Indemnity Benefit Duration and Obesity

Frank Schmid, Chris Laws, and Matthew Montero

Abstract

Motivation. There is mounting evidence of obesity contributing to the cost of workers compensation. Longitudinal studies by Duke University (Østbye, Dement, and Krause [5]) of its own employees—and by Johns Hopkins University (Pollack et al. [7]) of employees of a multi-site U.S. aluminum manufacturing company—point to substantially higher odds of injury for workers in the highest obesity category. Further, a 2011 Gallup survey (Witters and Agrawal [9]) of 109,875 full-time employees found that obese employees account for a disproportionately high number of missed workdays, thus causing a significant loss in economic output. Finally, an NCCI study (Laws and Schmid [4]) of workers compensation claims established that where claimants are assigned a comorbidity code indicating obesity, the medical costs of the claim are a multiple of what is observed otherwise. In the following study, using a methodology similar to the one employed by Laws and Schmid [4], but accounting for a possible immortal time bias, it is shown that the indemnity benefit duration of claimants with an obesity-related comorbidity indicator is a multiple of what is observed for comparable non-obese claimants.

Method. The study makes use of a matched-pairs research design, where obese claims are matched with comparable non-obese claims in the data set. Exact matching applies to all claim characteristics, except age at injury, where proximity matching is employed. The set of matched pairs is then analyzed using a semiparametric Bayesian Weibull proportional hazard model, the nonparametric component of which accounts for the possible nonlinear influence of age. Aside from age, an indicator variable signifying obesity is the only covariate in the model—this is because net of these two covariates (and duration, which serves as the dependent variable), the claims within each set of matched pairs are identical. The model is estimated by means of MCMC (Markov chain Monte Carlo simulation).

Results. The study shows that, based on Temporary Total and Permanent Total indemnity benefit payments, the duration of obese claimants is more than five times the duration of non-obese claimants, after controlling for primary ICD-9 code, injury year, U.S. state, industry, gender, and age. When Permanent Partial benefits are counted toward indemnity benefit duration as well, this multiple climbs to more than six.

Availability. The model was implemented in R (cran.r-project.org/) using the sampling platform JAGS (Just Another Gibbs Sampler, mcmc-jags.sourceforge.net/). JAGS was linked to R by means of the R package rjags (cran.r-project.org/web/packages/rjags/index.html).

Keywords. Obesity, Duration, Proportional Hazard Model, Semiparametric Model, Workers Compensation

1. INTRODUCTION

There is mounting evidence of obesity contributing to the cost of workers compensation. Longitudinal studies by Duke University (Østbye, Dement, and Krause [5]) of its own employees—and by Johns Hopkins University (Pollack et al. [7]) of employees of a multi-site U.S. aluminum manufacturing company—point to substantially higher odds of injury for workers in the highest obesity category. Further, a 2011 Gallup survey (Witters and Agrawal [9]) found that obese employees account for a disproportionately high number of missed workdays, thus causing a significant loss in economic output. Finally, an NCCI study (Laws and Schmid [4]) of workers compensation claims found that claimants with a comorbidity code indicating obesity experience medical costs that are a multiple of what is observed for comparable non-obese claimants.

In contrast to the mentioned studies, there is research that finds no evidence for a link between work-related injuries and obesity. Shaw et al. [8] "assess the effect of body mass index (BMI) on pain and function outcomes in the acute and sub-acute stages of work-related low back pain (LBP)," finding that "BMI is not a useful prognostic factor during the acute and sub-acute stages of work-related LBP." Similarly, Pollack and Cheskin [6], who review the early literature on the relation between workplace injuries and obesity, report only weak evidence for such a link.

In the following study, using a methodology similar to the one employed by Laws and Schmid [4], but accounting for a possible immortal time bias, it is shown that the indemnity benefit duration of claimants with an obesity-related comorbidity indicator is a multiple of what is observed for comparable non-obese claimants. The study makes use of a matched-pairs research design, where obese claims are matched with comparable non-obese claims in the data set. Exact matching applies to all claim characteristics, except age at injury, where proximity matching is employed. The set of matched pairs is then analyzed using a semiparametric Bayesian Weibull proportional hazard model, the nonparametric component of which accounts for the possible nonlinear influence of age. Aside from age, an indicator variable signifying obesity is the only covariate in the model—this is because net of these two covariates (and duration, which serves as the dependent variable), the claims within each set of matched pairs are identical. The model is estimated by means of Markov chain Monte Carlo simulation (MCMC). The study shows that, based on Temporary Total and Permanent Total indemnity benefit payments, the duration of obese claimants is more than five times the duration of non-obese claimants, after controlling for primary ICD-9 code, injury year, U.S. state, industry, gender, and age. When Permanent Partial benefits are counted toward duration as well, this multiple climbs to more than six.

1.1 Research Context

The Duke University study by Østbye, Dement, and Krause [5] is a comprehensive statistical analysis of the effect of obesity on the cost of workers compensation. This study makes use of a longitudinal data set, which was obtained by monitoring a cohort of 11,728 employees of Duke University and the Duke University Health System from January 1, 1997, through December 31, 2004. The cohort was defined by all employees who had at least one health risk assessment (HRA) during this time period; taking an HRA is voluntary and available to employees eligible for healthcare benefits. (Note that the number of members in the study may have shrunk over time due to employee termination or disability.) The members of this cohort were assigned to body mass index (BMI) categories based on the first HRA they participated in during the time of the study.

At the end of the eight-year time window, the number of workers compensation claims, the number of lost workdays, and the indemnity and medical costs related to workers compensation were tallied for each employee; then, this information was matched to the BMI category (and other characteristics) of the claimant. There are six BMI categories, ranging from underweight to recommended weight, overweight, and three classes of obesity. The highest level of obesity is class III, which comprises the morbidly obese, identified by a BMI of 40 or higher. The Duke University study finds that for the morbidly obese employees, the medical costs are 6.8 times the costs for employees of recommended weight. At the same time, an employee in this morbidly obese group is twice as likely to have a claim, while the number of lost workdays is almost 13 times higher. For obese classes II (BMI of at least 35 but less than 40) and I (BMI of at least 30 but less than 35), the medical costs per employee are (respectively) 3.1 and 2.6 times the medical costs for employees of recommended weight; the respective multiples for the number of claims read 1.9 and 1.5, while the respective multiples for the lost workdays are 8.3 and 5.3. (The numbers cited above rest on the bivariate analysis presented in Østbye, Dement, and Krause [5], Table 3.)

