ADVANCED ANALYTICAL
CONSULTING GROUP

BANKRUPTCY AND MASS WITHDRAWAL MODELING IN PIMS

October 1, 2019

Advanced Analytical Consulting Group, Inc.

Constantijn W.A. Panis, PhD
213-784-6400
stanpanis@aacg.com

Karthik Padmanabhan, MS, MBA
312-551-9001
karpad@aacg.com

SUMMARY

The Pension Benefit Guaranty Corporation (PBGC) insures millions of participants in private defined benefit (DB) pension plans against loss of some or all benefits in case their plan is unable to pay benefits. The PBGC uses a stochastic modeling system, the Pension Insurance Modeling System (PIMS), to project its future expected claims. PIMS consists of a version for single-employer plans (SE-PIMS) and a version for multiemployer plans (ME-PIMS).

The bankruptcy of the sponsor of a single-employer plan is an important trigger of plan failures, and thus of additional PBGC liabilities. Likewise, the mass withdrawal of all employers that participate in a multiemployer plan can cause or accelerate a plan's insolvency. This report reviews bankruptcy aspects of SE-PIMS and mass withdrawal aspects of ME-PIMS. Our interim reports on these two components identified numerous topics for potential review; this report focuses on priority areas.

For the bankruptcy aspects of SE-PIMS, we generally confirmed the PBGC's current and updated models for predicting bankruptcies. We found that bankruptcy rates of private and public firms are roughly equal, as assumed by SE-PIMS. Some issues were identified with the stability of SE-PIMS's future stochastic scenarios, but the distribution of bankruptcy probabilities—the key outcome—was remarkably stable. SE-PIMS assumes that the PBGC will recover, in bankruptcy proceedings, 5% of unfunded liabilities for both Chapter 11 reorganizations and Chapter 7 liquidations; we noted that historical recoveries tended to exceed 5% but found no compelling reason to adjust the recovery parameters. Finally, SE-PIMS assumes that a bankruptcy filing does not trigger a plan failure if the plan's assets are 80% or more of its liabilities; we found little historical support for this assumption, but the available data were incomplete and preclude us from recommending a change.

ME-PIMS predicts mass withdrawals through a complex set of equations involving numerous determinants. While the determinants plausibly relate to mass withdrawal decisions, we found no theoretical or empirical support for the model, except that the average predicted mass withdrawal rate is close to the historical rate. We analyzed withdrawals by individual employers, which ME-PIMS does not currently model. These withdrawals, which are concentrated in certain industries and poorly funded plans, are not insignificant and may warrant inclusion in ME-PIMS. ME-PIMS makes assumptions about the collectability of withdrawal liabilities; we attempted to validate those with historical collections, but were unable to because of data constraints. Finally, up until very recently, ME-PIMS simulated a subset of multiemployer plans; we explored potential issues with modeling the universe instead, and found the recent expansion of ME-PIMS to be sensible and scientifically sound.

CONTENTS

1. INTRODUCTION

In order to monitor the solvency of its insurance operations, the Policy, Research and Analysis Department (PRAD) of the Pension Benefit Guaranty Corporation (PBGC) developed a system of connected models—the Pension Insurance Modeling System (PIMS)—to forecast revenues, expenses, and related metrics. The PBGC contracted with Advanced Analytical Consulting Group, Inc. (AACG) to review bankruptcy aspects of the Single-Employer Pension Insurance Modeling System (SE-PIMS) and mass withdrawal aspects of its multiemployer counterpart, ME-PIMS. This document reports on our review of specific priority areas.

In SE-PIMS, we evaluated five aspects related to the modeling of potential bankruptcies of defined benefit (DB) plans. They are:

1. Review of the integrity of bankruptcy model estimates;
2. Analysis of simulated distributions of key variables for reasonableness and appropriate correlations;
3. Evaluation of whether bankruptcy rates of large private firms are similar to those of large public firms;
4. Analysis of historical bankruptcy recovery rates under Chapter 7 and 11 filings, and examine repeat bankruptcies; and
5. Evaluation of the assumptions triggering plan termination in bankruptcy.

In ME-PIMS, we evaluated four aspects related to the modeling of mass withdrawals of multiemployer plans. They are:

1. Review of the current model for unnecessary complexity and potentially suggest an alternative;
2. Investigation of whether withdrawals by individual contributing employers can be incorporated into ME-PIMS;
3. Evaluation of assumptions on the collectability of withdrawal liabilities; and
4. Identification of potential issues with modeling the entire population of multiemployer plans rather than a sample.

Section 2 of this report addresses bankruptcy aspects in SE-PIMS. Section 3 discusses mass withdrawal areas in ME-PIMS, and Section 4 concludes.

2. BANKRUPTCY MODELING IN SE-PIMS

Overview

According to recent Form 5500 filings, roughly 25,000 firms sponsor a PBGC-insured single-employer DB plan. While the Form 5500 contains financial information about these DB plans, they contain very little information on financial metrics that can predict the risk of bankruptcy of the sponsoring firms. A small subset of sponsors—mostly companies that have publicly traded equity or debt—publish their financial information in annual reports, corporate filings with the Securities and Exchange Commission (SEC), and other channels. Compustat culls from such sources and bundles financial records into a database. The PBGC merged that database with sources on bankruptcy filings, applied various inclusion criteria, and estimated a bankruptcy model. This process is reviewed below in the section on "Integrity of the Bankruptcy Model Estimates," starting on page 5.

SE-PIMS uses the bankruptcy model's parameter estimates to simulate future bankruptcy probabilities. Ideally, it would simulate bankruptcy probabilities for all sponsors of PBGC-insured DB plans, but it lacks the required financial information on most plan sponsors. Instead, it uses scaled copies ("partners") of firms for which financial information is available to represent firms without financial information. This approach implicitly assumes that small and private firms face approximately equal bankruptcy risks as large firms with publicly available financial records. This assumption is evaluated below in the section on "Bankruptcy Rates of Large Private Firms," starting on page 7.

The unit of observation in the bankruptcy component of SE-PIMS is a firm (DB plan sponsor). SE-PIMS simulates future values of assets, liabilities, sponsor bankruptcy filings, and other financial metrics for 1,434 firms (350 firms plus 1,084 partners). Financial metrics are a function of, among others, macroeconomic variables such as rates of return on stocks and bonds. SE-PIMS generates 5,000 stochastic scenarios for its key financial metrics—500 stochastic scenarios of macroeconomic variables, and 10 firm-specific scenarios for every macroeconomic scenario. The simulation horizon is 20 years, i.e., each scenario consists of 20 future annual values. We explore the simulated distributions, correlations, and time series properties in the section on "Simulated Distributions of Key Variables," starting on page 10.

For purposes of SE-PIMS, a bankruptcy is defined as the filing for protection under Chapter 7 or 11 of the U.S. Bankruptcy Code. By default, SE-PIMS assumes that the consequences for pension plans of a Chapter 7 liquidation are the same as for a Chapter 11 reorganization. That assumption is tested in the section on "Recovery Claims by Bankruptcy Type," starting on page 24.

Finally, the bankruptcy of a DB plan sponsor does not necessarily result in termination of the plan. SE-PIMS assumes that the plan will continue to fulfill its obligations despite a bankruptcy of its sponsor if assets cover 80% or more of liabilities. We evaluate that assumption in the section on "Assumptions Triggering Plan Termination in Bankruptcy," starting on page 28.

Integrity of the Bankruptcy Model Estimates

The Current Model

The bankruptcy model used in the current version of SE-PIMS is primarily based on Compustat data with annual corporate financial information, augmented with information on corporate bankruptcies from New Generation Research, Inc. ("New Generation Research"). The analysis took the following approach.

- The sample was restricted to companies that sponsored a DB plan.
- The analysis was restricted to bankruptcy filings in 1980 and later because the U.S. Bankruptcy Code underwent major changes in 1980. The most recent bankruptcy filings in the analysis date to 1998.
- Companies were included only if they reported 500 or more employees in at least two consecutive years, and only starting with the two years in which they first met that criterion, even if their workforce later shrank below 500 people.
- The analysis does not incorporate repeat bankruptcies; records after the first bankruptcy since 1980 are excluded from the analysis.
- Companies incorporated outside the United States are excluded from estimation.
- Companies that are subsidiaries of other companies are excluded from the analysis.
- A weighting scheme is applied to correct for the PBGC's finding that records are disproportionately often missing when companies approach bankruptcy.
- The model is a logistic regression model to explain the annual incidence of bankruptcies.

AACG replicated the PBGC's procedures and encountered a few issues, including the following:

- The timing of bankruptcies was not always correct in the analysis file. For example, a company that (according to a manual search) filed for bankruptcy in 1986 was included as filing for bankruptcy in 1997. The erroneous timing may have affected parameter estimates because explanatory variables changed over time. For example, the workforce exceeded 6,000 people up to 1984 and was reported as zero people during the 1990s.
- The data contained gaps in companies' time series that did not appear to be addressed by the weighting scheme.
- The weighting scheme accounts for employment category, which was imputed if missing. Under certain circumstances, that imputation appeared to be incorrectly implemented.

We did not attempt to quantify the implications of these issues. However, the next subsection suggests that the introduced biases was negligible.

We confirmed that the bankruptcy model estimates were applied accurately to the 2018 version of SE-PIMS.

An Updated Model

The current bankruptcy model was estimated on data from 1980 to 1998. During the subsequent 20 years, the average annual number of bankruptcy filings was lower than during the estimation period: the average annual number of corporate bankruptcy filings decreased from 61,640 in 1980–1998 to 36,056 in 1999–2018 (Statistical Abstract of the United States, 1980–2011; Administrative Office of the U.S. Courts, 1997–2018).[1]

Given a potential concern that a model based on 1980–1998 data would overstate recent bankruptcy filings, we analyzed bankruptcy data through 2017.

The PBGC provided us with Compustat data from 1988 through 2017, along with additional information on bankruptcy filings from New Generation Research, Research Insight, and a 2006 Corporate Tracker file. We merged the source files and prepared the data for estimation following the principles described in the previous section. The average annual bankruptcy probability, weighted to correct for non-random patterns of missing data, was 0.84% per year for 1988–2017, compared with 0.90% per year for 1980–1998 (based on the estimation sample of the current model). In other words, the long-term average bankruptcy rate among large DB sponsors has not changed much since 1980–1998, and any concern that the current bankruptcy model substantially overestimates bankruptcy risks appears to be misplaced.

The similarity of bankruptcy *rates* in 1980–1998 and 1988–2017 stands in contrast to the steep decline in *number* of filings after 1998. Part of the explanation lies in the fact that the number of publicly listed companies has decreased markedly, from 9,113 in 1997 to 5,734 in 2016.[2] Also, the number of filings relates to all corporations, whereas the relevant bankruptcy rate applies to the small subset of firms in the bankruptcy analysis, namely large corporations that sponsor a DB plan and publish financial information.

We emphasize that bankruptcy rates are volatile over time (Figure 1). The observation that the average 1980–1998 rate (0.90%) was close to the 1988–2017 average (0.84%) should therefore be interpreted with caution.

---

[1] Because of incomplete historical data on filings under Chapters 7 and 11, these numbers refer to all corporate filings, including those under Chapters 12 or 13.
[2] "America's Roster of Public Companies Is Shrinking Before Our Eyes," *Wall Street Journal*, January 6, 2017. Available at https://www.wsj.com/articles/americas-roster-of-public-companies-is-shrinking-before-our-eyes-1483545879.

Figure 1. Weighted Bankruptcy Rates among Large DB Plan Sponsors, by Year (1988–2017)

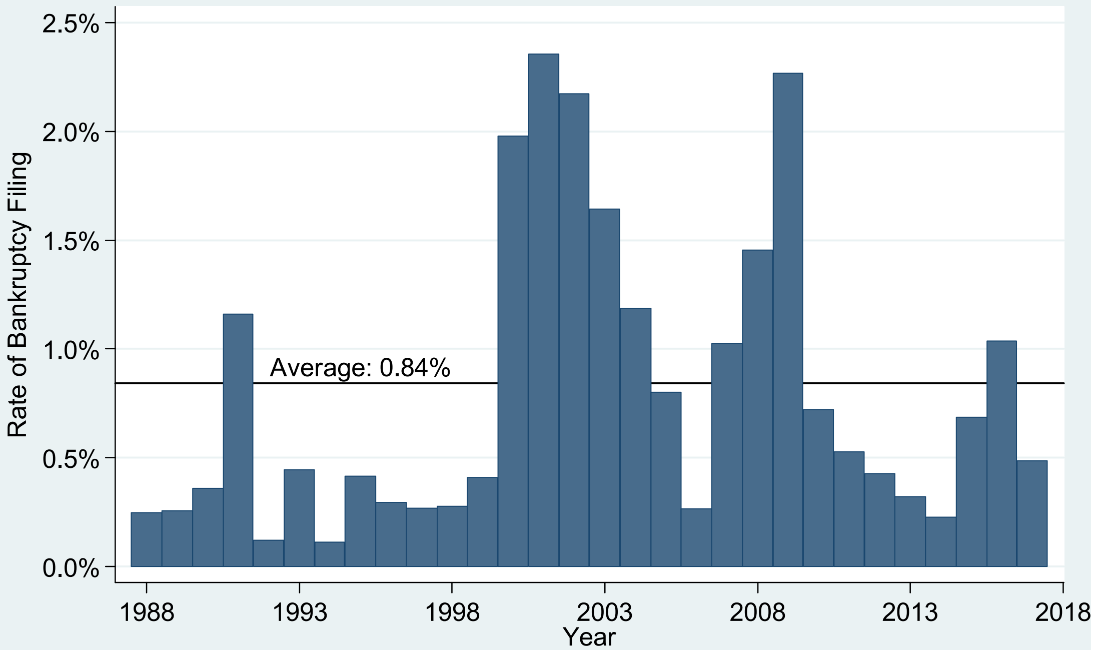

Source: AACG calculations based on Compustat data augmented with bankruptcy information; see text.

For comparison, the *FY 2018 PBGC Projections Report* (Table 4) documented an average bankruptcy probability of 0.5% during the simulation period.

We estimated a logistic regression model of the probability of a bankruptcy based on the 1988–2017 data. The results are qualitatively similar to those estimated on 1980–1998 data and used in the current SE-PIMS. All signs of estimates of parameters that the two specifications have in common are the same, and their statistical significance levels are very similar. Separately, we compared our estimates to those obtained by the PBGC, based on the same source data. The sample and results were very slightly different, presumably mostly because our algorithm to match data sources differs from the PBGC's.[3] For all practical purposes, we confirm the PBGC's updated model.

