Quantifying Financial Risk in Employer Self-Funded Health Insurance: A Monte Carlo Simulation Analysis

Leia McIntosh[i] Dr. Amanda Cook[ii], Dr. Tice Sirmans[iii]

[i]Student - Master of Science in Applied Statistics Bowling Green State University - Department of Applied Statistics & Operations Research, [ii]Faculty - Bowling Green State University Department of Economics , [iii]Faculty - Illinois State University Department of Finance, Insurance, and Law

INTRODUCTION

53% of Americans receive health insurance as an employee benefit, and 65% of insured employees are covered under a self-funded employer plan.[1] In self-funded arrangements, the financial risk of severe health shocks among employees is retained by employers. Low-probability, high-cost health events amplify this risk.

For a binomial distribution, it is a statistical "rule of thumb" that if the expected number of illnesses (n*p) exceeds 10, the distribution can be approximated by a normal distribution, where n is the number of trials (firm size) and p is the probability of success (illness prevalence). However, with low-probability illnesses, n must be sufficiently large for the normal distribution approximation to apply, and to allow for employers to bear the risk effectively. With an illness prevalence of 1%, firms must have at least 1,000 employees to take advantage of this predictability. While three-quarters of large firms of 500 employees or more are self-insured for health benefits, over 99.2% percent of U.S. firms have less than 1,000 employees.[2][3] This makes it challenging for a vast majority of firms to prepare for the financial consequences of employee health shocks.

In this study, we aim to quantify the financial risks faced by employers in self-funded health insurance plans, exploring how variations in illness prevalence, care costs, and firm size influence employers' ability to predict and manage severe healthcare expenses.

METHODS

To examine the variation between realized and expected per-person healthcare expenditures by varying illness prevalence, cost of care, and firm size, we employed Monte Carlo simulations. We considered variable firm sizes between 10 and 1,800 employees in three distinct models. In each model, the illness prevalence and cost of care parameters differed. The BG Data (Mild) and BG Data (Severe) models were calibrated to the most recently available Bowling Green State University health expenditure data, and the Theoretical model was calibrated to represent a low-probability high-cost health event.

For each model, we simulated 1,000 firms of each firm size. Within each firm, individual healthcare costs were determined based on the assigned illness probabilities and costs of care. These individual costs were aggregated to calculate the total firm expenditure, which was then divided by the number of employees to calculate the average per-person cost for a single firm. This process was repeated across all firms to generate a distribution of average per-person costs for each firm size. This procedure was carried out separately for each of the three models.

First, we visualized the relationship between firm size (n) and expected number of illnesses (n*p) across varying illness probabilities, highlighting where the expected number of illnesses is equal to 10 for each probability. Then, to assess the variation in between distributions, we analyzed the ratio of the standard deviation to the mean to quantify the spread relative to the expected value. Finally, we examined the difference between observed costs and expected costs at both the individual firm level and averaged across a small sample of firms. These quantities were visualized using bar charts and beeswarm plots.

RESULTS

Figure 1: Relationship Between Firm Size (n) and Expected Number of Illnesses Per Firm (n*p)

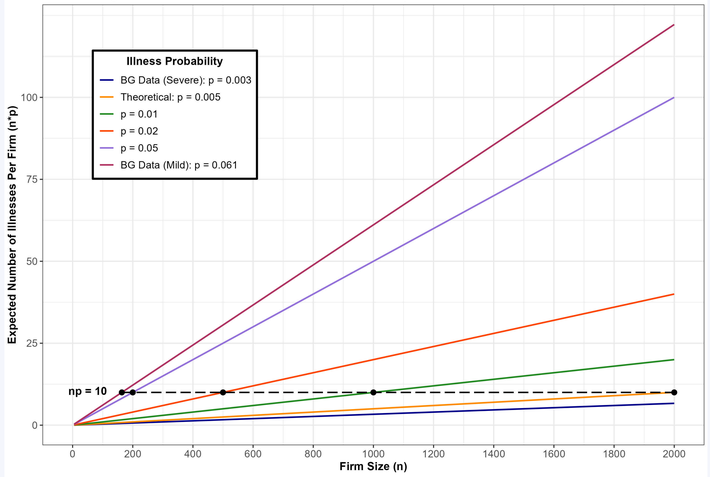

Figure 2: Ratio of Sample Standard Deviation to Sample Mean by Firm Size (Sample Size: 1,000 Firms)