Another way of presenting the findings of the Duke University study is on an approximate per-claim basis. Transforming the medical costs per 100 full-time equivalent (FTE) employees into costs per claim based on the reported means shows that this amount is 3.4 (obesity class III), 1.7 (obesity class II), and 1.7 (obesity class I) times the magnitude recorded for employees of normal weight. The corresponding numbers for indemnity read 5.5 (obesity class III), 3.4 (obesity class II), and 2.9 (obesity class I). Finally, the number of lost workdays per claim for obesity classes III, II, and I are 6.4, 4.5, and 3.5 (respectively) times the amount for claimants of recommended weight. (Here, too, the numbers are based on the bivariate analysis presented in Table 3 of Østbye, Dement, and Krause [5].)

The Johns Hopkins University study (Pollack et al. [7]) offers no direct evidence of an effect of obesity on lost workdays. The authors analyze the distribution and odds of occupational injury among 7,690 workers at eight plants of a U.S. aluminum manufacturing company from January 1, 2002, through December 31, 2004. The study finds that the odds of injury in the highest obesity group, as compared to the ideal body weight, runs at 2.21; injuries to the leg or knee were particularly prevalent among this very obese group.

Finally, the Gallup survey (Witters and Agrawal [9]) focused on "unhealthy days" per month, which are then converted into missing workdays at a rate of 0.31 missed workdays per unhealthy day. The survey was conducted from January 2, 2011, through October 2, 2011, covering 109,875

full-time employees. Employees of normal weight and no chronic conditions establish the baseline rate; the chronic conditions included in the survey comprise a past heart attack, high blood pressure, high cholesterol, cancer, diabetes, asthma, depression, and certain types of recurring pain. The baseline rate of unhealthy days per month runs at 0.34. At 0.36, this number is slightly higher for the group of overweight or obese and no chronic conditions. Similarly, the number of unhealthy days per month for those with one or two chronic conditions is higher for the overweight and obese (at 1.08) than those of normal weight (1.07). Finally, for employees with three or more chronic conditions, the numbers of unhealthy days per month read 3.51 for the overweight and obese and 3.48 for those of normal weight. Although correlation does not establish causation, it is of interest that among the full-time working population, of the employees who have one or two chronic conditions, there are twice as many obese or overweight employees than there are employees of normal weight. When it comes to the group of employees with three or more chronic conditions, the number of obese or overweight employees exceeds those of normal weight by a factor of more than three.

1.2 Objective

The effect of obesity on indemnity benefit duration is studied using nonparametric and semiparametric statistical approaches. Indemnity benefit duration is measured in two alternative ways. Both statistical analyses account for interval-censoring and right-censoring of duration. The nonparametric analysis rests on Kaplan-Meier plots; the semiparametric model employs a partially linear regression in a Weibull proportional hazard framework.

1.3 Outline

Section 2 offers a detailed description of how the data set was compiled, how obesity was identified, and how indemnity benefit duration was calculated. This is followed by a discussion of challenges that arise from matching by age, and by descriptive information of the matched pairs that enter the analysis. Section 3 then offers a nonparametric analysis based on Kaplan-Meier plots and a semiparametric analysis using a partially linear Weibull proportional hazard model. Section 4 concludes. An appendix offers technical details on the immortal time bias correction, the choice of the primary ICD-9 code, the conversion of benefit time into calendar time for the waiting period, and the calculation of duration where there are multiple matches; further, the appendix displays the JAGS computer code for the Weibull proportional hazard model.

2. THE DATA

We use a large data set of workers compensation claims provided by a set of insurance companies. The data comprise records from 40 states (AK, AL, AR, AZ, CO, CT, DC, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MD, ME, MI, MO, MS, MT, NC, NE, NH, NM, NV, OK, OR, RI, SC, SD, TN, TX, UT, VA, VT, WI, and WV) and 11 injury years (1998–2008). The transaction dates run through December 31, 2009, thus allowing for the accumulation of at least one year of indemnity and medical transactions for any given claim. Further, the data is evaluated as of June 30, 2010, thereby accommodating a reporting lag of six months. Claimants who are reported as being 15 years of age or younger are excluded.

A claim is categorized as "obese" if, for any diagnostic field (other than the primary one) on any medical transaction that occurs within 12 months after injury, the three leading digits of the ICD-9 code equal 278. The ICD-9 code 278, which is "non-reimbursable," denotes overweight, obesity, and other hyperalimentation. Note that the claim characteristic of interest rests on medical transactions. For instance, a claim whose indemnity benefits ends after one month without having acquired the ICD-9 code 278 may still qualify as obese if, within the following 11 months, there is a medical transaction with a 278 comorbidity indicator.

Allowing claims 12 months to acquire the ICD-9 code 278 creates "immortal time," that is, time during which the claim could not have ended—otherwise, the claim would not have been able to acquire the ICD-9 code 278. The problem that arises from there being immortal time may best be illustrated for the case where there is no difference between indemnity benefit and medical benefit durations. Assume that obesity has no impact on the duration of a claim and that any given claim on any given day has a certain probability of ending and, independent of that, a certain probability of assuming the ICD-9 code 278. The longer the duration of a claim is, the more time this claim has to acquire the code; as a result, the median duration of claims with an ICD-9 code 278 tends to exceed the median duration of claims that did not acquire this code. This difference in median durations is known as "immortal time bias."

In order to correct for the immortal time bias, it is stipulated that a non-obese claim be a potential match for an obese claim only if this non-obese claim had at least one medical transaction past the time span it took the obese claim to acquire the ICD-9 code 278. As shown in Appendix 5.1, this immortal time bias correction (ITBC) almost perfectly removes the bias, given the nature of the data set. This does not imply that the applied ITBC approach will perform well in general.

The data set offers no information on the body mass index (BMI), which is a standard measure of obesity. As a result, we are not able to differentiate among degrees of obesity. It can be assumed that the comorbidity indicator identifies the claimant as severely obese, thus putting this claimant into one of the higher obesity classes. This conjecture is supported by the fact that only 0.15 percent of the claims in the data set acquire the ICD-9 code 278 (within 12 months of the date of injury). By comparison, in the Duke University study (Østbye, Dement, and Krause [5]), the proportion of morbidly obese claimants amounts to 4.9 percent.

The small percentage of obese claims in the total number of claims, along with the very high total number of claims (4,821,562) is most suited for a matched-pairs research design. In such an analytical setting, each obese claim in the data set is matched with a non-obese claim of the same primary ICD-9 code (at the three-digit level), injury year, U.S. state, industry, gender, and age at injury. Except for age at injury, all matching criteria are categorical in nature, which allows for exact matching (thus obviating the need for propensity matching). At the same time, extending the concept of exact matching to age at injury (which is measured on a scale of integers) leaves many obese claims without matches. Thus, for age at injury, we use proximity matching.