Bankruptcy Rates of Large Private Firms

The bankruptcy model equation in SE-PIMS was estimated on data for firms (1) that sponsored one or more DB plans, (2) for which financial data were available, and (3) that employed 500 or more workers in at least two consecutive years. Large private firms typically lack published financial data and are thus excluded from the models

[3] The AACG algorithm attempts to increase the match rate by using a normalized version of company names, in addition to the keys used by the PBGC (CUSIP, EIN, and Global Vantage Key). Our final sample consisted of 36,895 records, compared with 36,291 in the PBGC estimation sample.

used to estimate bankruptcy rates in SE-PIMS. We investigated whether large private firms have significantly different bankruptcy probabilities than public firms.

Firms with 500 or More Employees

First, we compared bankruptcy rates for public and private firms with 500 or more employees. Bankruptcy data were sourced from New Generation Research, an industry source on corporate bankruptcies. Around 2012, New Generation Research updated its data collection procedures for private companies that expanded the number of bankruptcies as well the amount of information recorded about the bankruptcy. This allowed us to estimate bankruptcy rates for large private companies, albeit for a recent time period only.

In the case of a company with many subsidiaries, New Generation Research data contains records for the parent company and its subsidiaries. In these cases, we kept and used only one instance (with preference given to a public instance, if present) when calculating the bankruptcy rates.[4]

As Figure 2 shows, the rates for large private firms are roughly in line with bankruptcy rates of large public firms.[5]

[4] For evaluation purposes, we kept and included all instances of bankruptcy records for a company and its subsidiaries and used it to calculate bankruptcy rates. The resulting bankruptcy rates were too high to be plausible.
[5] The denominator used to calculate the private bankruptcy rate is obtained from a firm count maintained by the Census Bureau's Business Dynamic Statistics. A firm is defined as "a business organization consisting of one or more domestic establishments that were specified under common ownership or control. The firm and the establishment are the same for single-establishment firms." See https://www.census.gov/ces/dataproducts/bds/data.html. We subtracted the number of public companies (from Compustat) to obtain the number of private firms.

Figure 2. Bankruptcy Rates: Private and Public Companies with 500 or More Employees (1981–2017)

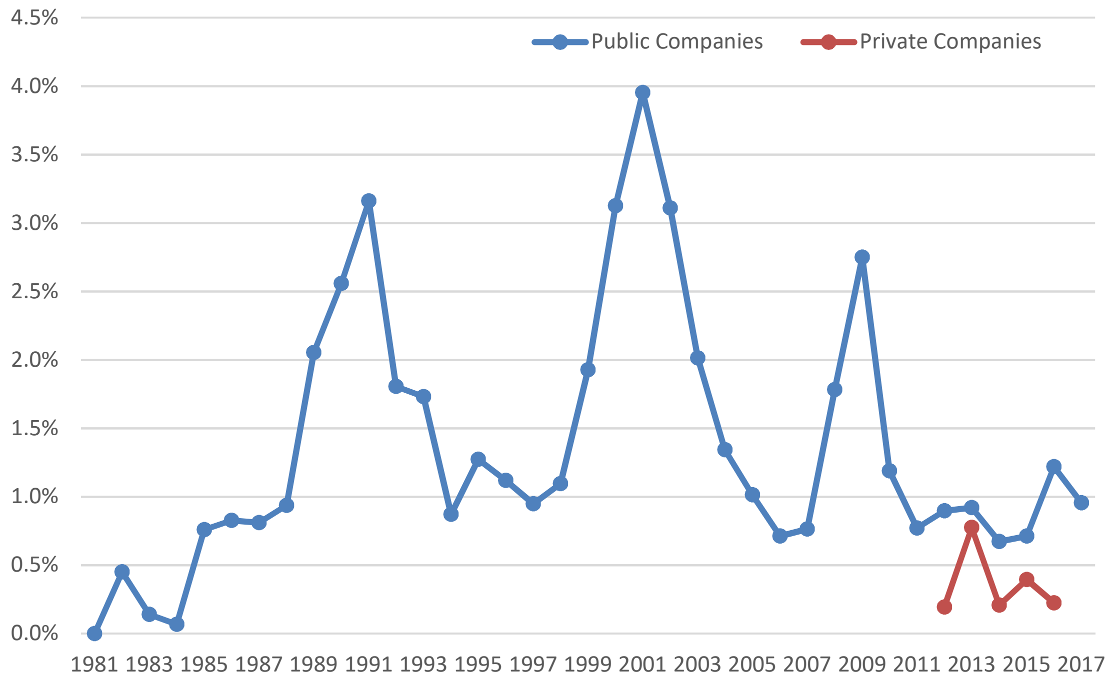

Sources: New Generation Research, Inc., Compustat, U.S. Census Bureau - Business Dynamics Statistics

Firms with 5,000 or More Employees

Next, we compared bankruptcy rates for public and private firms with 5,000 or more employees. As Figure 3 shows, the bankruptcy rates for very large private firms are very close to those for very large public firms.

Figure 3. Bankruptcy Rates: Private and Public Companies with 5,000 or More Employees (1981–2017)

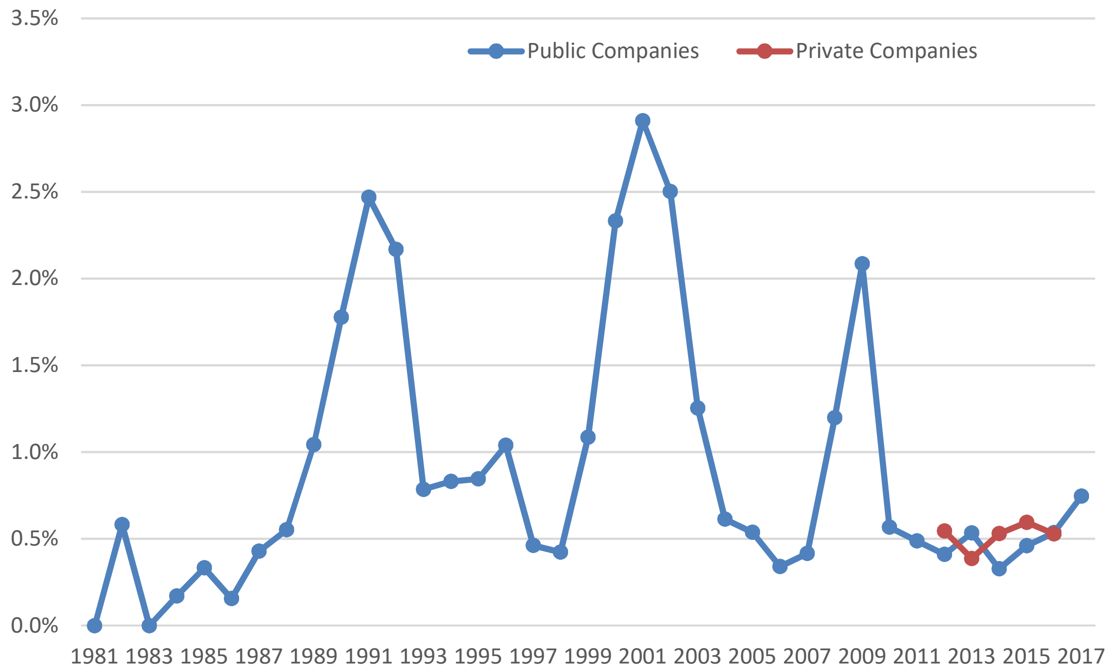

Sources: New Generation Research, Inc., Compustat, U.S. Census Bureau - Business Dynamics Statistics

On the basis of the results discussed above, there is no need to have separate bankruptcy rate assumptions for plans sponsored by private employers. However, the available time series for private firms is short and we recommend this analysis be updated in a few years to ensure bankruptcy rates for large private and public companies do not diverge in the future.

Simulated Distributions of Key Variables

The unit of observation in the bankruptcy component of SE-PIMS is a firm (DB plan sponsor). SE-PIMS simulates future values of assets, liabilities, sponsor bankruptcy filings, and other financial metrics for 1,434 firms (350 firms plus 1,084 partners). Financial metrics are a function of, among others, macroeconomic variables such as rates of return on stocks and bonds. SE-PIMS generates 5,000 stochastic scenarios for its key financial metrics—500 stochastic scenarios of macroeconomic variables, and 10 firm-specific scenarios for every macroeconomic scenario. The simulation horizon is 20 years, i.e., each scenario consists of 20 future annual values. This section summarizes the distributions, correlations, and time series properties of simulated macroeconomic and firm-specific values.

Simulated Macroeconomic Variables

It is our understanding that only the 30-year Treasury yield and the rate of return on stocks are fully stochastic in the economic projections, and that all other values (such as inflation and rates of return on Treasury bonds, corporate bonds, and plan assets) are derived from those two variables.

Figure 4 shows the historical yield on 30-year Treasury bonds from 1977 through 2018 (green line) and the distribution of projected yields through 2038 (blue lines).[6]

Figure 4. Historical and Distribution of Projected Yield on 30-Year Treasury Bonds

Sources: Board of Governors of the Federal Reserve System (historical data), SE-PIMS (projections). Yield on 30-year Treasury bonds (20-year bonds from Feb. 2002 to Feb. 2006).

By construction, projected yields start close to recent historical yields and waver out over time, because (the natural logarithm of) the yield on 30-year government bonds is modeled as a difference from its previous value (*SE-PIMS System Description*, page 2-9 and footnote 1 on page 5-13):

$$\Delta ln(i_t) = \alpha_i + \beta_i \Delta ln(i_{t-1}) + \varepsilon_{i,t},$$

where both $\alpha_i = 0$ and $\beta_i = 0$ in simulations. We regressed simulated values of $\Delta ln(i_{t-1})$ on $\Delta ln(i_t)$ and confirmed that neither the estimated α_i nor β_i is statistically distinguishable from zero in SE-PIMS simulations. With $\alpha_i = \beta_i = 0$, the log-yield follows a random walk.

[6] Historical data are monthly, projected data are annual. Percentiles of projected yields are calculated within each year, i.e., the depicted blue lines are not illustrative of any single scenario. From February 2002 through February 2006, the U.S. Treasury did not sell any bonds with a 30-year maturity; instead, the figure depicts the yield on 20-year bonds for that period.

At the time these simulations were generated, the yield on 30-year Treasury bonds was near historic lows. One may be concerned that yields drift even lower, into unchartered territory, in a sizable number of scenarios. Indeed, yields have in fact decreased further—for example, to 1.98% on August 15, 2019—and they are lower yet in some other developed economies—for example, the German government issued 30-year bonds at a negative yield in August 2019.[7] The current generating equation (random walk in log-yields) does not support negative yields.

Few investors purchase a newly issued 30-year bond and hold it until maturity. To represent annual bond returns, SE-PIMS converts yields into one-year returns—see Figure 5.[8]

Figure 5. Historical and Distribution of Projected Annual Returns on 30-Year Treasury Bonds

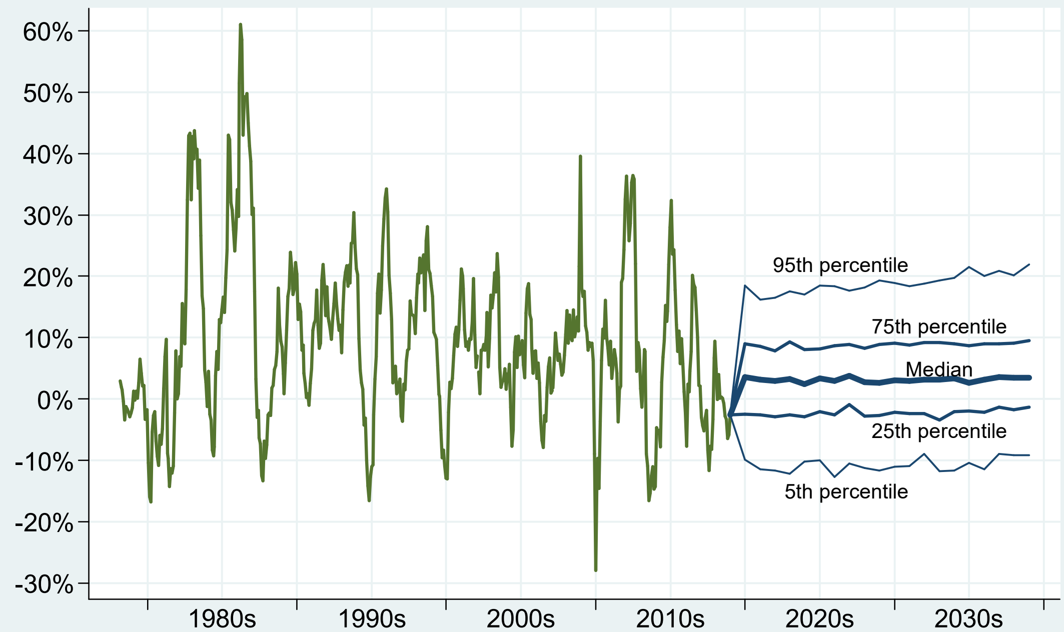

Sources: Board of Governors of the Federal Reserve System (historical data), SE-PIMS (projections).
Return on 30-year Treasury bonds (20-year bonds from Feb. 2002 to Feb. 2006).

Figure 5 shows the historical annual returns on 30-year Treasury bonds and the distribution of projections in SE-PIMS. Even at the 95th percentile, simulated rates of

---

[7] "Germany for First Time Sells 30-Year Bonds Offering Negative Yields," *Wall Street Journal*, 21 August 2019. Available at https://www.wsj.com/articles/germany-for-first-time-sells-30-year-bonds-offering-negative-yields-11566385847.

[8] Consider a $1 bond that was purchased at par when the yield was r_0. One year later, the bond pays a coupon of r_0 and its market value is equal to the net present value of 29 remaining annual coupon payments of r_0 plus a $1 principal repayment after 29 years, all discounted at the prevailing yield r_1. The rate of return is thus the coupon payment (r_0) plus the amount by which the market value after one year exceeds the $1 purchase price. We confirmed that SE-PIMS correctly converted yields into annual returns.

return frequently fall short of historical values. Rates of return on existing bonds are boosted when yields fall, but since current yields are already low, the potential for sizable yield decreases (and correspondingly high rates of return) is limited. In other words, we have no concern over the relatively low rates of return on bonds during the simulation period.

Turning from government bonds to corporate stocks, Figure 6 shows the historical annual returns (including dividends) of a portfolio of stocks that mimic the Standard & Poor's 500 (S&P 500) index from 1986 to 2018 (brown line), and the distribution of projected returns through 2038 (blue lines).