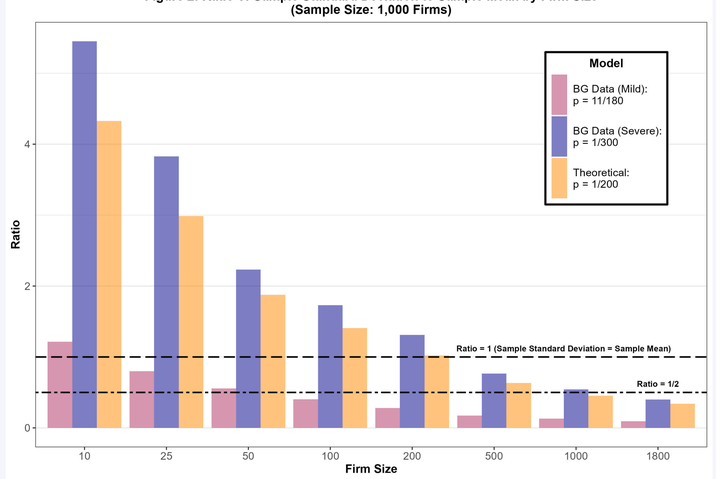

Figure 3: Differences Between Expected and Observed Annual Per-Person Costs by Firm Size (Sample Size: 5 Firms)

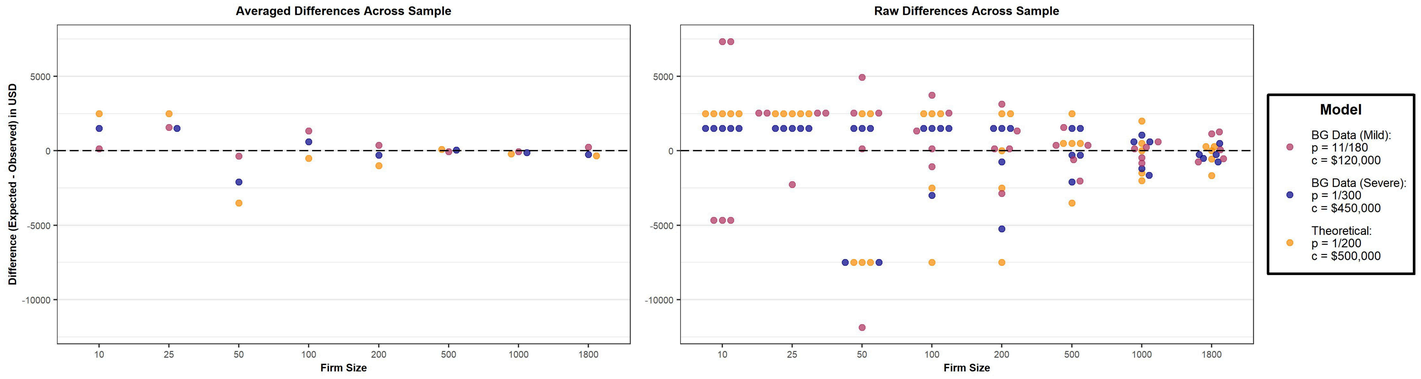

KEY FINDINGS

Figure 1 illustrates the theoretical relationship between minimum firm size and various illness probabilities (i.e., the expected number of illnesses), highlighting the point at which using the normal distribution becomes an appropriate approximation. The plot includes probabilities for our three models, along with a few benchmark values for comparison. The figure shows that firms must often be extremely large for the normal approximation to hold, especially for rare illnesses. Notably, BGSU, which has approximately 1,800 employees on its healthcare plan, is not large enough to reasonably predict and account for healthcare expenses related to the BG (Severe) illness probability.