Proximity matching rests on the concept of the nearest neighbor. When exact matching for age is not feasible (for lack of exact matches when using integers for years of age), researchers often resort to matching by age bracket. In matching by age bracket, an obese claim is paired with a non-obese claim that belongs to the same (for instance) five-year age bracket (subject to being identical by all exact-matching characteristics). A disadvantage of matching by age bracket is that many claims are not matched with the closest neighbor. For instance, a 20-year-old obese claimant may be matched (within the 20–24 age bracket) with a 24-year-old non-obese claimant but is prevented from being matched with an otherwise identical 19-year-old claimant. Proximity matching avoids this problem by looking for the nearest neighbor. At the same time, it may not be appropriate to have the concept of the nearest neighbor rest on the simple age difference. Because aging is a nonlinear process, it may be preferable to match an obese 25-year-old claimant with a 20-year-old instead of a 30-year-old. Similarly, matching an obese 55-year-old with a non-obese 60-year-old may be more appropriate than matching this person with a non-obese 50-year-old. For this reason, we use a sigmoid function to create a fuzzy set for youngness; the sigmoid function reads $1/(1+\exp(-\sigma \cdot (h-45)))$, where h is the age at injury and σ was chosen to equal 0.12. (For the concept of fuzzy sets, see, for instance, Kasabov [2].) Chart 1 depicts the fuzzy set; the degree of youngness is set to 50 percent at age 45 and then, in an "S-shaped" manner, it gets near zero and 100 percent as

the years of age approach 100 and zero, respectively. The nearest neighbor of an obese claim to an otherwise identical non-obese claim is defined by the smallest difference (in absolute value terms) in the degree of youngness.

Chart 1: Degree of Youngness, Defined by Sigmoid Function

In establishing the primary ICD-9 code, there is an order of priority based on (1) providers, (2) the time in the duration of a claim at which the medical service was provided, and then (3) the paid amount. The logic of establishing the primary ICD-9 code is detailed in Appendix 5.2.

Matching by industry relies on an NCCI industry classification; the five industries comprise Manufacturing, Contracting, Office and Clerical, Goods and Services, and Miscellaneous.

To summarize, we define as neighbors to a given obese claim the set of non-obese claims that match exactly based on primary ICD-9 code, injury year, U.S. state, industry, and gender; only non-obese claims are considered that have at least one medical transaction following the acquisition of the ICD-9 code 278 by the pertinent obese claim. In this matching process, a given non-obese claim may be used as a neighbor to more than one obese claim. Among the thus identified set of neighbors, the nearest neighbor is chosen based on the degree of youngness.

This nearest neighbor may not be unique because of ties in the youngness distance, in which case there are multiple matches. In the event of multiple non-obese matches for a given obese claim, a synthetic claim is generated using the median duration from the set of multiple matches. The algorithm for calculating the median indemnity benefit duration for a set of multiple matches is provided in Appendix 5.3.

2.1 Indemnity Benefit Duration

We calculate indemnity benefit duration (duration, for short) based on recorded indemnity transactions. The transactions are from the payment categories Temporary Total and Permanent Total; claims with transactions in the category Fatal are excluded. In a sensitivity analysis, we count transactions in the Permanent Partial category toward duration as well.

The duration of indemnity benefits may be censored. Censoring occurs when the duration is only partially known. In the event of right-censoring, the duration is only known to be greater than a certain value. In the case of interval-censoring, the duration falls into a given interval, as is the case with a claim that does not breach the waiting period. Note that for the purpose of this study, we count the waiting period toward the indemnity benefit duration even where the retroactive period has not been reached.

We employ two concepts of duration. Duration concept I aggregates compensated time intervals based on the "from-date" and "through-date" information associated with benefit payments. If more than 20 percent of the transactions associated with a given claim miss the from-/through-date information, then the claim is categorized as right-censored at the waiting period.

Duration concept II measures the time interval between injury date and most recent through-date. Transaction dates substitute for through-dates where the latter are missing. On one hand, duration concept II appears to be less sensitive to missing observations; on the other hand, if a claimant returned to work intermittently, duration concept II fails to capture this information.

For claims without indemnity transactions, it is assumed that the claim did not breach the waiting period. Such claims are treated as interval-censored on the waiting period. For claims with indemnity transactions, benefit duration is established by means of the concept of dormancy; only transactions prior to the date of dormancy are considered. The date of dormancy is the first transaction date that is not followed by another transaction within 180 days; reopenings past the date of dormancy are not accounted for. The latest date for which such a transaction-free time interval can be established is 180 days prior to December 31, 2009 (which is the latest transaction

date considered in this study). Claims for which no 180-day transaction-free interval can be established are categorized as right-censored.

Both duration concepts measure calendar time (as opposed to number of missed workdays). Where there are transactions of indemnity payments available, measuring calendar time is straightforward; this is because the time intervals covered by the from-/through-date sets of recorded transactions tend to refer to calendar time rather than the number of compensated workdays; this calendar year information can then be linked to the date of injury. Measuring duration when there are no indemnity transactions is more demanding. Appendix 5.4 offers an algorithm for measuring the elapsed calendar time that leads up to the end of the waiting period.

2.2 Matching-Toward-the-Center Problem

Matching by the degree of youngness tends to pair old claimants with younger claimants and young claimants with older claimants; this is because the numbers of claimants on the edges of the age distribution tend to be comparably small. For instance, within the age bracket 60–64, the number of observations tends to decline with age. Thus, a claimant at the center of this age bracket has a comparatively high chance of being "matched down the age distribution," as opposed to being "matched up"; this is simply because the number of potential matches within this bracket is greater at lower ages than at higher ages. Conversely, young claimants tend to be "matched up the age distribution." Note that this matching toward the center (of the age distribution) is not unique to proximity matching, but it is also characteristic of the traditional approach of matching by the age bracket. Although proximity matching by age, as applied in this analysis, mitigates the matching-toward-the-center-problem, it turns out not to eliminate the problem entirely.

In Laws and Schmid [4], the dependent variable is the (natural) logarithm of the ratio of obese to non-obese claim costs; age at injury serves as a covariate. In such a regression design, matching toward the center of the age distribution affects the interpretation of the regression results for age (and age alone). As pointed out by the authors, if the medical costs of workers compensation claims increase with age, then the influence of age on the cost of obesity is underestimated for young claimants and overestimated for old claimants. As a consequence, the estimated effect of age may be distorted on the edges of the age distribution, taking on an "S-shaped" form.