Figure 6. Historical and Distribution of Projected Annual Returns on Stocks

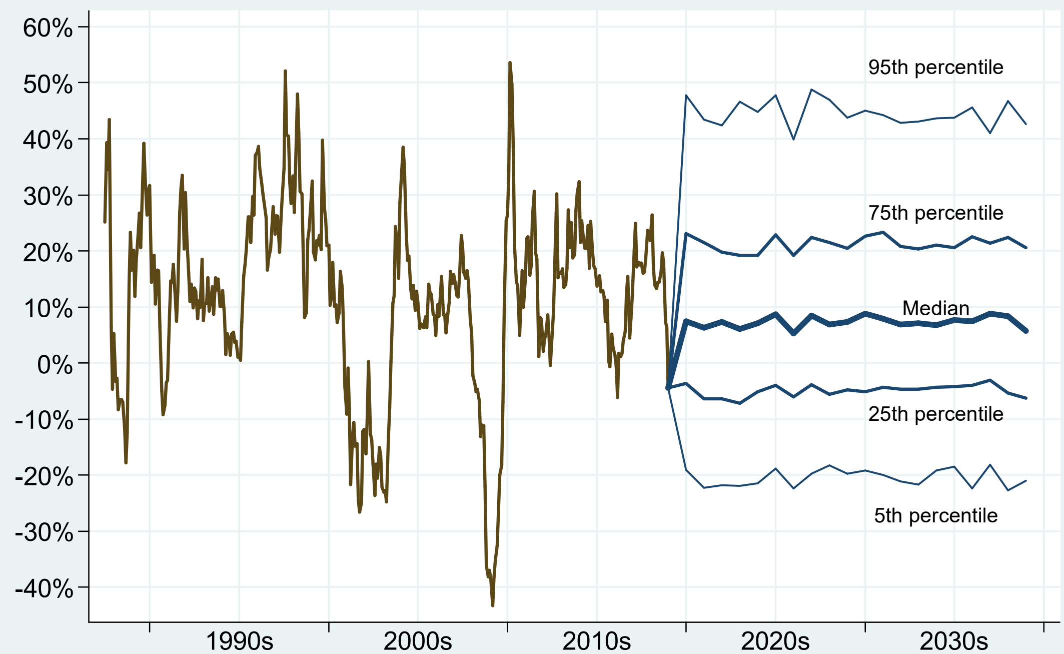

Sources: Bloomberg, Cboe (historical data), SE-PIMS (projections).

SE-PIMS models annual real returns on stocks as an autoregressive process (*SE-PIMS System Description*, page 2-10 and footnote 1 on page 5-13):

$$ln(1 + s_t) = \alpha_s + \beta_s ln(1 + s_{t-1}) + \varepsilon_{s,t},$$

where $\alpha_s = 0.0652$ and $\beta_s = 0$. (At these values, simulated values should randomly fluctuate around $e^{0.0652}$-1=6.7%.) We regressed simulated values of $ln(1 + s_{t-1})$ on $ln(1 + s_t)$ and confirmed that the estimated β_i is not statistically significant from zero in SE-PIMS simulations. Our estimate of α_s is 0.0708, and statistically different from 0.0652 at 5% significance level.

Simulated Firm-Level Variables

Bankruptcy probabilities are a function of (lagged) ratio of corporate assets over liabilities, ratio of cash flow (net of pension contributions) over assets, employment, funded ratio of the company's DB plan(s), and indicators for activity in the financial or utilities sectors. The industry indicators are time-invariant; future values of the other explanatory variables are simulated in SE-PIMS, appropriately transformed (logged/lagged), multiplied by parameter estimates, and converted into bankruptcy probabilities.

Table 1 shows the parameters applied to calculate bankruptcy probabilities. These were obtained from a logistic regression, as described on page 5.[9]

Table 1. Bankruptcy Probability Parameters

Intercept	-4.6623
Log(assets$_{t-1}$/liabilities$_{t-1}$)	-0.9197
Log(assets$_{t-2}$/liabilities$_{t-2}$)	0.2406
Log(employment$_{t-1}$)	-0.2503
Log(employment$_{t-1}$) x Log(assets$_{t-1}$/liabilities$_{t-1}$)	-0.0373
Log(employment$_{t-1}$) – Log(employment$_{t-2}$)	-0.6676
Log(funded ratio$_{t-1}$)	-0.1302
Cash flow$_{t-1}$/Assets$_{t-1}$	-3.8125
Cash flow$_{t-2}$/Assets$_{t-2}$	-2.0676
Financial industry	-2.7136
Utilities industry	-1.6721

For any single simulation, denote the sum of products of parameters and explanatory variables by $\beta' X$. The probability that a firm will file for bankruptcy in a given year is $\left(1 + exp(-\beta' X)\right)^{-1}$.

Special adjustments are made to particularly high bankruptcy probability estimates (which are capped at 20%) and to the probabilities of 14 firms (and their partners), whose intercepts were modified. We replicated the bankruptcy probability calculations for 2021–2038 and found that the 99.6% of the results matched those in SE-PIMS within 0.01%.[10] Almost all discrepancies of more than 0.01% were limited to just four firms (and their partners); a potential explanation is that their calculations were adjusted in a way that is unknown to us.[11] The cap at 20% affected about 0.03% of the calculations.

To determine whether a firm is simulated to enter bankruptcy, SE-PIMS converts probabilities into binary outcomes (bankruptcy or no bankruptcy). The standard way of doing this is by comparing the probability to a variable between 0 and 1, randomly drawn from a uniform distribution; if the probability exceeds the random variable, a

[9] The *SE-PIMS System Description* (page 6-13) documents these parameters with opposite signs. As printed in Table 1, a positive (negative) parameter may be interpreted as increasing (decreasing) the bankruptcy probability when the explanatory variable increases in value. For example, the parameter on employment is negative, indicating that more employees translate into lower bankruptcy risks.
[10] Calculations for 2018–2020 could not be verified because the required lagged values were not available in our data extract.
[11] The four firms are represented by IDs 580, 816, 1076, and 1271.

bankruptcy is simulated; otherwise, the firm continues operating. As a result, the average bankruptcy probability should equal the average share of firms that are simulated to file for bankruptcy (up to random variation). We found this to be approximately true for 2021 and later years, but not for 2020. Averaged over all 1,434 firms/partners and all 5,000 scenarios, i.e., over more than 7 million outcomes annually, the average fraction that SE-PIMS projected to file for bankruptcy was 0.55% in 2020, compared with an average bankruptcy probability of 0.49% (Figure 7). It is unclear to us what may have caused this discrepancy.

Figure 7. Average Bankruptcy Probabilities and Simulated Bankruptcy Rates (2018–2038)

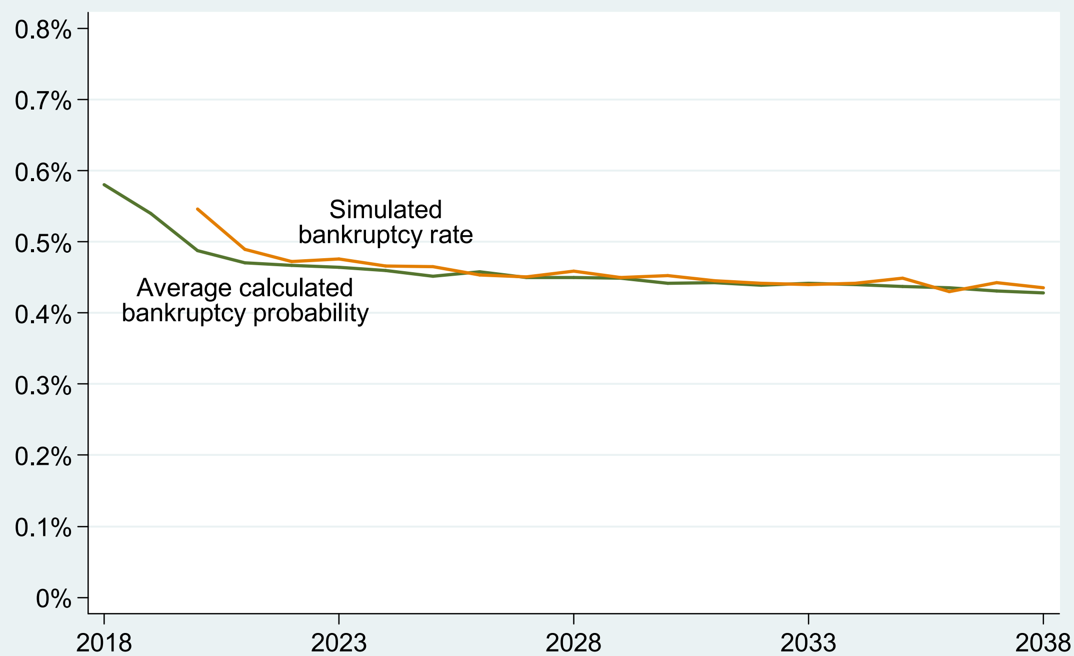

Source: SE-PIMS for 2018, specification 127.

Figure 7 also demonstrates that the average bankruptcy probability and rate gradually decrease over time. This is consistent with expectations, because high-risk firms face elevated bankruptcy rates and disproportionately drop out of the simulations, leaving a mix of relatively low-risk firms in later years.

The green line in Figure 7 shows average bankruptcy calculations for 1,434 firms/partners and 5,000 scenarios. Consider just three out of the 1,434 firms/partners (Figure 8). The firms differ in their 2018 credit rating, as compiled by Compustat. SE-PIMS simulated the highest average bankruptcy probabilities for the firm with the worst credit rating (BB-), followed by the firm rated at BBB, and the lowest probabilities for the firm with the best credit rating (AA+). This is, of course, consistent with expectations and suggests that the bankruptcy model responds appropriately to financial metrics. (The model does not control for credit rating itself; see Table 1).

Figure 8. Average Bankruptcy Probabilities for Three Illustrative Firms (2018–2038)

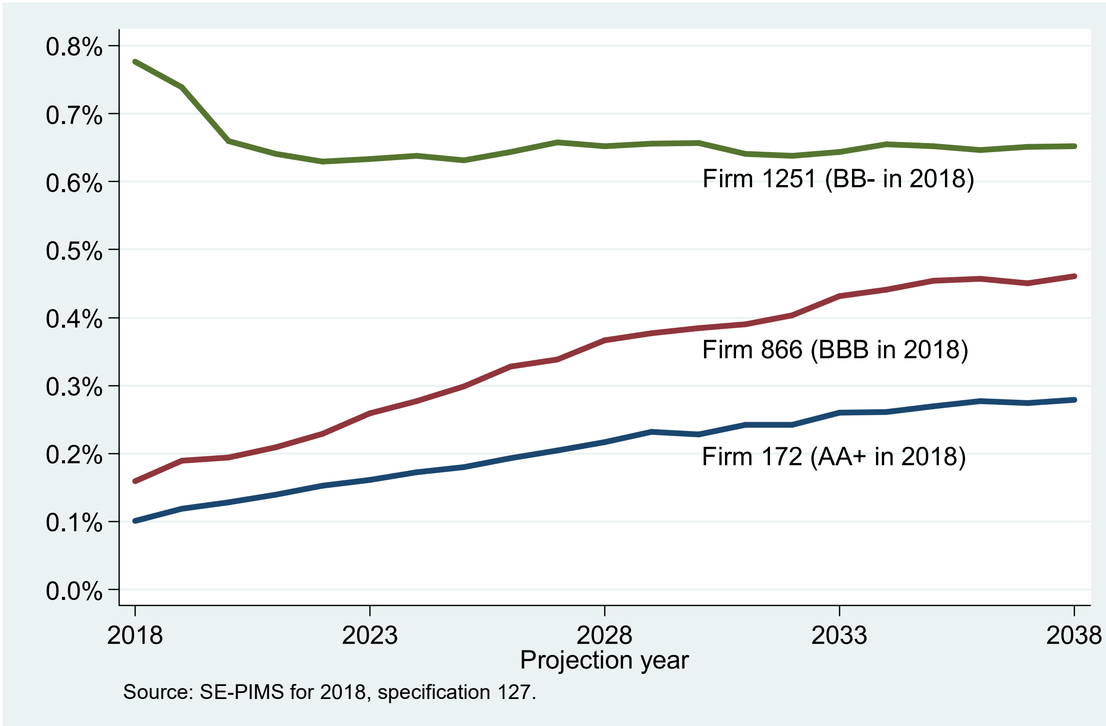

Figure 8 further demonstrates that average bankruptcy probabilities are converging over time. This is, again, consistent with expectations. The credit ratings are based on the firms' financial status in 2018. SE-PIMS models the ratios of equity to liabilities and of cash flow to assets as mean reversion processes, so that financially weak (strong) firms become stronger (weaker) over time, leading to a convergence of average financial strength and bankruptcy probabilities.

More generally, Figure 9 shows average bankruptcy probabilities of all firms with known credit rating, by rating category. (The vertical axis uses a logarithmic scale.) As expected, bankruptcy probabilities tend to correlate with credit rating, and converge over time.

Figure 9. Average Bankruptcy Probabilities, by Credit Rating (2018–2038)

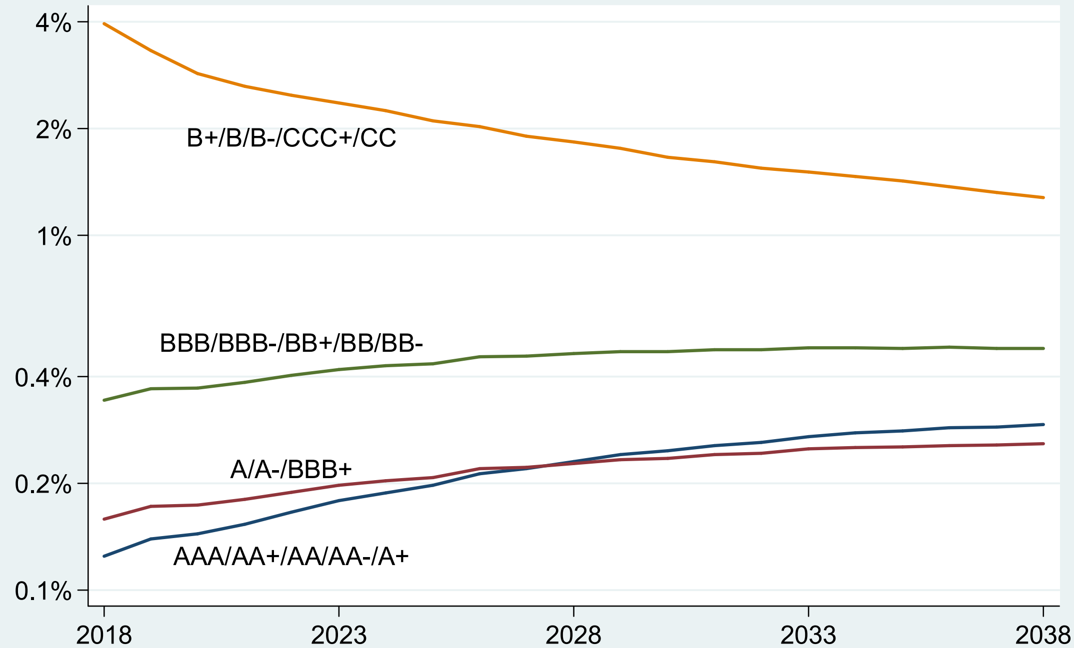

Source: SE-PIMS for 2018, specification 127. Probabilities are weighted and displayed on a logarithmic scale.