Figure 2 highlights the importance of predictability in financial risk assessment by showing how the standard-deviation-to-mean ratio varies across firm sizes for each model. Predicting financial risk becomes manageable when this ratio is ½ or smaller, allowing an approximate 95% confidence interval to determine whether the true mean per-person cost is significantly different from zero. BG Data (Severe) takes the longest to reach this standard, crossing the ratio = ½ threshold at 1,800 employees, followed by Theoretical at 1,000 employees. In contrast, BG Data (Mild) reaches this threshold at only 100 employees. Although it is expected that rarer illnesses are harder to predict financially, this visual quantifies how unreliable the expected value becomes as illness probability decreases and, in turn, variation increases.

Figure 3 plots the difference between expected and observed health expenditures. On the left, we represent the average deviations between anticipated health expenditures and the observed expenditures across five-firm samples. On the right, we plot the individual deviations for each of the five firms. The left plot suggests that averaging across five firms smooths out variation, giving a false sense of predictability. From this plot, firms approximately stabilize to a difference of zero (perfectly predicted health expenditures) at a firm size of 200. The right plot reveals significant fluctuations, with many firms facing observed costs severely exceeding predicted costs. Additionally, the right plot shows that firms do not stabilize around zero until reaching closer to 1,000 employees. This highlights a common misconception: firms often focus on averages and expected values, when, in reality, individual firm costs remain highly variable.

SIGNIFICANCE

It is crucial to understand the practical implications of illness predictability on firms. As discussed previously, 99.2% of firms in the United States have fewer than 1,000 employees, and with 1,000 employees, a firm can reasonably predict an illness occurring at a 1% probability. Predicting illnesses with less than 1% prevalence therefore requires increasingly larger firms. For example, consider the nation's leading cause of death: heart disease. In 2023, the U.S. saw 680,981 deaths from heart disease, with a population of approximately 338 million (per CDC data).[4][5] Based on this data, heart disease deaths occurred at a rate of 0.201% in the general population. Using a normal approximation, a firm would need 4,976 employees to reasonably predict and bear this risk. Similarly, BGSU, which is considered to be a very large employer, is still not mathematically large enough to account for its realized severe illness prevalence in recent years.

Additionally, from an individual firm standpoint, even if a five-year average is able to smooth out health shocks, it's important to recognize that approximately 50% of private-sector businesses fail within that timeframe.[6] Extreme yearly fluctuations in health expenditures, consequently, can have lasting impacts on business survival. The moral dilemma of choosing to prioritize employee coverage or personal/business expenditures is one that a business owner should never be forced to face. To further illustrate this concept, Figure 3 may also be interpreted as representing five business years for the same firm.

In summary, firms are often sold self-insuring plans under the premise that they are "sufficiently large" and are therefore able to mitigate the financial risk of severe employee health shocks. In this study, we quantified when firms can and cannot effectively manage risk, relating illness prevalence to firm size. Additionally, we demonstrated why the concept of being "sufficiently large" is a widely inaccurate metric, with the ability to quickly and severely impact a firm's financial stability and success.

REFERENCES

[1] Kaiser Family Foundation. (2023, October 18). 2023 *Employer Health Benefits Survey – Section 10: Plan funding.* https://www.kff.org/report-section/ehbs-2023-section-10-plan-funding/

[2] U.S. Bureau of Labor Statistics. (2024, March). *Number of business establishments by size of establishment in selected private industries.* https://www.bls.gov/charts/county-employment-and-wages/establishments-by-size.htm

[3] Employee Benefit Research Institute. (2024, August 9). *New research finds increasing number of self-insured health plans in small and medium-sized businesses but a decreasing number in large companies since passage of the Affordable Care Act of 2010.* https://www.ebri.org/content/new-research-finds-increasing-number-of-self-insured-health-plans-in-small-and-medium-sized-businesses-but-a-decreasing-number-in-large-companies-since-passage-of-the-affordable-care-act-of-2010

[4] National Center for Health Statistics. (2024, December). *Death rates for leading causes of death, 2021.* U.S. Centers for Disease Control and Prevention. https://www.cdc.gov/nchs/products/databriefs/db521.htm#fig4

[5] U.S. Census Bureau. (n.d.). *U.S. and world population clock.* Retrieved March 1, 2024, from https://www.census.gov/popclock/

[6] U.S. Bureau of Labor Statistics. (2024, March). *Table 7. Survival of private sector establishments by opening year, total private* [Data file]. https://www.bls.gov/bdm/us_age_naics_00_table7.txt