In this study, and unlike in Laws and Schmid [4], for each pair of obese/non-obese claims, both observations enter the statistical analysis. As a result, the estimated effect of age is not adversely

affected by the matching-toward-the-center problem as long as the influence of age is properly accounted for in the regression equation.

2.3 Descriptive Characteristics

We start out with a data set of 4,821,562 claims, of which 7,145 (or 0.1482 percent) are categorized as obese. We are able to match 6,528 obese claims (or 91 percent) to non-obese claims, among which we identified 487,603 potential matches before correcting for the immortal time bias and applying proximity matching by the degree of youngness. Correcting for the immortal time bias reduces the number of usable obese claims to 6,435; the number of potential non-obese matches drops to 356,423. After proximity matching by the degree of youngness, the number of potential matches reduces to 14,929, which translates into an approximate 2.3 non-obese matches per obese claim.

Table 1 details the number of matched pairs by injury year. The non-obese claim of a matched pair may be synthetic where there are multiple matches, as discussed above.

Table 1: Numbers of Matched Pairs by Injury Year

Injury Year	Number of Matched Pairs
1998	333
1999	319
2000	391
2001	472
2002	550
2003	548
2004	591
2005	678
2006	759
2007	841
2008	953
Total	6,435

Chart 2 offers the relative claim frequency by U.S. state, pooled over the studied 11 injury years. Florida and Texas are the most highly represented states, whereas Rhode Island and West Virginia are the least highly represented. The representation of a state in the data set depends primarily on the size of its labor force, but it also depends on the combined market share of the insurance companies that contribute to the analyzed data set. Another contributing factor is the share represented by the self-insured.

Chart 2: Relative Frequency of Matched Pairs by State, 1998–2008

Chart 3 presents a distribution of age at injury of the set of studied (obese and non-obese) claims; again, the observations are pooled over the 11 injury years. The median age at injury is in the neighborhood of 40.
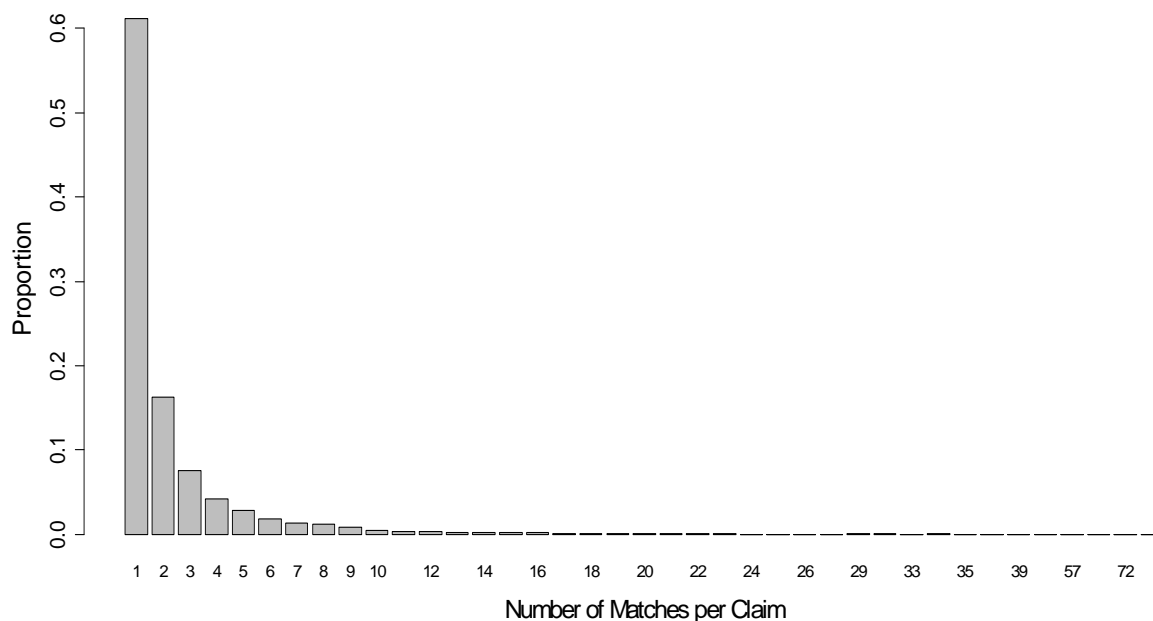
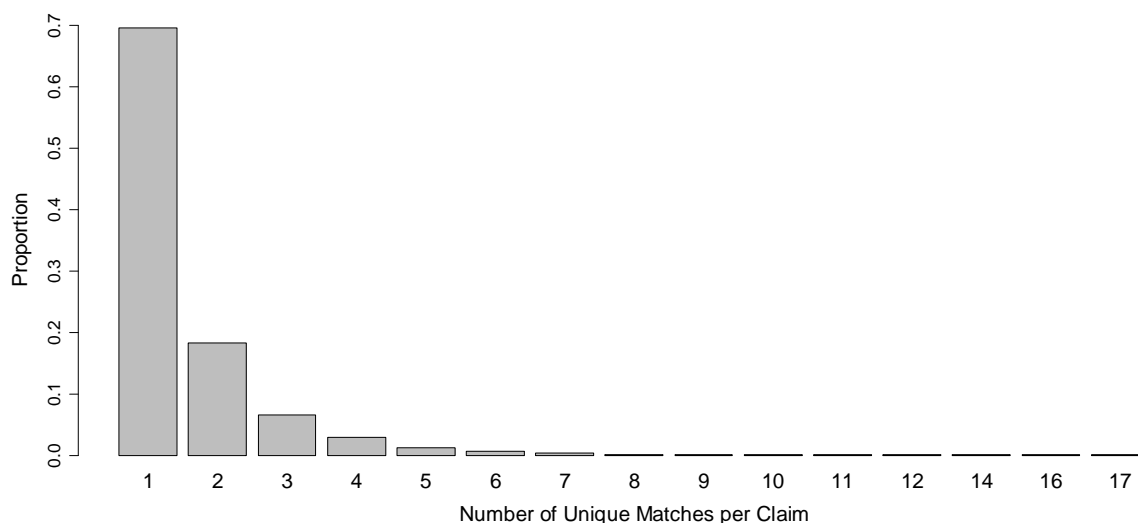**Chart 3:** Relative Frequency of Claims by Age at Injury, 1998–2008

Chart 4 displays the relative frequency of the number of matches per obese claim. As mentioned, on average, there are 2.3 matches per obese claim, but, as the chart illustrates, this distribution is highly skewed. The majority of obese claims (around 60 percent) have only one match; obese claims with more than 10 matches are rare. As discussed, where matches are not unique, the median duration is calculated using the algorithm detailed in Appendix 5.3.