While it is expected that the average bankruptcy probability of firms with different initial financial strengths converge over time, microsimulation models generally aim to preserve distributions in their future scenarios. Figure 10 illustrates the distribution of simulated bankruptcy probabilities. (The vertical scale uses a logarithmic scale.) Apart from a decline of the 5th and 95th percentiles during the early part of the simulation period, the distribution of bankruptcy probabilities is remarkably stable.

Figure 10. Distribution of Simulated Bankruptcy Probabilities (2018–2038)

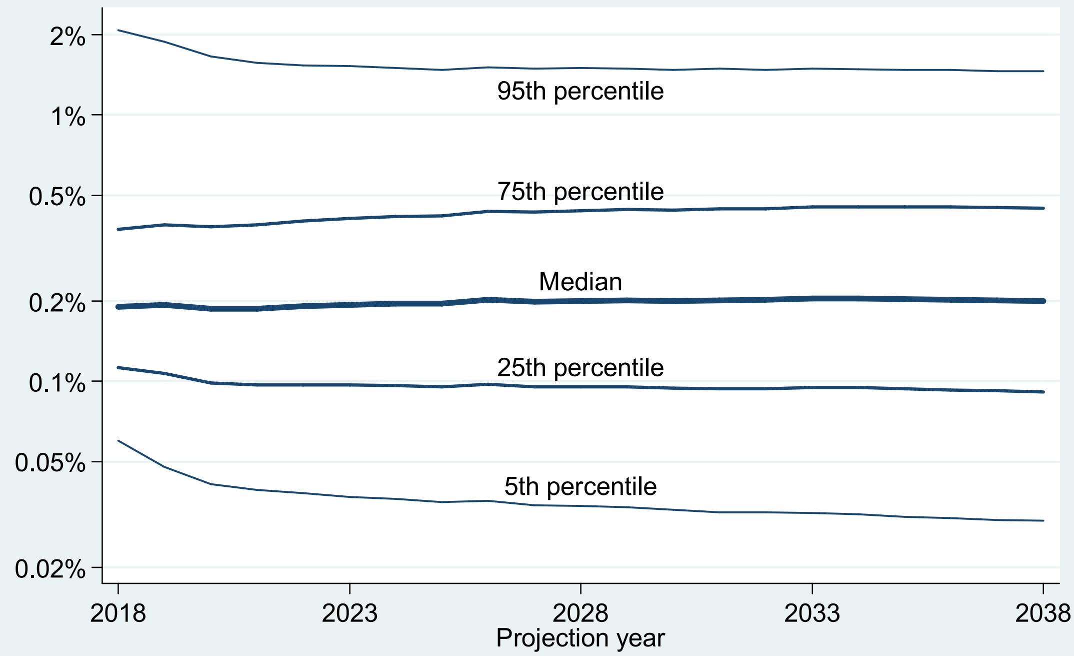

Source: SE-PIMS for 2018, specification 127.

The declining rates at the 5th and 95th percentile pose a puzzle. Since SE-PIMS's initial conditions reflect the economy's particular state around 2017, one may expect to see shifting distributions as the model evolves to stochastically cover a more representative set of states. However, the economy was generally robust in 2017, and initial bankruptcy rates may have been below their long-term average. The pattern in Figure 10, with declining rates at the 5th and 95th percentiles, runs counter to that argument.

Similar to Figure 10's distribution of bankruptcy probabilities, Figure 11 shows distributions of simulated values for other financial metrics:

1. Ratio of corporate assets over liabilities,
2. Employment (in 1,000 workers),
3. Corporate equity (stock market capitalization),
4. Pension underfunding (consolidated for firms that sponsor multiple plans),
5. Pension plan funded ratio (consolidated),
6. Corporate liabilities,
7. Ratio of corporate cash flow (before pension contributions) to assets, and
8. Ratio of pension contributions to assets.

The vertical axes use logarithmic scales (except see below). Most panels show the 5th, 25th, 50th, 75th, and 95th percentiles during the simulation period (blue lines). The percentiles are calculated separately for each year, i.e., the lines do not represent specific simulated paths.

Figure 11. Distribution of Simulated Financial Metrics (2018–2038)

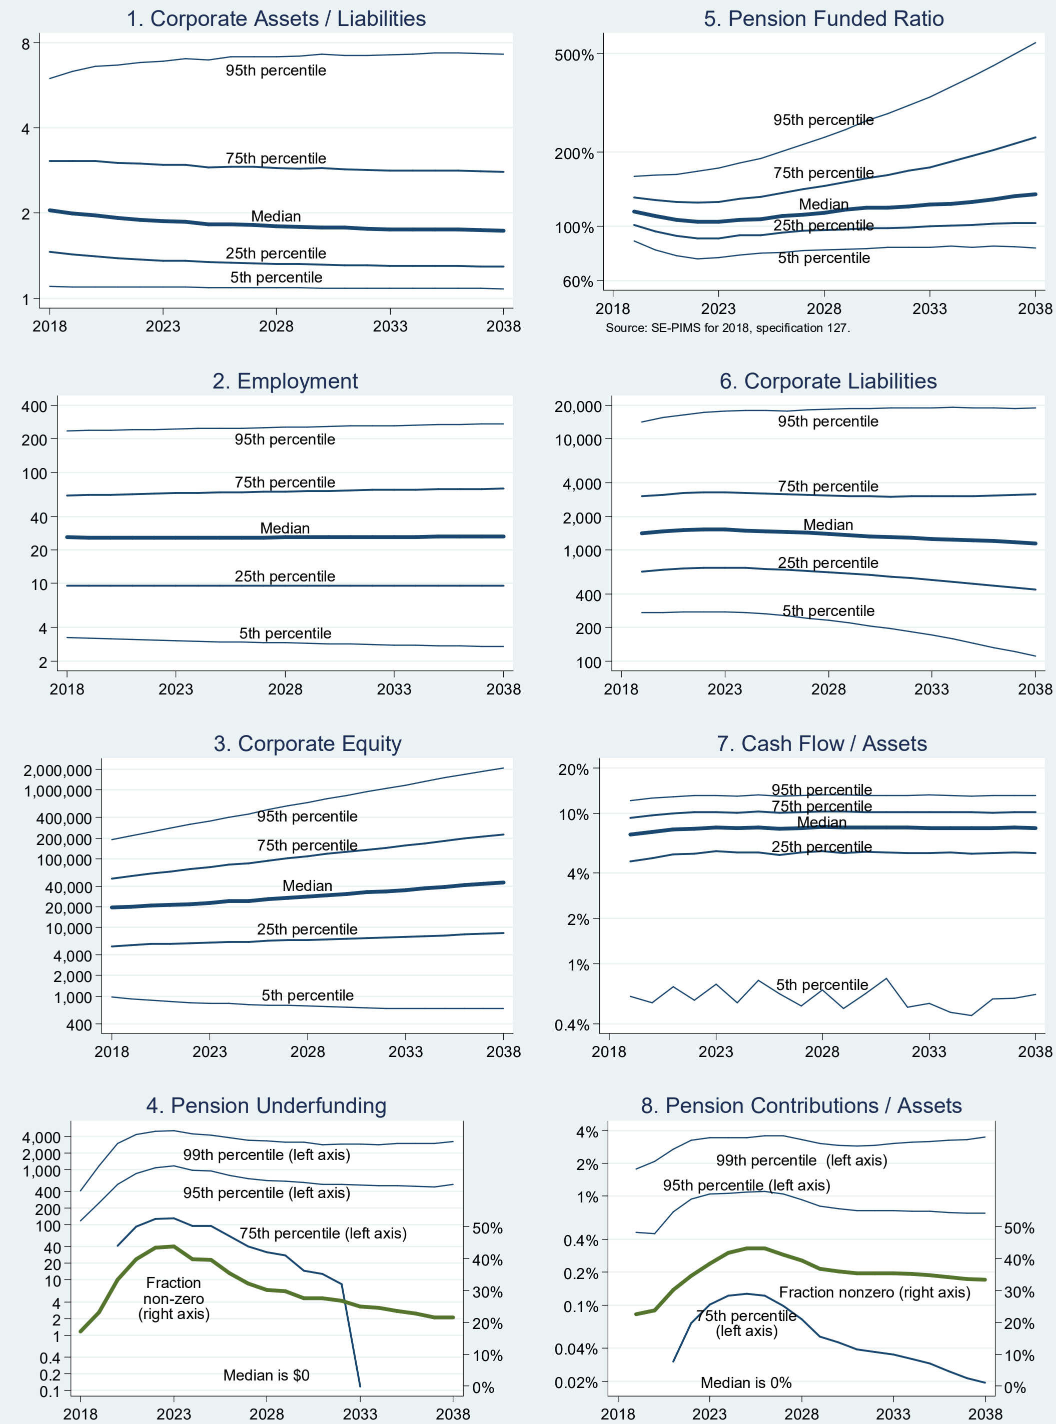

Source: SE-PIMS for 2018, specification 127.

In each year, about 7 million values of the pertinent metric enter into the calculation of percentiles.[12] Most values of pension underfunding (panel 4) and the ratio of pension contributions to assets (panel 8) are zero, so the medians of their distributions are zero. In addition to the 75th and 95th percentiles, these two panels also show the 99th percentile, along with a green line that depicts the fraction of values that are not zero. The green lines' scales are on the right axes; unlike all left axes, those scales are linear rather than logarithmic.

Figure 11 serves to gauge the stability of each metric's distribution over the simulation period. The distributions need not remain constant. For example, as time proceeds, weaker firms will have been removed from simulations due to a bankruptcy. Also, the 2018 or 2019 starting year is close to recent history, and recent history is not necessarily representative of longer periods.

Most simulated values maintain their distribution quite well in the 20-year simulation period, with two notable issues. First, corporate equity (panel 3) generally increases, which may be due to the specification of its generating process (*SE-PIMS System Documentation*, pages 5-7 and 5-13):

$$\Delta ln(E_t) = \alpha_E + \beta_E \Delta ln(E_{t-1}) + \varepsilon_{E,t}$$

where α_E=0.0306 and β_E=0.0206. In words, the equity growth rate is weakly persistent over time (positive but small β_E), with an upward drift in equity of about 3% per year. That drift translates into an upward-shifting distribution of corporate equity over the simulation period. The upward shift may reflect general economic growth, though its primary driver (α_E) is not correlated with macroeconomic variables. (The macroeconomic portion of the stochastic innovation term, $\varepsilon_{E,t}$, is highly correlated with the residual in stock market returns, but its mean is zero.) Insofar we are aware, corporate equity does not play any direct role in bankruptcy modeling, and we did not investigate further.

Second, the distribution of pension funded ratio (panel 5) generally deteriorates until 2023 and improves in later years. We are unsure what drives this pattern. The early years may be driven by SE-PIMS's assumption that firms make only the minimal required pension contribution, and that contribution is generally small in the near future because it is our understanding that many firms have accumulated credit balances in recent years. Improvements in plan funding in later years are more difficult to explain. By 2028, a quarter of plans are predicted to have funded ratios in excess of 145% (and in excess of 228% by 2038).

We understand that the concept of funded ratio in Figure 11 differs from the metric used for purposes of plan termination after a sponsor bankruptcy. The latter tends to be lower, in part because it tends to assume a lower discount rate of future benefits. That said, the two metrics presumably move in tandem, and it appears worthwhile to explore whether the funded ratio for termination purposes is projected to follow a similar pattern as that for funding purposes. A generally upward trajectory may

[12] To be precise, 7,170,000 values (1,434 firms/partners times 5,000 scenarios) in the early years and fewer in later years because firms drop out after their path involves a bankruptcy filing. By 2038, 6,510,400 values remain for each metric.

suggest that SE-PIMS understates claims due to plan terminations following bankruptcies.

The pattern in funded ratios is consistent with the distribution of pension underfunding (panel 4); the fraction of firms with underfunding initially increases from 17% in 2018 to 44% in 2023 and declines thereafter. It is also consistent with the distribution of the ratio of pension contributions to assets (panel 8), and with the *FY 2018 PBGC Projection Report*, which anticipates a decline in new claims after 2023.[13]

Correlations between Macroeconomic and Firm-Level Variables

We calculated correlation coefficients between select pairs of macroeconomic and firm-level variables.

Stock market returns and bankruptcy risks. We expect a negative correlation between bankruptcy risks and the rate of return on a broad basket of stocks. However, the contemporaneous correlation between bankruptcy risk and rate of return on stocks is 0.0063 in the simulation data, i.e., positive and very small.[14] We also calculated annualized 10-year trailing average rates of return and found a correlation with bankruptcy risks of -0.1116; negative, as expected.[15]

Stock market returns and the ratio of cash flow to assets. We expect a positive correlation between cash flow and stock market returns. However, the contemporaneous correlation was -0.2090 in the simulation data, and the correlation with annualized 10-year trailing average rates of return on stocks was -0.1651. This finding goes counter to our expectations; see below.

Rate of return on pension assets and pension funded ratio. SE-PIMS assumes that DB plan assets earn a rate of return that is a function of returns on bonds and stocks. (Ultimately, SE-PIMS exogenously generates only the yield on 30-year Treasury bonds and stock market returns.) We expect a positive correlation between the rate of return on pension assets and plans' funded ratios. Indeed, the contemporaneous correlation was positive (0.1631). More importantly, the correlation with annualized 10-year trailing average rates of return on pension assets was also positive (0.4409) and stronger than the contemporaneous correlation, as expected.

