cancer-statistics/costofcancer.

PENCHANSKY, R. and THOMAS, J. W. (1981). The concept of access: Definition and relationship to consumer satisfaction. *Med. Care* **19** 127–140.

ROSENBAUM, P. R. and RUBIN, D. B. (1983a). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **45** 212–218.

ROSENBAUM, P. R. and RUBIN, D. B. (1983b). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70** 41–55. MR0742974

SCHLESSELMAN, J. J. (1978). Assessing effects of confounding variables. *Am. J. Epidemiol.* **108** 3–8.

SCHRAG, D., MITRA, N., XU, F., RABBANI, F., BACH, P. B., HERR, H. and BEGG, C. B. (2005). Cystectomy for muscle-invasive bladder cancer: Patterns and outcomes of care in the Medicare population. *Urology* **65** 1118–1125.

TIAN, L. and HUANG, J. (2007). A two-part model for censored medical cost data. *Stat. Med.* **26** 4273–4292. MR2405353

US CENSUS (2008). Annual social and economic (ASEC) Supplement; 2008. Available at http://www.census.gov/hhes/www/cpstables/032009/hhinc/new02_001.htm.

VANDERWEELE, T. J. (2008). Sensitivity analysis: Distributional assumptions and confounding assumptions. *Biometrics* **64** 645–649, 671. MR2432439

VANDERWEELE, T. J. and ARAH, O. A. (2011). Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. *Epidemiology* **22** 42–52.

WANG, H. and HEITJAN, D. F. (2008). Modeling heaping in self-reported cigarette counts. *Stat. Med.* **27** 3789–3804. MR2526609

WANG, L. and KRIEGER, A. M. (2006). Causal conclusions are most sensitive to unobserved binary covariates. *Stat. Med.* **25** 2257–2271. MR2240099

WILLAN, A. R., LIN, D. Y., COOK, R. J. and CHEN, E. B. (2002). Using inverse-weighting in cost-effectiveness analysis with censored data. *Stat. Methods Med. Res.* **11** 539–551.

YABROFF, K. R., LAMONT, E. B., MARIOTTO, A., WARREN, J. L., TOPOR, M., MEEKINS, A. and BROWN, M. L. (2008). Cost of care for elderly cancer patients in the United States. *J. Natl. Cancer Inst.* **100** 630–641.

YANAGAWA, T. (1984). Case-control studies: Assessing the effect of a confounding factor. *Biometrika* **71** 191–194. MR0738341

E. A. HANDORF
BIOSTATISTICS AND BIOINFORMATICS FACILITY
FOX CHASE CANCER CENTER
TEMPLE UNIVERSITY HEALTH SYSTEM
333 COTTMAN AVENUE
PHILADELPHIA, PENNSYLVANIA 19111
USA
E-MAIL: elizabeth.handorf@fccc.edu

J. E. BEKELMAN
DEPARTMENT OF RADIATION ONCOLOGY
PERELMAN SCHOOL OF MEDICINE
UNIVERSITY OF PENNSYLVANIA
3400 SPRUCE STREET
PHILADELPHIA, PENNSYLVANIA 19104
USA
E-MAIL: bekelman@uphs.upenn.edu

D. F. HEITJAN
N. MITRA
DEPARTMENT OF BIOSTATISTICS AND EPIDEMIOLOGY
PERELMAN SCHOOL OF MEDICINE
UNIVERSITY OF PENNSYLVANIA
423 GUARDIAN DRIVE
PHILADELPHIA, PENNSYLVANIA 19104-6021
USA
E-MAIL: dheitjan@upenn.edu
        nanditam@upenn.edu