Chart 4: Relative Frequency of Numbers of Matches per Obese Claim

Finally, Chart 5 depicts a frequency distribution for the number of unique matches per obese claim. The multiple matches shown in Chart 4 are not necessarily unique in the sense that they may have different durations (but are otherwise identical). In Chart 5, a match is unique if there is no other (non-obese) claim that is identical in all respects (disregarding the type of censoring). For instance, two claims with durations of two weeks, one of them being right-censored, do not count as only one observation in the chart. Chart 5 indicates that in the vast majority of cases, an obese claim has a unique non-obese match.

Chart 5: Relative Frequency of Numbers of Unique Matches per Obese Claim

3. THE STATISTICAL ANALYSIS

Conceptually, the statistical problem of quantifying the effect of obesity on the duration of indemnity benefits is one of survival, that is, continuation of benefits. The hazard in this survival framework manifests itself in the termination of benefits. In a first step, the data is analyzed by means of Kaplan-Meier plots. Following this nonparametric analysis, a semiparametric Weibull proportional hazard model is estimated. In both analyses, the obese claims enter the statistical model along with their non-obese match. Also, both statistical frameworks account for interval-censoring and right-censoring. Finally, in both analytical settings, benefit duration is measured in calendar time.

3.1 Kaplan-Meier Plots

The Kaplan-Meier estimator is a simple way of computing a survival curve, which is a graph that depicts the survival experience (i.e., continuation) of indemnity benefit payments by claim. The percentage of surviving claims is displayed on the vertical axis. The horizontal axis depicts calendar time since injury.

Chart 6 presents the Kaplan-Meier plots for duration concept I. The vertical axis of the plot indicates the proportion of claims that survive any given day displayed on the timeline (which is on

the log10 scale). The reading on the vertical axis of the starting point of the survival curve indicates the proportion of claims that do not breach the waiting period. As shown, only about 25 percent of obese claims do not breach the waiting period, whereas more than 50 percent of the non-obese claims do not breach the waiting period.

The waiting period varies by state, from a minimum of three days to a maximum of seven days. After translating lost workdays into calendar time (see Appendix 5.4), the upper bound of this interval grows to 11 days. This is because a worker whose regular workweek runs from Monday through Friday, and whose first missed workday is a Friday, will have missed eleven calendar days by the time he will have missed seven workdays. In the Kaplan-Meier plot, during the waiting period, the hazard of the affected, interval-censored claims is constant.

The whiskers at the bottom of the Kaplan-Meier plots indicate the location of non-censored claims ("Observed") and censored claims ("Censored"); censored claims may be either interval-censored (in which case the whiskers indicates the end of the waiting period) or right-censored (at the waiting period or later).

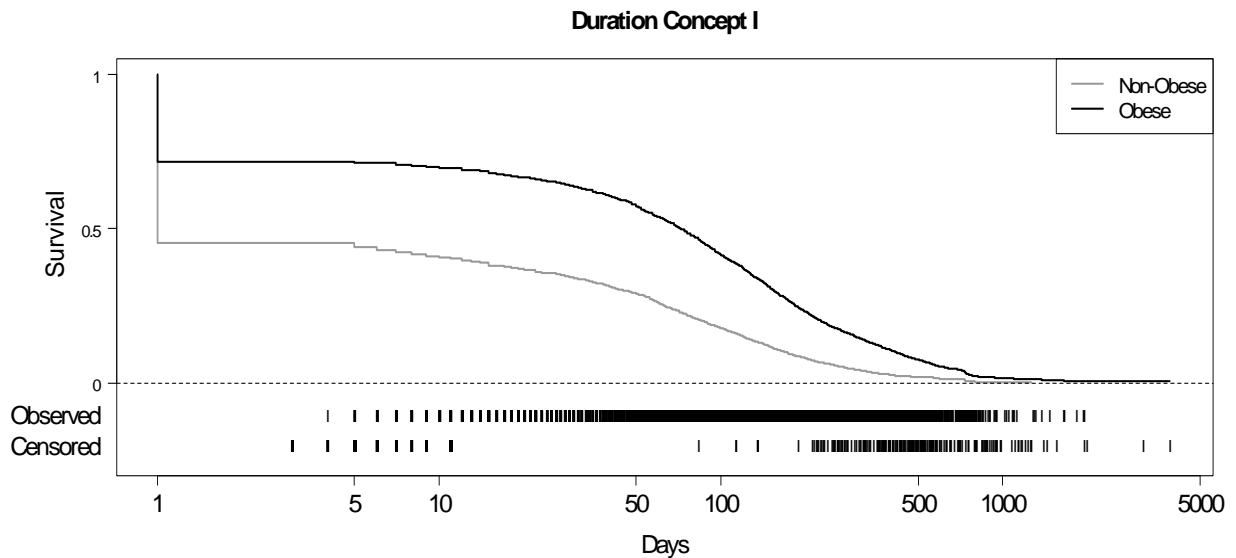**Chart 6:** Kaplan-Meier Plot for Duration Concept I

Chart 6 shows that for any duration displayed on the horizontal axis, the proportion of claimants still collecting indemnity benefits is considerably higher for the obese than for the non-obese. To

the degree that obesity correlates with age, the vertical distance between the survival curves may be distorted slightly by the matching-toward-the-center effect.

Chart 7 displays the survival curves for duration concept II. These curves look very similar to those obtained using duration concept I.

Chart 7: Kaplan-Meier Plot for Duration Concept II

In an effort to ascertain the impact that the immortal time bias would have on the results if left uncorrected, we applied the Kaplan-Meier estimator to the data without the immortal time bias correction (ITBC). Chart 8 displays (for duration concept I) the survival curves with and without ITBC. As mentioned, the ITBC correction reduces the data set to 6,435 obese claims and 14,929 non-obese claims. Without the ITBC, the number of obese claims is slightly larger at 6,528; by comparison, there are 19,397 matching non-obese claims. Not surprisingly, the Kaplan-Meier estimators for the two sets of obese claims are nearly identical to the point of being indistinguishable in the chart. For the non-obese matches, however, the survival curves are distinctly different and, considering the log10 scale, this difference is nonnegligible.

Chart 8: Kaplan-Meier Plot for Duration Concept I, Immortal Time Bias Correction

3.2 Weibull Proportional Hazard Model

The purpose of the proportional hazard model is to quantify the effect of obesity on the mean (and median) duration while controlling for the influence of age. Controlling for age is necessitated by the matching-toward-the-center effect.