Rate of return on pension assets and pension underfunding. We expect a negative correlation between rates of returns on pension assets and plans' underfunding. The

---

[13] See Figure 13 in the *FY 2018 PBGC Projections Report*, available at https://www.pbgc.gov/sites/default/files/fy-2018-projections-report.pdf.
[14] The correlation is based on 136,507,182 records (up to 20 simulation years of 1,434 firm/partners and 5,000 scenarios). Arguably, the number of degrees of freedom is only one-tenth of this number, because macroeconomic variables are generated for only 500 scenarios and replicated 10 times. Either way, the number of observations is huge and all reported statistics in this section are highly statistically significant.
[15] Based on simulated values in 2028 and later; we used only simulated values—which start with 2019—to calculate 10-year trailing average rates. The number of records (73,376,024) is again very large.

contemporaneous correlation was -0.1290 and the correlation with annualized 10-year trailing rates was -0.1871; negative and stronger than the contemporaneous correlation, as expected.

Rate of return on pension assets and ratio of pension contributions to assets. We expect a negative correlation between rates of returns on pension assets and firms' contributions to pensions. The contemporaneous correlation was -0.0078 (negative but small) and the correlation with annualized 10-year trailing rates was -0.1259; negative and stronger than the contemporaneous correlation, as expected.

In sum, the correlations between macroeconomic and firm-level metrics that we explored are generally consistent with expectations, except for the correlation between stock market returns and the ratio of cash flow to assets. This exception may be the result of the autoregressive way in which the ratio of cash flow to assets is modeled (*SE-PIMS System Description*, page 5-7 and 5-13):

$$CFA_t = \alpha_{CFA} + \beta_{CFA} CFA_{t-1} + \gamma_{CFA} F + \varepsilon_{CF\ \ ,t}$$

where CFA_t represents the ratio of cash flow to assets in year t and F is an indicator for the financial services industry. In words, the ratio depends on the firm's ratio in the prior year and on a stochastic innovation term, $\varepsilon_{CFA,t}$. The macroeconomic portion of that term is negatively correlated with the stochastic term in stock market returns; the correlation coefficient is -0.76 (*SE-PIMS System Description*, page 5-14). This explains the overall negative correlation, but the pattern goes counter to our expectations.

Incomplete Stochastic Variation

Finally, we observed a lack of variation among certain variables in the simulation data. In principle, we expect all firm-specific metrics to vary across all simulation scenarios. However, employment, pension underfunding, pension funded ratio, and corporate liabilities values appear to be constant across partners of a firm and across cycles within a scenario. Consider Table 1, which shows select variables of a randomly selected firm (51) and a randomly selected group of scenarios (212-1 through 212-10). The firm enters both as itself (51-1) and for two partners (51-2, 51-3). Note that all variables take on the same value for all partners and simulation cycles (except when a firm is simulated to have filed for bankruptcy, such as firm 51-1 in scenario 212-5 and firm 51-2 in scenario 212-8). While only 2019 and 2020 are shown, this lack of variation occurs in all years. Other metrics, such as the ratio of corporate assets to liabilities, do vary across partners and scenarios. It is our understanding that this lack of stochastic variation is a design choice to speed up certain computation-intensive actuarial calculations in SE-PIMS.

Table 2. Illustrative Simulation Values of Select Variables

Firm-Partner	Scenario-Cycle	Employment 2019	Employment 2020	Underfunding 2019	Underfunding 2020	Funding ratio 2019	Funding ratio 2020	Liabilities 2019	Liabilities 2020
51-1	212-1	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-2	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-3	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-4	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-5	52.73	.	72.29	.	0.878	.	590.33	.
51-1	212-6	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-7	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-8	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-9	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-1	212-10	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-1	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-2	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-3	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-4	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-5	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-6	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-7	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-8	52.73	.	72.29	.	0.878	.	590.33	.
51-2	212-9	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-2	212-10	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-1	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-2	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-3	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-4	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-5	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-6	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-7	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-8	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-9	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04
51-3	212-10	52.73	59.35	72.29	90.79	0.878	0.845	590.33	584.04

Source: SE-PIMS for 2018, specification 127.

Recovery Claims by Bankruptcy Type

Chapter 7 versus Chapter 11

While SE-PIMS has the capability to model a Chapter 7 or a Chapter 11 bankruptcy separately, the probability of a Chapter 11 bankruptcy is set to zero in its 2018 implementation. However, the historical number of Chapter 11 reorganizations is sizeable, and the fraction of total filings that were under Chapter 11 fluctuated between 17% and 34% (see Figure 12).[16]

Figure 12. Chapter 7 and Chapter 11 Bankruptcy Filings over Time

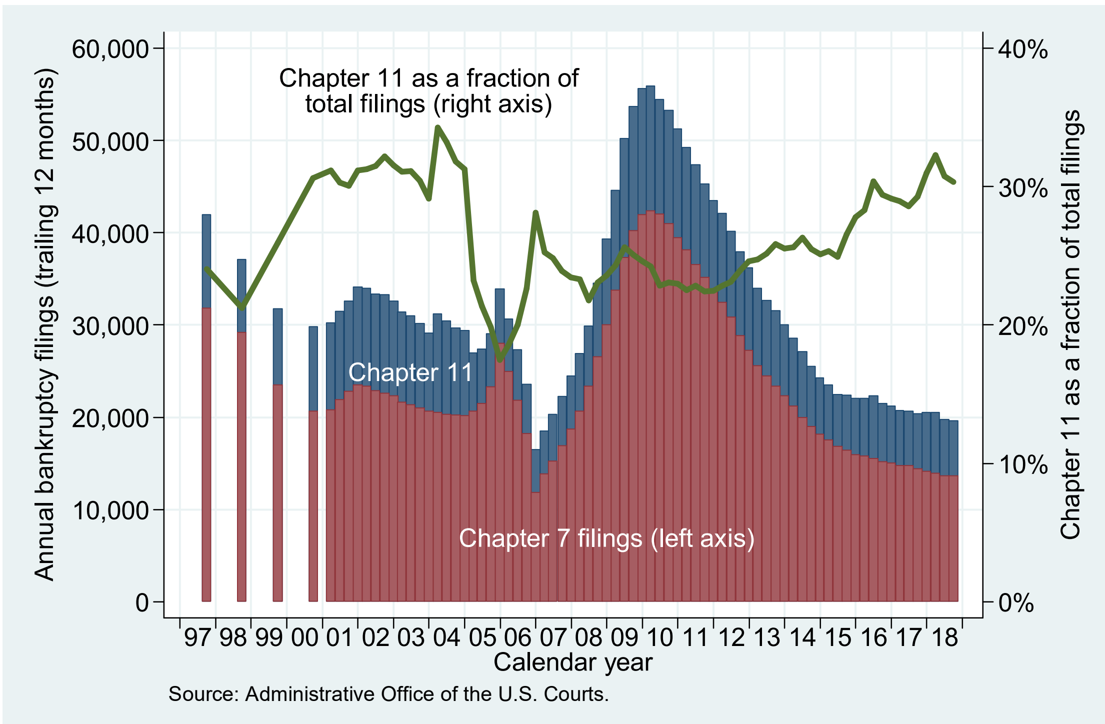

Chapter 11 bankruptcies involve a reorganization and can have different implications for pension plans compared to a Chapter 7 liquidation. We analyzed differences in recovery rates for DB pension plans after filings under Chapter 7 and Chapter 11. The analysis is based primarily on two data sources. First, the PBGC provided us with a listing of terminated and trusteed plans, their termination date, unfunded liabilities, and the amounts that could be recovered in bankruptcy proceedings for "Due and Unpaid Employer Contribution" and unfunded "Employer Liability." (It is our understanding that the latter represents an unsecured claim in bankruptcy, whereas

---

[16] Figure 12 is based on all corporate bankruptcies, not just bankruptcies of companies that sponsor a DB plan. It is our impression that DB plan sponsors tend to be larger than other corporations, and that Chapter 11 is more common among larger than among smaller companies. In other words, Chapter 11 reorganizations may be more common among DB plan sponsors than Figure 12 suggests. Also see Table 3, where we identified 62 filings under Chapter 7 and 164 under Chapter 11.

the former has a higher level of priority.) Second, the PBGC provided us with a database from New Generation Research with information on corporate bankruptcies (company name, EIN, year of filing, and type of filing—Chapter 7, 11, or other). To improve the match rate between the two data sources, we also used Form 5500 filings, which provide a cross-walk between plan and sponsor names.

We restricted the list of terminated plans to 622 plans that were trusteed in 2010–2018. Of these, 62 were matched to a Chapter 7 filing and 164 to a Chapter 11 filing. The bankruptcy filing type of the remaining 396 plans could not be determined because they matched to multiple bankruptcies of different filing types (11 plans) or did not match (385 plans). Table 3 summarizes the PBGC's recoveries in bankruptcy proceedings, by bankruptcy type.

Table 3. Recoveries in Bankruptcy Proceedings, by Type of Bankruptcy Filing (2010–2018)

	Chapter 7	Chapter 11	Unknown	All
Number of trusteed plans	62	164	396	622
Average recovery	$0.8m	$2.5m	$0.3m	$1.0m
Average unfunded liability	$4.8m	$25.6m	$5.5m	$10.7m
Mean recovery rate				
Unweighted	6.8%	9.6%	5.7%	6.8%
Weighted by unfunded liability	17.0%	9.8%	6.2%	9.0%
Median recovery rate				
Unweighted	0.0%	1.1%	0.0%	0.0%
Weighted by unfunded liability	1.5%	5.1%	0.7%	2.3%

Sources: PBGC, New Generations Research, Form 5500 filings.

Both the average recovery and the average unfunded liability were smaller for plans whose sponsor liquidated under Chapter 7 than for plans whose sponsor reorganized under Chapter 11. The average recovery rate as a percentage of unfunded liability was also lower under Chapter 7 (6.8%) than under Chapter 11 (9.6%). However, the pattern reverses when weighting the recoveries by unfunded liability: 17.0% under Chapter 7 and 9.8% under Chapter 11. Put differently: the PBGC recovered 17.0% of the unfunded liabilities under Chapter 7, compared with 9.8% under Chapter 11. Ultimately, the weighted recovery rate is more relevant to the PBGC's balance sheet than the unweighted average rate.

The analysis is sensitive to the experiences with large plans. For example, the higher weighted fraction under Chapter 7 is due to a single plan whose recoveries amounted to 49% of that plan's unfunded liabilities and more than doubled the weighted recovery rate of plans under Chapter 7.

Median recovery rates tend to be low because the recovery was zero for many plans. For example, the PBGC did not recover anything from 55% of plans under Chapter 7, so the median recovery was 0.0%.

The average recovery rates for all 622 terminated plans were 6.8% (unweighted) and 9.0% (weighted).

In sum, while the recovery rates in Chapter 7 liquidations may differ from those in Chapter 11 reorganizations, our analysis presents no compelling case for changing

SE-PIMS's current practice of applying the same recovery rate to all plans that are projected to be trusteed after a bankruptcy. Separately, SE-PIMS currently assumes that 5% of unfunded liabilities will be recovered. This is a somewhat conservative estimate in light of the 9.0% average weighted recovery among all 622 trusteed plans. To place the difference in perspective: SE-PIMS assumes that the PBGC will absorb 95% of unfunded liabilities, and our findings suggest that the net claim was in fact 91% of unfunded liabilities in recent years.

Repeat Bankruptcies

Separately, AACG investigated repeat bankruptcies by the same firm. Based on bankruptcy data from New Generation Research, we identified 216 companies that filed for bankruptcy twice between 1981 and 2018 (Figure 13).

Figure 13. Bankruptcies by Year for Companies that Filed for Bankruptcy Twice

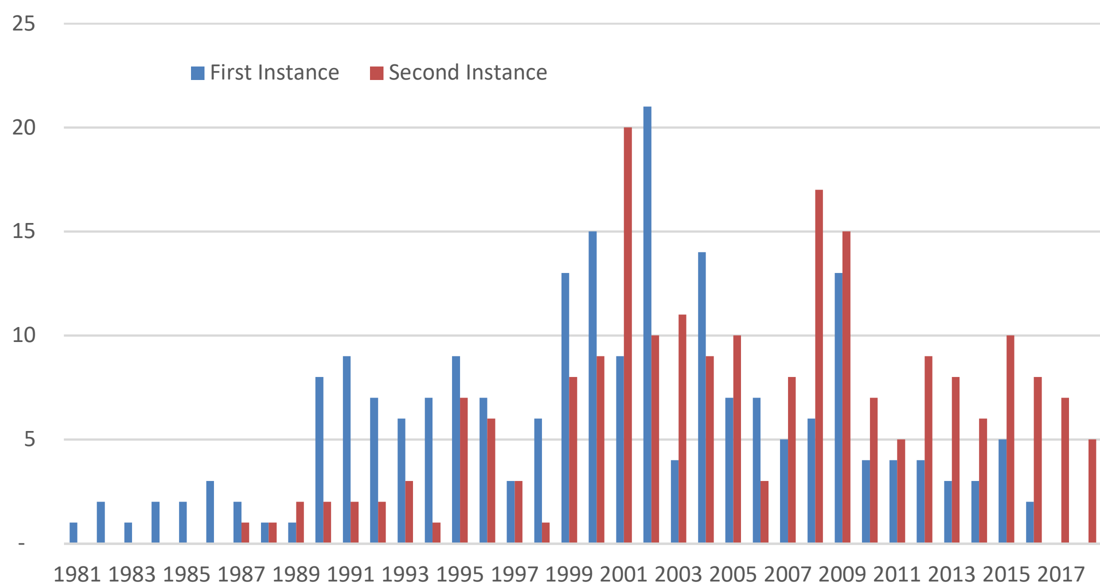

Source: NEW GENERATION RESEARCH, INC.

Of these companies, 184 companies were public when they first filed for bankruptcy and 32 were private. A vast majority (209 out of 216 companies) filed under Chapter 11, and only two companies under Chapter 7. (Five companies filed under Chapter 15, under which a representative of a corporate bankruptcy proceeding outside the United States can access the U.S. court system.)

Figure 14. Number of Public and Private Companies by Bankruptcy Filing Instance (1981–2018)

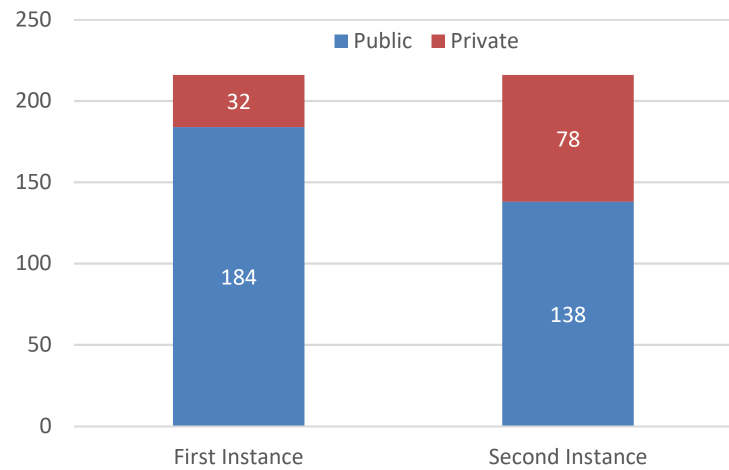

Source: NEW GENERATION RESEARCH, INC.

When these same companies filed for bankruptcy again, only 138 were public and 78 were private (Figure 14). Most repeat bankruptcies (198 out of 216) again filed under Chapter 11, while 13 companies filed under Chapter 7 (Figure 15).

Figure 15. Type of Bankruptcy by Bankruptcy Filing Instance

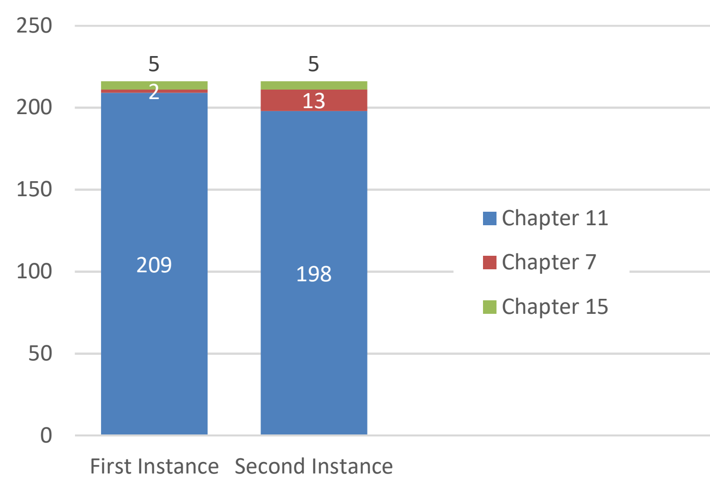

Source: NEW GENERATION RESEARCH, INC.

We identified 17 companies that filed for bankruptcy three times. All were public when they first filed for bankruptcy. One company was private when it refiled, and seven companies were private by the time they filed for the third time. All filed under Chapter 11 of the Bankruptcy Code in every instance.

Ideally, the analysis would tell whether second or third bankruptcies are more or less likely to occur than first bankruptcies. Unfortunately, the sample is too small to reliably answer that question, especially since our interest is limited to companies that sponsor a DB plan. Even if SE-PIMS were to model subsequent bankruptcies, their contribution to the parameter estimation would be restricted further because their DB plan could have terminated after the first filing or because they became privately held, so that their financial information would no longer be public. In sum, we envision little benefit from expanding SE-PIMS to account for repeat bankruptcies.

Assumptions Triggering Plan Termination in Bankruptcy

When a plan sponsor enters bankruptcy in an SE-PIMS scenario, its plan(s) may or may not be trusteed, depending on the plans' funded ratios (*SE-PIMS System Description*, page 4-4). If the market value of assets of all plans sponsored by the firm is 80% or more of the combined termination liabilities for all the plans, no claim is assumed. Otherwise, PBGC terminates and trustees the plans, i.e., it assumes the plans' assets and liabilities. We explored to what extent the 80% rule matches recent experiences.

The PBGC provided us with a list of plans that were terminated after their sponsor filed for bankruptcy, and another list of plans that survived their sponsor's bankruptcy.[17] Information on terminated plans included assets and liabilities, so their funded ratio was readily determined. For plans that survived a bankruptcy, unfunded liability was available, but not their funded ratio.

In the context of the 80% rule, assets and liabilities are measured for termination purposes as opposed to for funding purposes, as reflected on Form 5500 filings. An important difference lies in the discount rate applied to future benefit payments. The discount rate for termination purposes tends to be lower than that for funding purposes, resulting in higher liabilities and a lower funded ratio.

The lists with terminated and surviving plans were matched to Form 5500 filings just prior to the approximate date of bankruptcy. We imputed the funded ratio of surviving plans as follows.

First, we assumed that the PBGC's metric of assets for termination purposes is close to the market value of assets reported on Form 5500 filings (Schedule SB, line 2a). Since liabilities are equal to the sum of assets and unfunded liabilities, we imputed the funded ratio of surviving plans as assets/(assets+unfunded liabilities).[18] Figure 16 shows the distribution of funded ratio (for termination purposes) for 300 trusteed plans and 105 consolidated surviving plans. The boxes delimit the 25th, 50th, and

---

[17] More precisely: a list of companies that filed for bankruptcy and whose plans survived. We understand that in a small number of cases, one or more plans survived the bankruptcy, and one or more plans were terminated and trusteed. However, no information was provided on which plans survived and which did not. Since such mixed results are rare, we treated all plans of corporate parents on the list as having survived.

[18] Recall that information on unfunded liabilities for termination purposes is provided for corporate parents and not for their individual plans. We therefore consolidated assets from Form 5500 filings to the level of the sponsor, and used consolidated plans as the unit of observation.

75th percentiles; the "whiskers" denote adjacent values, and the single dots reflect individual outlier plans.[19] The red vertical line is at 80% funding.

Figure 16. The Distribution of Funded Ratios around Bankruptcy for Trusteed and Surviving Plans (2009–2018); Imputation Based on Unfunded Liability and Form 5500 Assets

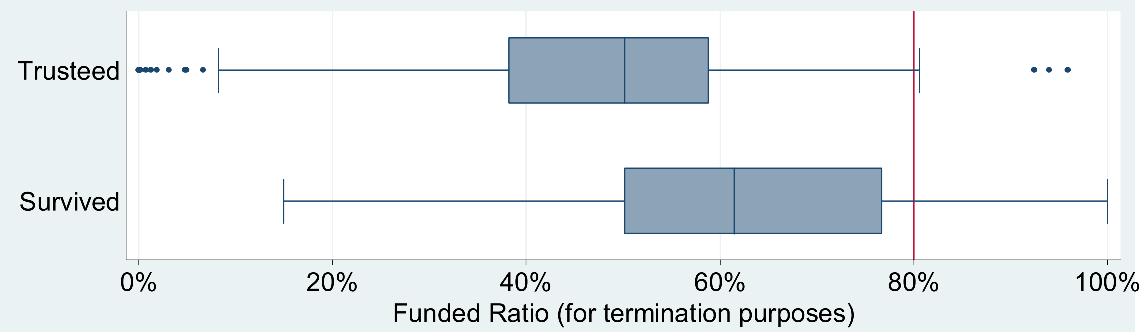

Source: PBGC, Form 5500 filings.

As expected, almost all trusteed plans had a funded ratio under 80%. However, most plans (81%) that survived their sponsor's bankruptcy also had an (imputed) funded ratio below 80%. While surviving plans were generally better funded than trusteed plans, most surviving plans would be misclassified based on the 80% rule.

Second, we adopted a regression-based imputation approach. Using complete data for terminated plans, we estimated a linear regression model of funded ratio for termination purposes and applied the parameter estimates to consolidated surviving plans to impute their funded ratios. Explanatory variables included the "adjusted funding target attainment percentage" (Schedule SB, line 15), assets (Schedule SB, line 2a), total funding target (Schedule SB, line 3d(3)), and unfunded liabilities for termination purposes (from the PBGC); all monetary figures were converted into per participant amounts. (Several alternative model specifications gave very similar results.) See Figure 17.

---

[19] Adjacent values are the most extreme values within 150% of the interquartile range from the nearer quartile. Tukey's Rule, a commonly used guideline in statistical analysis, states that statistical outliers are values outside the adjacent values. (John W. Tukey, 1977. *Exploratory Data Analysis*. Addison-Wesley.)

Figure 17. The Distribution of Funded Ratios around Bankruptcy for Trusteed and Surviving Plans (2009–2018); Imputation Based on Regression Model

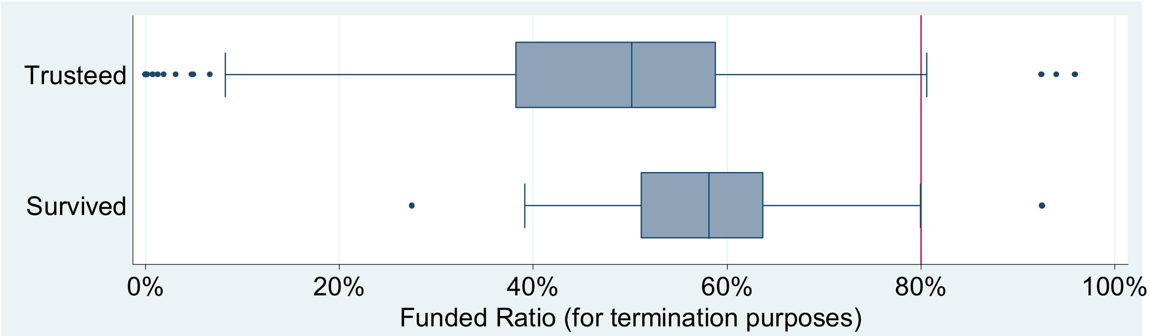

Source: PBGC, Form 5500 filings.

The regression-based imputation also resulted in funded ratios among consolidated surviving plans that rarely exceed 80%, suggesting that the decision to terminate and trustee a plan is more complex than a simple rule based on an 80% funded ratio threshold.

Separately, we were surprised to find that the funded ratios of surviving plans were generally low. Of course, neither imputation approach discussed above is perfect. It may be worthwhile for the PBGC to revisit the issue based on non-imputed data.

Finally, we explored whether other metrics better distinguish trusteed and surviving plans, but found that both adjusted funding target attainment percentage (from Schedule SB, line 15) and per participant unfunded liability (from the PBGC) have distributions that overlap in a manner similar to that depicted in Figure 16.

3. MASS WITHDRAWAL MODELING IN ME-PIMS

Overview

At a very high level, ME-PIMS incorporates mass withdrawals as follows.

- Plans may require financial assistance from the PBGC when they become insolvent. A plan is considered insolvent when its assets are depleted and it can no longer make benefit payments.
- Mass withdrawals may trigger or accelerate insolvency. A mass withdrawal occurs when all participating employers stop participating in the plan at the same time.
- ME-PIMS "books" a financial responsibility within 10 years before it is expected to become insolvent for an ongoing plan or within 20 years for a plan in mass withdrawal.
- For each plan, ME-PIMS calculates a probability of mass withdrawal based on six factors that relate to the size and financial status of the plan.
- Once a plan is flagged as facing a mass withdrawal, the plan is frozen—no further accruals due to service or increases in benefits units, no new entrants, and a roll-back of benefits. The roll-back affects active, retired and term-vested participants. The plan will be owed withdrawal liability payments from formerly participating employers, but some of those payments may not be collectable.
- The PBGC books a financial responsibility within 10 or 20 years before an expected insolvency, but provides financial assistance only upon insolvency. It is possible that the plan's finances improve (in reality or in simulation), in which case the claim will be unbooked.

This section reviews several areas related to mass withdrawal modeling in ME-PIMS.

First, the mathematical formula that ME-PIMS uses to predict mass withdrawals is complex and not based on a statistical model of historical mass withdrawals. We discuss the model and suggest an alternative in the section "Model Complexity," starting on page 32.

Second, ME-PIMS currently models mass withdrawals, but not withdrawals by individual employers. We analyzed individual withdrawals in historical data; see the section on "Withdrawals by Individual Contributing Employers," starting on page 35. (Separately, our companion report on potential contagion effects of mass withdrawals discusses factors that employers may consider when deciding whether to withdraw from a plan.)

Third, employer withdrawals generally trigger a withdrawal liability, which is payable in installments. ME-PIMS makes certain assumptions about the collectability of withdrawal liability payments. We attempted to validate these assumptions in the section on "Collectability of Withdrawal Liabilities," starting on page 41.

Finally, ME-PIMS simulates a sample of multiemployer plans, rather than the universe. We identified potential issues and obstacles to incorporating all plans in the section on "Modeling the Population Rather Than a Sample," starting on page 41.

Model Complexity

ME-PIMS calculates mass withdrawal probabilities based on six metrics that relate to plan size and financial status:

S1. Plan size (tied to the number of active and total participants),
S2. Ratio of active to inactive participants,
S3. Ratio of assets to benefit payments and expenses,
S4. Ratio of the market value of assets to vested liabilities,
S5. Ratio of current year to previous year contributions, and
S6. Ratio of credit balance to employer contributions.

Those metrics are time-varying and must themselves be forecast in order to predict mass withdrawals. An algorithm transforms the six metrics into mass withdrawal probabilities such that those probabilities, on average, mimic historical mass withdrawal rates. (For purposes of this section, ignore the possibility that future rates may differ from historical rates.) Smaller plans and plans in weaker financial health are projected to experience a mass withdrawal at higher rates than larger plans on solid footing.

Like SE-PIMS, ME-PIMS is a system to forecast, ultimately, liabilities on the PBGC's balance sheet. Its accuracy is enhanced if it can pinpoint the plans that will experience a mass withdrawal. However, it makes little difference whether ME-PIMS misidentifies plan *A* for mass withdrawal when really plan *B* will experience it; what matters is the size of the liability that mass withdrawals create for the PBGC. Some of the six metrics that are used to calculate a mass withdrawal probability relate to that liability, but none capture it directly or fully. For example, one of the metrics is the year-over-year percentage increase in contributions to the plan, which seems to bear no relationship to the financial assistance that the plan may require.

Potential financial assistance is perhaps best measured by a plan's total unfunded vested liability. That metric is not used to calculate mass withdrawals, possibly because its predictive ability is low. This raises the question whether the modeling system gains much from calculating mass withdrawal probabilities based on factors that may have little correlation with potential financial assistance.

We do not know how well the six factors correlate with total unfunded vested liability. If they correlate poorly, perhaps the model may be simplified into a uniform probability of mass withdrawal that is applied to all plans. Or perhaps the sole factor determining that probability should be a plan's total unfunded vested liability.

A similar issue exists in the context of bankruptcies in SE-PIMS. In that model, several explanatory variables are based on Compustat data, which are available for only a subset of plan sponsors. Simplification of the bankruptcy model in SE-PIMS may therefore remove data constraints. That potential benefit does not extend to ME-PIMS, because all explanatory variables are based on Form 5500 filings. Still, the current model's complexity does not necessarily translate into greater accuracy of the PBGC's expected financial assistance payments.