A proportional hazard model implies that the hazard in one group (e.g., the obese) is, on the timeline (calendar time since injury), a constant proportion of the hazard in the other group (e.g., the non-obese); this proportion is known as the hazard ratio. The baseline hazard as a function of time is assumed to follow a Weibull distribution.

As specified, the Weibull proportional hazard model accommodates interval-censoring and right-censoring. Further, the model is semiparametric in that it consists of a parametric component (akin to a standard, linear regression model) and a nonparametric component. The linear component includes only one explanatory variable, which is the indicator variable for obesity. The purpose of the nonparametric component is to capture the potentially nonlinear influence of age.

The model is estimated by means of MCMC, using the sampling platform JAGS in R. The JAGS code is displayed in Appendix 5.5.

Following Dellaportas and Smith [3], the hazard function can be written as

$$\lambda(t; \boldsymbol{X}) = \lambda_0(t) \cdot \exp\left(\boldsymbol{X\beta} + f(z)\right) , \tag{1}$$

where $\lambda_0(t)$ is an unknown function of time and $\boldsymbol{X\beta} + f(z)$ is a semiparametric regression model. The matrix $\boldsymbol{X}$ comprises the covariates at time t that are included in the parametric component of the partially linear model, and z represents the single covariate that enters the nonparametric component. The nonparametric component is represented by a linear spline with knots at ages 25, 35, 45, 55, and 65. For the purpose of parameter identification, the spline is centered on age 45.

For $\lambda_0(t) = \rho t^{\rho-1}$, we obtain the Weibull hazard function

$$\lambda(t; \boldsymbol{X}) = \rho t^{\rho-1} \cdot \exp\left(\boldsymbol{X\beta} + f(z)\right) , \tag{2}$$

where $\rho > 0$ is the shape parameter of the Weibull distribution.

With the cumulative hazard given by $\Lambda(t) = t^\rho$, we can write the Weibull density function as

$$f(t) = \rho t^{\rho-1} \cdot \exp\left(\boldsymbol{X\beta} + f(z) - t^\rho \exp\left(\boldsymbol{X\beta} + f(z)\right)\right) . \tag{3}$$

Repeatedly substituting the scale parameter λ for the generalized semiparametric regression model $\exp\left(\boldsymbol{X\beta} + f(z)\right)$ delivers the Weibull density function and its familiar form:

$$\begin{aligned} f(t) &= \rho t^{\rho-1} \cdot \exp\left(\boldsymbol{X\beta} + f(z) - t^\rho \lambda\right) \\ & \rho t^{\rho-1} \cdot \exp\left(\boldsymbol{X\beta} + f(z)\right) \cdot \exp\left(-t^\rho \lambda\right) \\ & \rho t^{\rho-1} \cdot \lambda \cdot \exp\left(-t^\rho \lambda\right) . \end{aligned} \tag{4}$$

The mean of the Weibull distribution is known to read

$$\left(\frac{1}{\lambda}\right)^{\frac{1}{\rho}} \cdot \Gamma\left(1 + \frac{1}{\rho}\right) , \tag{5}$$

where Γ is the gamma function. Similarly, the median of the Weibull distribution equals

$$\left(\frac{1}{\lambda}\right)^{\frac{1}{\rho}} \cdot (\ln(2))^{\frac{1}{\rho}} . \tag{6}$$

Most interestingly, the ratio of the mean to the median is solely a function of the shape parameter ρ and, thus, independent of the hazard λ and the associated generalized semiparametric regression model $\exp\left(\boldsymbol{X\beta} + f(z)\right)$. By implication, any multiplicative impact of a covariate on the mean equals the multiplicative impact on the median.

The effect of obesity is captured by an indicator variable and manifests itself in a change in intercept; this indicator variable is the only covariate in the parametric part of the partially linear regression model.

Let $\hat{\beta}_1$ be the estimated intercept associated with obese claims, and let $\hat{\beta}_0$ be the corresponding value for non-obese claims. In this case then, the ratio of mean (and median) durations of obese claims to non-obese claims (or, equivalently, the multiplicative impact of obesity on mean and median durations) equals

$$\frac{\left(\dfrac{1}{\exp\left(\hat{\beta}_1 + \hat{f}(z)\right)}\right)^{\frac{1}{\hat{\rho}}}}{\left(\dfrac{1}{\exp\left(\hat{\beta}_0 + \hat{f}(z)\right)}\right)^{\frac{1}{\hat{\rho}}}} = \exp\left(\hat{\beta}_0 - \hat{\beta}_1\right)^{\frac{1}{\hat{\rho}}} , \tag{7}$$

where $\exp\left(-\left(\hat{\beta}_0 - \hat{\beta}_1\right)\right)$ is the associated hazard ratio.

Chart 9 depicts the multiplicative impact of obesity. As shown in the chart, the mean and median durations of obese claimants are more than five times the average durations of non-obese claimants. When Permanent Partial benefit payments are counted toward indemnity benefit duration as well, this multiple climbs to more than six.

Chart 10 displays the hazard ratio. When excluding Permanent Partial benefit payments, the probability of the indemnity benefits of an obese claim ending on any given day is a little more than half the probability of this event occurring for a non-obese claim. Including Permanent Partial benefit payments in the analysis has little impact on the estimated hazard ratio.

Finally, Chart 11 depicts the effect of age on indemnity duration for the set of matched pairs analyzed in this study. The effect of age is net of the effect of obesity. As mentioned, the spline is centered on age 45 for the purpose of parameter identification; as a result, the credible intervals are relative to a claimant with an age at injury of 45. (For the difference between credible intervals and confidence intervals, see Carlin and Louis [1].) The chart shows that a claimant whose age at injury falls into the age bracket 45 to 65 tends to have indemnity benefit durations that are about twice the indemnity durations of claimants who sustain a workplace injury at age 25. As a caveat, the research design employed in this study aims at discerning the effect of obesity on indemnity benefit duration, not the effect of age. For the purpose of arriving at definitive statements on the effect of age, it is preferable to use a more comprehensive data set, as opposed to the employed matched pairs.