Mass withdrawal probabilities are based on six factors that are combined into a probability. Page 50 of *Key Differences Between SE-PIMS and ME-PIMS* describes the process as follows. Consider six factors *S1* through *S6*, each of which is a transformation of financial status or number of plan participants. One of those factors, *S2*, is a transformation of funded ratio (*FR*), which also enters the algorithm directly. The mass withdrawal probability (MWP) is given by:

Term1 = min (1 , (*S1* + *S2*+ *S3*) / 3.5 + *S4* / 20)

Term2 = max (*S1, S2, S3*) * 0.4

MWP = max (Term1, Term2) * *S5* * *S6*

If (*FR* >= 0.9 + *S1*) then MWP = 0.

MWP = MWP ^ (2 - MWP)

MWP = MWP / max (1 , 100 * max (0.25 - MWP , 0))

MWP = MWP * 0.5

The algorithm to combine and transform the six factors is not immediately transparent. It appears to be increasing in all six factors, each of which it itself constructed such that it increases with financial health or decreases with plan size. The first four lines combine the six factors into an intermediate number, which we here label as "mass withdrawal risk score." The last three lines transform the mass withdrawal risk score into a final mass withdrawal probability. Figure 18 numerically illustrates the transformation of risk score into probability, i.e., it represents the last three formulae above. The orange line at 0.4% denotes the weighted average annual mass withdrawal probability cited by French (2017)[20] and confirmed by our own analysis based on mass withdrawals in 2011–2018.

[20] Darren French (2017). "ME-PIMS Mass Withdrawal Assumptions." PBGC-PRAD memorandum dated March 2, 2017. Available at https://www.pbgc.gov/sites/default/files/me-pims-masswithdrawalassumptions.pdf.

Figure 18. Transformation of Mass Withdrawal Risk Score into a Probability

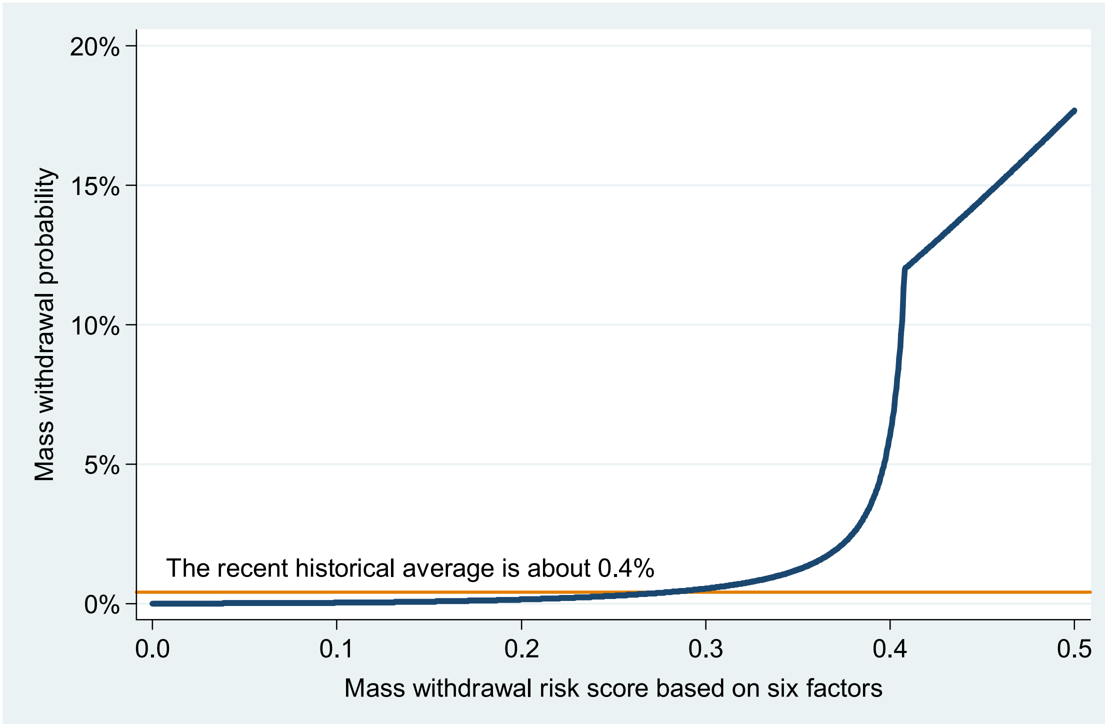

The algorithm was "retro-fit" to PBGC data indicating which ME plans were booked from 2003 to 2008, "so as to yield overall [mass withdrawal] probabilities on the order of ½% or less, with significantly higher than average probabilities for the plans that were actually booked during that period." (PBGC 2011, p. 49). We did not have access to simulation data to verify that simulated mass withdrawal rates were approximately 0.5%. However, we identified 40 mass withdrawals in 2011–2018, which translates into a rate of approximately 0.4%. This suggests that the average target rate remains approximately correct.

An important issue is whether the transformation is superior to a statistical model, such as a simple logit model. The "retro-fit" presumably involved extensive calibration of the functional form and scaling factors, and it is difficult to assess the six factors' relative contributions and trade-offs. In addition, the risk score transformation into a probability is much steeper above the average probability (0.4%) than below it. In other words, around the mean, positive changes in any of the six factors affect the probability more strongly than negative changes.

The properties of the algorithm are unclear. For example, it is unclear whether funded ratio is statistically significantly related to mass withdrawal probability. In contrast, the statistical properties of a logit model are well established. A statistical model may also be better equipped to evaluate other factors highlighted in other research that may influence mass withdrawals. For example:

- As argued above, PBGC's potential financial assistance payments are most directly affected by the total unfunded vested liability, suggesting that ME-PIMS may become more accurate if the mass withdrawal probability model controls for total unfunded vested liability.

- A statistical model can readily test for non-linear effects. According to Kiska et al. (2017),[21] "CBO expects that employers in severely underfunded plans are more likely to withdraw than employers in less underfunded plans because required contributions, particularly contributions to reduce underfunding, will place a larger financial burden on those employers. At the same time, CBO expects that employers in fully funded plans have a greater probability of withdrawing and switching to other forms of retirement benefits than do employers in underfunded plans that withdraw because of the plans' distress." These hypotheses may be tested by controlling for funded ratio squared.
- According to French (2017): "there appears to be some fairly clear and possibly predictive relationship between the number of contributing employers and the likelihood of pre-insolvency mass withdrawal." Again, this hypothesis is readily tested in a statistical model.
- French (2017) also noted differences in mass withdrawals across industries, which suggests industry indicators may improve the model.

Our companion report on contagion effects of mass withdrawals discusses several factors that, at least theoretically, should play a role in employers' withdrawal decisions. We recommend that future research identify metrics to capture those factors in a statistical model.

Withdrawals by Individual Contributing Employers

ME-PIMS does not currently model withdrawals by individual participating employers. Participating employers are assumed to either collectively withdraw or collectively remain in the plan.

A company that withdraws from a plan faces withdrawal liability payments. In principle, these payments aim to cover the unfunded vested liability attributable to the departing employer. A departing employer thus effectively shrinks the total unfunded vested liability of the plan. Ignoring individual withdrawals would therefore systematically overstate future liabilities. However, the withdrawal liabilities are merely assessed or estimated to be assessed and may not be collected in full. This issue is addressed in the next section.

Plans need to report on their Form 5500 filings the number of participating employers. We seek to explore the intensity of individual withdrawals over time and recommend whether the ME-PIMS model needs to be adjusted to accommodate such individual withdrawals.

Using data from the Form 5500 filings, we use two fields that contain information on withdrawals by individual employers. These are (1) the number of employers who withdrew during the preceding plan year (Schedule R, line 16a) and (2) the aggregate amount of withdrawal liability assessed or estimated to be assessed against such withdrawn employers (Schedule R, line 16b).

---

[21] Wendy Kiska, Jason Levine, and Damien Moore (2017). "Modeling the Costs of the Pension Benefit Guaranty Corporation's Multiemployer Program." Congressional Budget Office Working Paper 2017-04. Available at https://www.cbo.gov/system/files/115th-congress-2017-2018/workingpaper/52749-pbgcwp.pdf.

As Figure 19 shows, between 2010 and 2017, an annual average of 2,961 employers withdrew and were (estimated to be) assessed approximately $2.23 billion in withdrawal liabilities in the preceding year.

Figure 19. Amount of Withdrawal Liabilities (Estimated to be) Assessed and Number of Withdrawing Employers in Preceding Plan Year (2010–2017)

To place individual and mass withdrawals in context: The average annual number of individually withdrawing employers was 2,961, compared with 51 employers that withdrew annually as part of a mass withdrawal. The average annual withdrawal liabilities were $2.2 billion for individually withdrawing employers and $85 million for employers that withdrew collectively. In other words, individual withdrawals have been a much more important source of leakage from multiemployer plans than mass withdrawals in recent years.

In addition to individual employers leaving multiemployer plans, we also notice that there is a separate, general decline in the number of contributing employers to ME plans. A review of data captured on the Form 5500[22] shows this decline is larger than what can be explained by the number of individual employers withdrawing. This implies that employers are leaving multiemployer plans for reasons that are not classified as a "withdrawal."

Figure 20 shows change in contributing employers after accounting for individual withdrawals in the main industry sectors (Construction, Retail, Transportation, and

[22] Form 5500, line 7: Enter the total number of employers obligated to contribute to the plan (only multiemployer plans complete this item).

Other) between 2011 and 2017. Especially in the Construction industry, many employers appear to have left plans without having been reported as individual withdrawals; also see below.

Figure 20: Change in Contributing Employers Net of Individual Withdrawals for Major Industries, 2011–2017

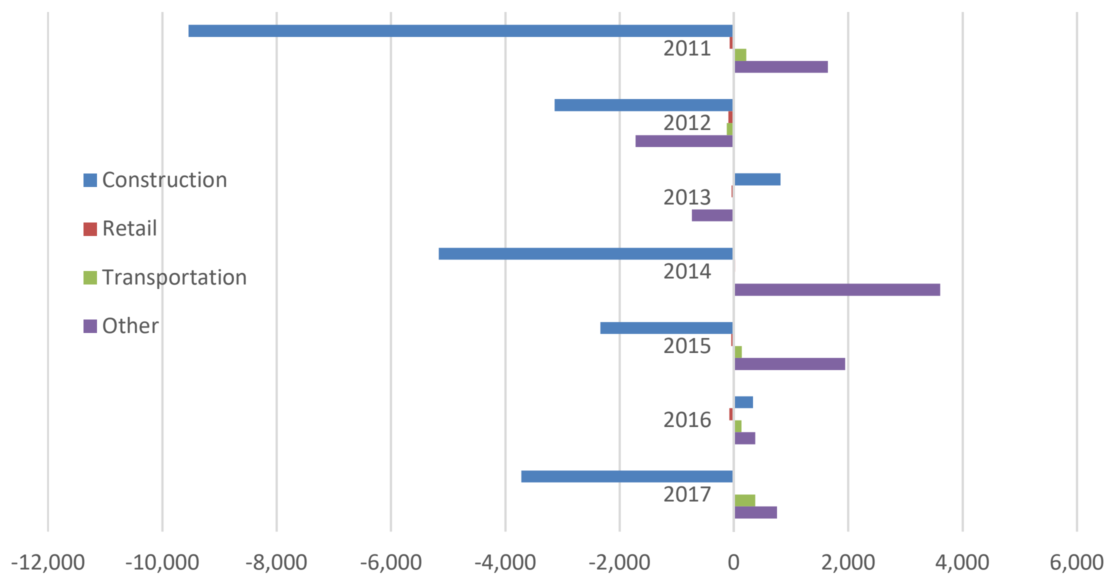

As a fraction of outstanding unfunded liabilities, assessed withdrawal liabilities averaged 0.43% over 2010–2017 (Figure 21). Recall that approximately 0.4% of multiemployer plans experienced a mass withdrawal each year. In that light, individual withdrawal liabilities are sizeable, and an argument can be made for incorporating them into ME-PIMS.

Figure 21. Ratio of Withdrawal Liabilities to Unfunded Liabilities, 2010–2017

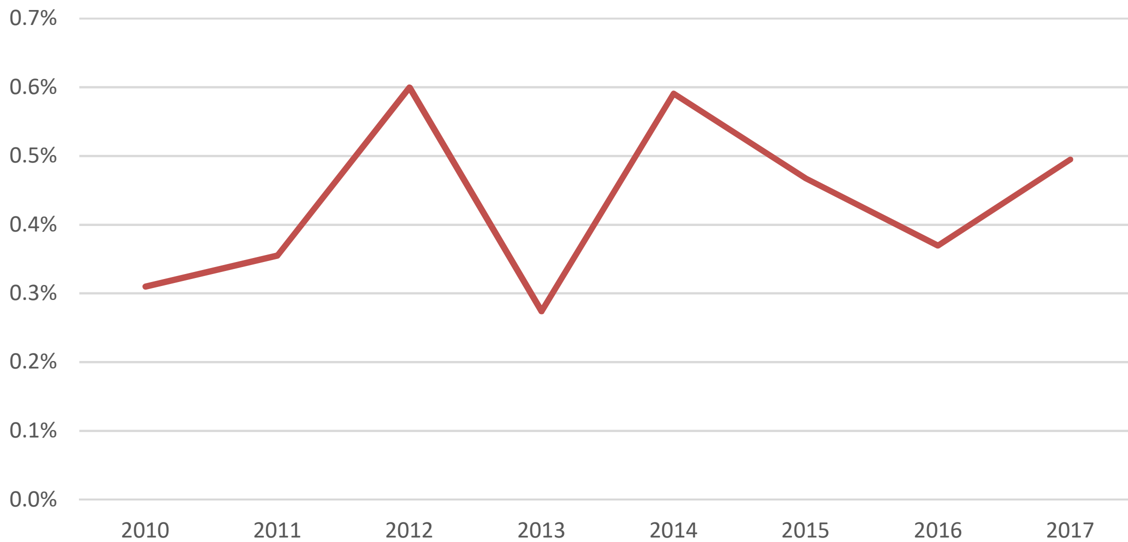

Withdrawals by individual employers correlate with industry and financial risk zone, as discussed next.

Withdrawals by individual employers are concentrated in certain industries. Between 2010 and 2016, withdrawals by individual employers from Truck Transportation, Retail Trade, Food Beverage Tobacco, Printing and Related, and Paper Manufacturing accounted for 64% to 84% of the total individual withdrawals (Figure 22). (In 2017, a large withdrawal by a company in the mining sector dominated and pushed this value down to 27%.) The Construction sector, which is subject to a different set of withdrawal liability rules, is conspicuously absent from Figure 22. The Construction industry accounted for almost 50% of the total unfunded liabilities in 2010–2017 but its withdrawal liabilities were low (2% of the total in 2017). This is presumably because the Construction industry and the Entertainment industry have special withdrawal and withdrawal liability rules.[23]

[23] Employee Retirement Income Security Act of 1974 (ERISA) §4203.

Figure 22. Proportion of Withdrawal Liabilities in Selected Industries (2010–2017)

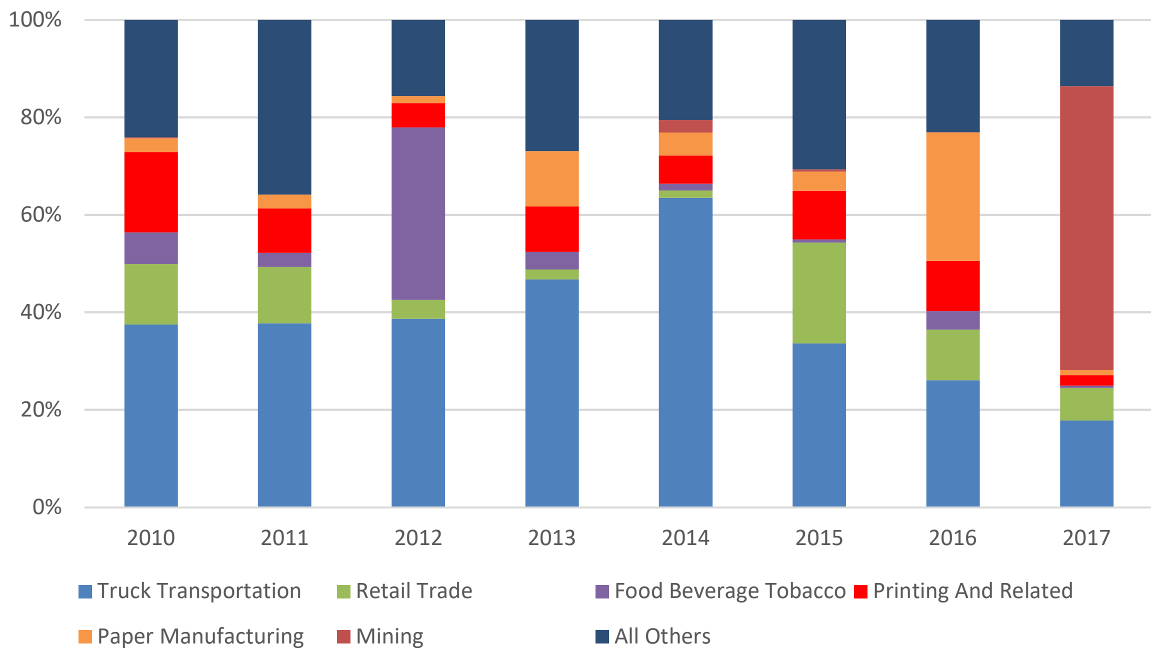

Separately, individual employers appear to be much more likely to withdraw from financially weak than from stable plans. Figure 23 shows that Critical (red) and Critical and Declining (purple) plans account for approximately 30% of current liabilities and 35% of unfunded liabilities during 2010–2017. However, these two risk categories account for over 80% of the withdrawals by individual employers during the same time period (Figure 24).

Figure 23. Proportion of Unfunded Liabilities by Plan Risk, 2010–2017

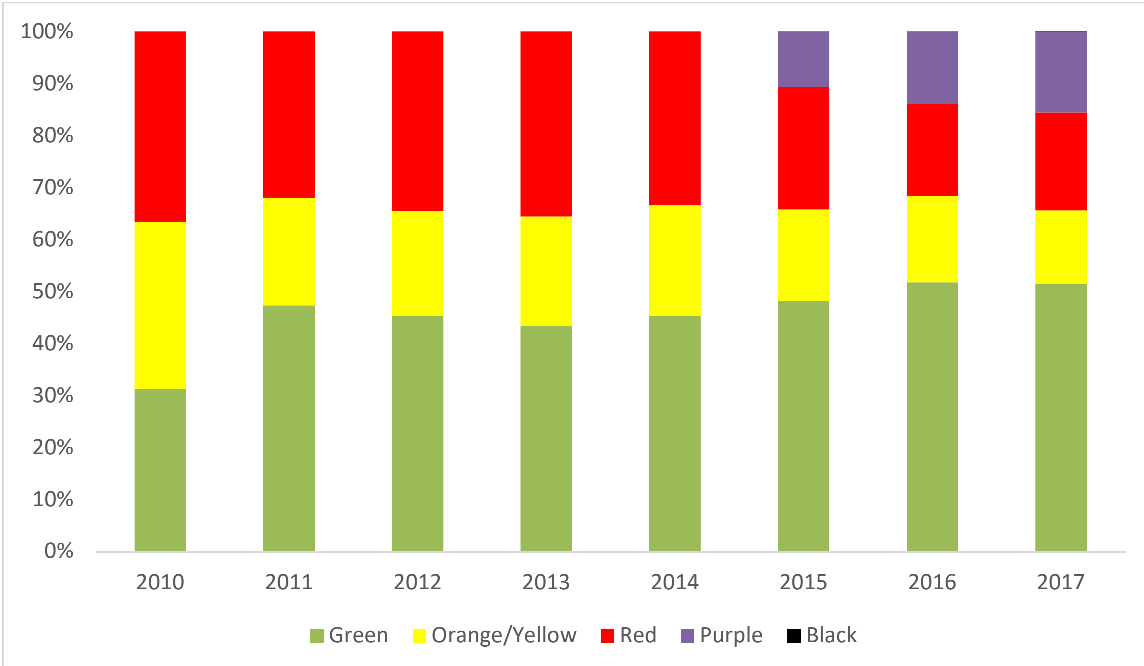

Figure 24. Proportion of Withdrawal Liabilities by Plan Risk, 2010–2017

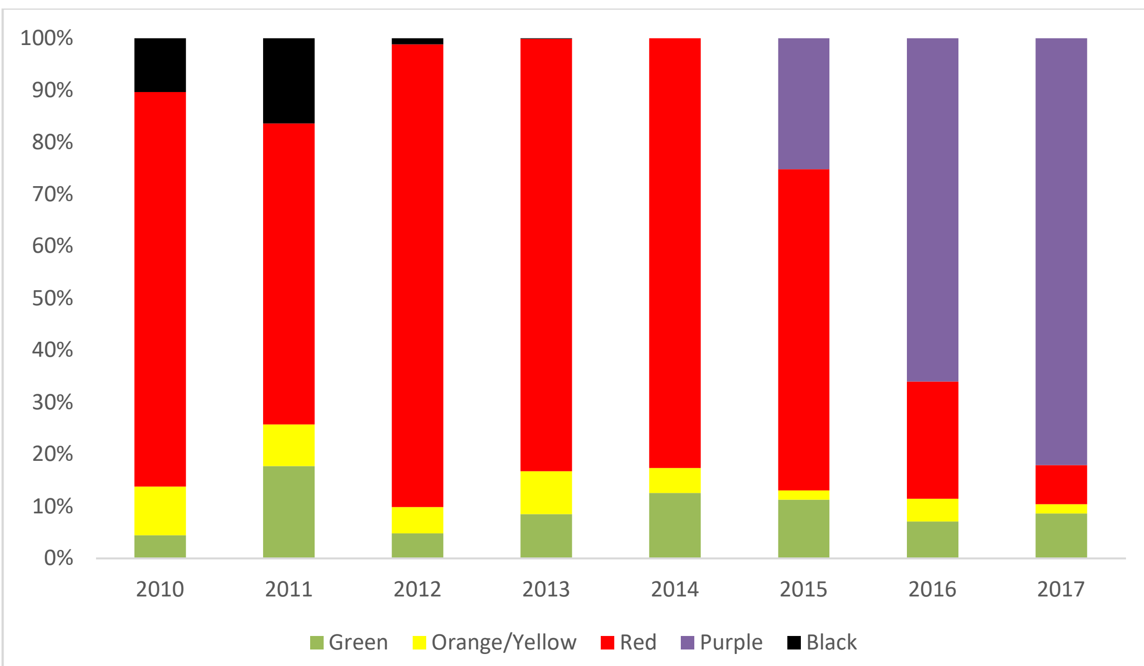

In summary, individual withdrawals are a sizeable drain on the contribution base of multiemployer plans, and their occurrence is related to industry and plans' financial status. If withdrawal liabilities were collected in full, the solvency of plans would not be dented by the departure of employers. However, if only a portion of unfunded liabilities are collected from formerly participating employers, individual withdrawals can hurt a plan's stability. The next section addresses this issue.

Collectability of Withdrawal Liabilities

For plans in mass withdrawal, ME-PIMS calculates a withdrawal liability that serves to pay off the plan's unfunded vested liability. The withdrawal liability is typically payable in installments that are roughly equal to historical contributions. ME-PIMS assumes a collectability rate of 40%, which is applied to the entire payment schedule. In addition, ME-PIMS assumes a rate of decay in the annual payment over the ensuing years, modifying the schedule downward accordingly. The assumed decay-rate is the average annual rate of decline, if any, in the actively employed population of the plan's collective sponsorship over the 5-year period immediately preceding mass withdrawal, and updated annually for future years in the simulation (*Key Differences Between SE-PIMS and ME-PIMS*, pp. 51–52).

We attempted to validate the collection assumptions in ME-PIMS. Unfortunately, the data currently at our disposal do not support this analysis. Schedule R asks about contributions by employers that contribute at least 5% of the total plan contributions. In rare cases, the filing indicated (on line 13e(2)) that the contribution represents a withdrawal liability payment. However, to our knowledge there is no complete accounting of withdrawal liability payments on the Form 5500.

It is our understanding that the PBGC recently conducted a survey among multiemployer plans to gain insight into withdrawal liability payments. Its results may be expected to shed light on collectability assumptions in ME-PIMS and, ultimately, the importance of modeling individual withdrawals in ME-PIMS. The issue affects not just plans with individual withdrawals, but also plans that experienced a mass withdrawal and possibly even insolvent plans.

Modeling the Population Rather Than a Sample

In 2017, ME-PIMS simulated future events for a sample of 321 multiemployer plans, which were weighted to represent the entire universe of PBGC-insured multiemployer plans (*FY 2017 PBGC Projections Report*). SE-PIMS is similarly based on a sample of DB plans, in part because little financial information is available for the majority of insured single-employer plans and their sponsors. In contrast, ME-PIMS relies predominantly on Form 5500 filings for its predictive modeling, i.e., data availability is presumably a lesser concern.

Analyses of a population based on a sample invariably require extrapolation, which introduces additional uncertainty. We explored whether ME-PIMS can include the population of multiemployer plans, i.e., approximately 1,400 plans. Among the obstacles to modeling the entire universe are the following.

- Not all information required by ME-PIMS is available in electronic format. All fields on the Form 5500 and its schedules are available, but PDF attachments are not. Some information that is important to actuarial calculations needs to be manually entered from PDF attachments. For example, pursuant to line 8 of Schedule MB, additional information may need to be attached regarding expected benefit payments over the next ten years, the numbers of active participants by age and years of service, average compensation or cash balances for each age/service category, retirement rates by age, and other factors that affect actuarial calculations.

- Even with sufficient data entry resources, the required information cannot be obtained for smaller plans because the requirements for detailed information generally apply to large plans only.[24]

As described in the *FY 2018 PBGC Projections Report* (which was released in August 2019), ME-PIMS was recently recoded and expanded to include all multiemployer plans. To solve the issue of incomplete data, plans were grouped by major industry sectors and by categories of the ratio of active to inactive participants. Detailed information was manually entered for a sample of 348 plans; the remainder of the plans were assumed to have similar provisions, age/service distributions, average salaries or cash balances, et cetera, as plans of the same industry and active/inactive category in the sample.

While we are unable to verify the implementation of this imputation process or evaluate potential unforeseen complications, we believe that the approach taken to extend ME-PIMS to the universe of multiemployer plans is sensible and scientifically sound.

[24] Such information is also not filed by plans that did not submit a Schedule MB. We found that all or almost all such plans are receiving financial assistance from the PBGC, i.e., the PBGC presumably has access to the required information.

4. CONCLUSION

This document reports on our review of bankruptcy aspects in SE-PIMS and mass withdrawal aspects in ME-PIMS. While sometimes hampered by limited available data, we found several issues that raise questions about embedded assumptions or relationships among simulated metrics. Some issues may be important; many others likely have only modest implications for the results generated by SE-PIMS and ME-PIMS. We did not identify any issues that, with a reasonable degree of confidence, would generate a sizeable bias in a known direction.

SE-PIMS and ME-PIMS are undergoing continuous development. Some of the issues raised in this report have already been addressed. For example, mass withdrawals are no longer assigned a stochastic probability of occurrence prior to plan insolvency, and ME-PIMS already simulates the entire universe of multiemployer plans (*FY 2018 PBGC Projections Report*).

Our review emphasizes issues that may have implications for projections of new claims to the PBGC. Much less emphasis was placed on the numerous cases in which we encountered something that raised questions but that, upon closer inspection, was in fact carefully thought through and properly addressed by SE-PIMS or ME-PIMS. Despite potential issues, the care and thoughtfulness with which the models have been developed is to be commended.

DISCLAIMER

The views, opinions, and/or findings contained in this report should not be construed as an official Government position, policy or decision, unless so designated by other documentation issued by the appropriate governmental authority.

We call your attention to the possibility that other professionals may perform procedures concerning the same information or data and reach different findings than Advanced Analytical Consulting Group, Inc. (AACG) for a variety of reasons, including the possibilities that additional or different information or data might be provided to them that was not provided to AACG, that they might perform different procedures than did AACG, or that professional judgments concerning complex, unusual, or poorly documented matters may differ.

This document contains general information only. AACG is not, by means of this document, rendering business, financial, investment, or other professional advice or services. This document is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action. Before making any decision or taking any action, a qualified professional adviser should be consulted. AACG, its affiliates, or related entities shall not be responsible for any loss sustained by any person who relies on this publication.

This document is the Part 2 Report of the Bankruptcy Component of Contract 16PBGC18C0023 (Pension Insurance Modeling System MAP-21 Peer Review Service).