Chart 9: Multiplicative Impact of Obesity on Mean Duration; Exclusive and Inclusive of Permanent Partial Indemnity Benefits; Duration Concept I and II; 80 Percent Credible Intervals
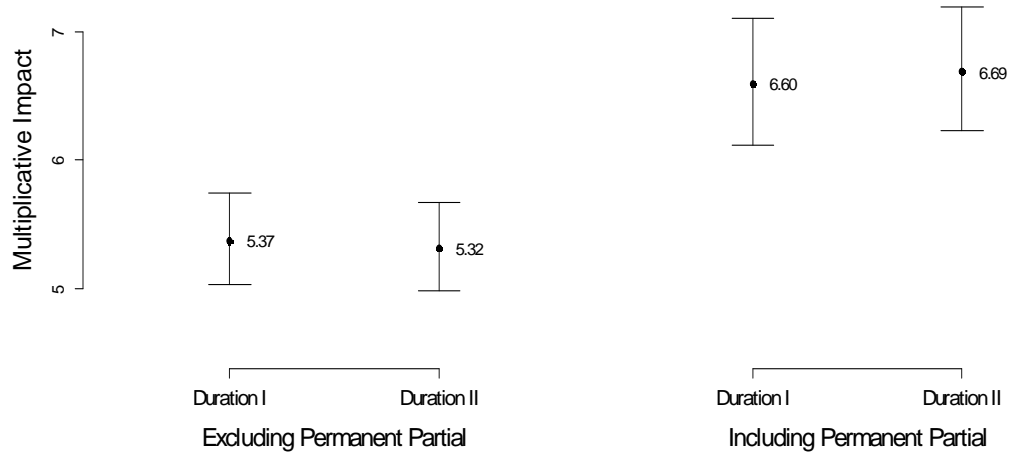

Chart 10: Hazard Ratio; Exclusive and Inclusive of Permanent Partial Indemnity Benefits; Duration Concept I and II; 80 Percent Credible Intervals

Chart 11: Effect of Age on Mean Duration, Duration Concept II, 80 Percent Credible Intervals

4. CONCLUSION

The study finds that obesity contributes in significant ways to the length of time during which claimants receive indemnity benefits. Indemnity duration was measured based on Temporary Total and Permanent Total indemnity benefit payments; in a sensitivity analysis, Permanent Partial benefits were counted toward indemnity benefit duration as well. Two concepts of aggregating observed indemnity benefit transactions into duration were employed, with little difference in the results. Further, the statistical analysis accounted for the presence of interval-censoring and right-censoring, both in the nonparametric framework of Kaplan-Meier plots and in the Bayesian semiparametric proportional hazard model.

The statistical analysis showed that claimants with a comorbidity indicator pointing to obesity have an indemnity benefit duration that is more than five times the value of claimants who do not have this comorbidity indicator but are otherwise comparable. Inclusive of Permanent Partial indemnity payments, this multiple climbs to more than six.

Clearly, the limiting factor in this study is the lack of information on the body mass index of the claimant. On one hand, it can be argued that the analysis overestimates the effect of obesity if the assignment of the comorbidity indicator, the ICD-9 code 278, is related to the arrival of obesity-related medical complications, as opposed to the condition of obesity. Medical complications provide an incentive to the physician to assign this "non-reimbursable" ICD-9 code. From this perspective, a selection bias may be introduced, as a result of which claimants who acquire this comorbidity indicator may disproportionately belong to the highest obesity category—the morbidly obese; this conjecture is supported by the fact that only 0.15 percent of the claims in the data set acquire the ICD-9 code 278 (within 12 months of the date of injury). On the other hand, a case can be made that the effect of obesity is underestimated. This is because many of the claimants who are categorized as non-obese in this study may, in fact, be overweight or obese, thereby diminishing the measured contribution of obesity to duration.

Despite the limitations of the data set employed in this study, the results obtained for the effect of obesity on indemnity duration are close to what has been established by Duke University for the morbidly obese. Based on the reported means in the Duke University study, for the morbidly obese, the number of lost workdays per claim amounts to 6.4 times the value observed for claimants of recommended weight. By comparison, for duration concept I [II], the multiples established here equal 5.4 [5.3] (Temporary Total and Permanent Total) and 6.6 [6.7] (Temporary Total, Permanent Total, and Permanent Partial).

5. APPENDIX

5.1 Simulation for Validating the Immortal Time Bias Correction (ITBC)

A Monte-Carlo simulation was carried out to investigate the validity of the immortal time bias correction described earlier. The simulation maintains the primary features of the underlying data-generation process under the assumptions that, first, the data-generating process for benefit duration is the same for all claims and, second, the probability of acquiring a 278 ICD-9 code is governed solely by the duration of a claim. The simulation assumes that (1) claims only generate ICD-9 codes while they are open, (2) claims are only considered open over their indemnity benefit duration, (3) there is a fixed probability, π, that on any day that a claim is open, it will produce a 278 ICD-9 code, and (4) the acquiring of a 278 ICD-9 on day x is independent of the acquiring of a 278 on day y, provided that $x \neq y$. In this simulation, the parameters were calibrated to reflect important characteristics of the data and the estimated statistical model.

The algorithm generates pairs of obese and non-obese claims as follows:

(1) Draw the indemnity benefit duration and the time to the first occurrence of a ICD-9 code 278 for a claim classified as obese

    (a) Randomly generate an indemnity benefit duration from a Weibull distribution and round up to the nearest integer

    (b) Randomly draw the time to the first ICD-9 code 278 from a geometric distribution with parameter π (Note: This step b is independent of step a)

    (c) If the time to the first ICD-9 code 278 is greater than the minimum of the drawn indemnity benefit duration or one year, then reject the draw and return to step 1; otherwise, accept the draw and proceed to step 2

(2) Draw a suitable non-obese claim to pair with the obese claim generated in step 1

    (a) Randomly generate an indemnity benefit duration from a Weibull distribution and round up to the nearest integer

    (b) Randomly draw the time to the first ICD-9 code 278 from a geometric distribution with parameter π (Note: This step b is independent of step a)

    (c) If the time to the first ICD-9 code 278 is less than the minimum of the drawn indemnity benefit duration or one year, then reject the draw and return to step 2; otherwise, accept the draw and return the obese/non-obese pair

Under the aforementioned simulation, the average duration for claims flagged as obese is 98.5 days; the average duration for claims flagged as non-obese is 98.2 days.

Carrying out the simulation without the immortal time bias correction is equivalent to all generated claims passing step 2c. Without the immortal time bias correction, the average duration for claims flagged as obese is 98.5 days and the average duration for claims flagged as non-obese is 9.5 days.

The simulation result demonstrates that the proposed correction removes the bulk of the immortal time bias.

5.2 Logic of Establishing Primary ICD-9 Code

The primary ICD-9 code of a given claim is the first ICD-9 code associated with reported medical transactions (using a valid service date) based on the following order of priority:

(1) Paid amount must be greater than zero

(2) Provider Group must be one of the following four:

> 01 (Physician), 04 (Hospital, Ambulatory Surgical Center, X-Ray, Lab), 05 (Clinic), 07 (Non-Medical)

When no such transaction exists, transactions from all other providers are considered:

> 02 (Chiropractor), 03 (Therapist), 06 (Pharmacies and Durable Medical Equipment Center), 08 (Other)

(3) Transactions are selected from the first remaining

(4) service date, subject to data cleansing considerations

(5) From the ICD-9 codes resulting from steps 1 through 3, the ICD-9 codes associated with the highest paid amount over the (heretofore recorded) life of the claim are selected

> The paid amount of a transaction is assigned in full to any ICD-9 code associated with the transaction

(5) If there are ties after applying the mentioned criteria, an ICD-9 code is randomly selected from the set of tied codes

5.3 Determining the Median Duration among Multiple Non-Obese Matches

Assume that an obese claim has multiple matching non-obese claims. In this case, the matching non-obese claim assumes the median duration of the set of multiple matches. This median duration is computed by first sorting the matching claims by duration. Then, in the event that the number of matching claims is odd, all claims are identified that have the same duration as the central value. If more than one such claim exists, a choice is made among these claims based on the following hierarchy, listed in descending order: right-censored, non-censored, interval-censored. In the event that the number of matching claims is even, the two central values are identified. Then, all claims with a duration that equals the longest duration of these two values are identified. If more than one such claim exists, a choice is made among these claims based on the following hierarchy, listed in descending order: right-censored, non-censored, interval-censored.

5.4 Converting the Waiting Period into Calendar Time

Waiting periods are stated in numbers of workdays. These workdays were converted into calendar time by assuming employees work five consecutive days a week. The conversion algorithm reads as follows:

(1) In the event that the day of injury is between (and inclusive of) Monday and Friday, the days off work are Saturday and Sunday; else, it is Tuesday and Wednesday

(2) Determine the shortest vector of consecutive days of the week, starting with the day of injury, such that the length of the vector, excluding the days off work, equals the waiting period

5.5 JAGS Code (Weibull Proportional Hazard Model)

```
## Defining is.censored (from JAGS Manual)
## X ~ dinterval(failure.time, waiting.period.or.last.observed)
## Left Censored X = 0; Right Censored X = 1; Not Censored X = NA
model
{
for(i.obs in 1:n.obs){
   failure.time[i.obs] ~ dweib(r,mu[i.obs])
   is.censored[i.obs] ~ dinterval(failure.time[i.obs], waiting.period.or.last.observed[i.obs])

    mu[i.obs] <- exp(alpha[is.obese[i.obs]+1] + f[i.obs])

    f[i.obs] <- age[i.obs] * beta[1] +
        (age[i.obs] >= knot[1])*(age[i.obs] - knot[1])*b[1] +
            (age[i.obs] >= knot[2])*(age[i.obs] - knot[2])*b[2] +
                (age[i.obs] >= knot[3])*(age[i.obs] - knot[3])*b[3] +
                    (age[i.obs] >= knot[4])*(age[i.obs] - knot[4])*b[4] +
                        (age[i.obs] >= knot[5])*(age[i.obs] - knot[5])*b[5]

    median[i.obs] <- (log(2) * 1/mu[i.obs])^(1/r)
}

##Spline
##Random regression coefficients corresponding to the truncated polynomial functions
for(i.knot in 1:n.knots){
    b[i.knot] ~ dnorm(0, tau.b)
}

##Fixed regression coefficient corresponding to the 'plus' functions
beta ~ dnorm(0,1.0E-2)

##Priors
r ~ dexp(1)
alpha[1] ~ dnorm(0,0.01) ##Not Obese
alpha[2] ~ dnorm(0,0.01) ##Obese

tau.b ~ dgamma(1.0E-3,1.0E-3)
sigma.b <- 1/sqrt(tau.b)

obese.control <- alpha[2] - alpha[1]
}
```

Acknowledgment

6. REFERENCES

[1]   Carlin, Bradley P., and Thomas A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis.* 2nd ed., Boca Raton: Chapman & Hall/CRC, 2000.

[2]   Kasabov, Nicola K., *Foundations of Neural Networks, Fuzzy Systems, and Knowledge Engineering*, Cambridge, MA: MIT Press, 1996.

[3]   Dellaportas, P., and Adrian F.M. Smith, "Bayesian Inference for Generalized Linear and Proportional Hazards Models via Gibbs Sampling," *Applied Statistics* **42**(3), 1993, pp. 443–459.

[4]   Laws, Chris, and Frank Schmid, "Reserving in the Age of Obesity," November 1, 2010, www.ncci.com/documents/research-age-of-obesity.pdf.

[5]   Østbye, Truls, John M. Dement, and Katrina M. Krause, "Obesity and Workers' Compensation: Results From the Duke Health and Safety Surveillance System," *Archives of Internal Medicine* **167**, 2007, pp. 766–773.

[6]   Pollack, Keshia M., and Lawrence J. Cheskin, "Obesity and Workplace Traumatic Injury: Does the Science Support the Link?" *Injury Prevention* **13**(5), 2007, pp. 297–302.

[7]   Pollack, Keshia M., Gary S. Sorock, Martin D. Slade, Linda Cantley, Kanta Sircar, Oyebode Taiwo, and Mark R. Cullen, "Association Between Body Mass Index and Acute Traumatic Workplace Injury in Hourly Manufacturing Employees" (with discussion), *American Journal of Epidemiology* **166**, 2007/2008, pp. 204-211 & **167**, pp. 123-124.

[8]   Shaw, William S., Torill H. Tveito, Mary Jane Woiszwillo, and Glenn Pransky, "The Effect of Body Mass Index on Recovery and Return to Work Following Onset of Work-Related Low Back Pain," *Journal of Occupational and Environmental Medicine* **54**(2), 2012, pp. 192–197.

[9]   Witters, Dan, and Sangeeta Agrawal, "Unhealthy U.S. Workers' Absenteeism Costs 153 Billion," Gallup, October 17, 2011, www.gallup.com/poll/150026/unhealthy-workers-absenteeism-costs-153-billion.aspx.

Abbreviations and notations

BMI, Body Mass Index
FTE, Full-Time Equivalent
HRA, Health Risk Assessment
ITBC, Immortal Time Bias Correction
ICD-9, International Classification of Diseases, Ninth Revision
LBP, Low back pain
MCMC, Markov chain Monte Carlo simulation
NCCI, National Council on Compensation Insurance

Biographies of the Authors

Frank Schmid, Dr. habil., was, at the time of writing, a Director and Senior Economist at the National Council on Compensation Insurance, Inc.
C**hris Laws** is a Research Consultant at the National Council on Compensation Insurance, Inc.
Matthew Montero is a Senior Actuarial Analyst at the National Council on Compensation Insurance, Inc.