

**A REVIEW OF THE PENSION BENEFIT GUARANTY CORPORATION PENSION
INSURANCE MODELING SYSTEM**

September 11, 2013

Jeffrey R. Brown, University of Illinois and NBER
Douglas J. Elliott, Brookings Institution
Tracy Gordon, Brookings Institution
Ross Hammond, Brookings Institution

Acknowledgements: The research team is grateful to Jeffrey Gianattasio and Katherine Waldock for exceptional research assistance. We thank Heitor Almeida, Stephen Dimmock, Dirk Hackbarth, George Pennacchi, Joshua Pollet, Jim Poterba and Scott Weisbenner for helpful conversations and comments. We are especially grateful to PBGC staff for their tremendous efforts in helping to answer our technical questions about the PIMS model, and to Paul Davies of the Social Security Administration for coordinating these interactions. We are also grateful to the U.S. Social Security Administration for funding this project through grant #5RRC08098400-05-00 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER.

Executive Summary

In December, 2012, the Social Security Administration (SSA) invited proposals from the Retirement Research Consortium to “independently review and evaluate the data, assumptions, and methods underlying models of the Pension Benefit Guaranty Corporation’s (PBGC) pension plan insurance programs and related models of pension funding and sustainability.” In response to this request, a team of researchers affiliated with the National Bureau of Economic Research (NBER) and the Brookings Institution prepared this analysis of PBGC’s Pension Insurance Modeling System (PIMS).

Our analysis suggests that the PIMS model was, in many ways, “state-of-the-art” when it was created approximately two decades ago. The use of stochastic simulation tools is a clear improvement over the deterministic model used previously. Among other benefits, a stochastic simulation model helps interested parties understand that there is a distribution of possible outcomes, not just an average outcome – a fact that is especially important for a program that is largely insuring against extreme events. It is also clear that the professional staff at PBGC has a deep understanding of both the capabilities and the limitations of the model. It is our impression that PBGC staff is committed to the principle that the PIMS model should be as unbiased as possible and insulated from political considerations.

However, several key components of the model have not been revised to reflect the availability of new tools, new insights from the academic literature, or even new data. PIMS has developed into a considerably more important tool for policymakers than was initially envisioned, but resources for PIMS have not risen commensurately, and budget and staffing constraints appear to have limited PBGC’s ability to keep the model up-to-date.

Our review also highlights three features of the existing governance system for overseeing PIMS: (i) some of the model documentation is internally inconsistent and outdated, (ii) the process for updating data and model parameters appears, at least to external observers, *ad hoc*, and (iii) there does not appear to exist any publicly-available, systematic inventory of the robustness checks that have been performed. Indeed, to the extent that methods or assumptions are tested, this fact is not documented in any central location, making it difficult to assess which features of the model are most critical. Other long-term models that are important to federal programs – such as the actuarial models underlying the report of the Trustees of the Social Security and Medicare programs – regularly undergo an external review by a technical panel of outside experts, a process that has led to continual improvement of those models over time.

A key finding of our review is that the limited treatment of correlated risk factors arising from the macroeconomic environment is likely to substantially understate the degree of fiscal risk to PBGC's insurance programs. This may be one reason that actual PBGC results have come out much below PIMS' median projections. In the PIMS model, there are very few avenues through which broader macroeconomic factors can operate directly on the distribution of potential future losses. In reality, however, macroeconomic factors directly affect many of the key drivers of PBGC's finances: for example, during an economic downturn, it is reasonable to expect more plan sponsors to experience financial distress and more plans to be underfunded. Consequently, the distribution of possible loss exposure has much "fatter tails" (that is, the probability of extreme losses is much greater) than is currently captured by the PIMS model. This matters because PBGC and other insurers have an asymmetric exposure to fat tails, being hurt more by the negative extremes than they are aided by the positive extremes.

Although our analysis focuses narrowly on the PIMS model, rather than broader policy questions about the pension insurance program, it is worth stressing that these extreme negative events are most likely to occur in states of the world in which the broader U.S. economy is relatively weak, which means that it would be a particularly economically painful time for the nation to have to address an underfunded pension insurance program. Recognizing the true economic costs of these correlated risks and how they affect the broader fiscal position of the U.S. government, therefore, has potentially important implications for program design, the average level of premiums, the question of whether to risk-adjust premiums, and other important policy parameters which are well beyond the scope of this narrow technical review of the PIMS model. Our review provides a number of specific observations about the model that could be used to guide future revisions to the model in this respect, particularly with regard to the modeling of the bankruptcy and financial market processes.

Table of Contents

Background.....	1
The Project	1
Pension Benefit Guaranty Corporation.....	2
Pension Insurance Modeling System.....	7
General Modeling Considerations.....	11
Financial Market Assumptions.....	16
Bankruptcy Frequencies and Severities.....	23
Other Modeling Choices Affecting Claims on PBGC.....	34
Issues Regarding the Composition of the Universe of Covered Plans.....	34
Contributions by plan sponsors.....	37
Pension liabilities	38
Discount rates	42
Projections of PBGC premiums	44
Effects of guarantee limitations on claims on PBGC	45
Special factors in ME PIMS.....	46
Peer review of Multiemployer PIMS	48
Historical predictive performance of PIMS.....	49
PIMS as a planning tool.....	52

Background

The Project

In December, 2012, the Social Security Administration (SSA) invited proposals from the Retirement Research Consortium to “independently review and evaluate the data, assumptions, and methods underlying models of the Pension Benefit Guaranty Corporation’s (PBGC) pension plan insurance programs and related models of pension funding and sustainability.” In response to this request, the National Bureau of Economic Research (NBER) submitted a proposal in January 2013 to assist with a review of economic, financial, and other aspects of PBGC’s single-employer and multiemployer models. A revised version of this proposal was accepted and funded by SSA in February 2013.

Over the past five months, the research team – consisting of researchers from the NBER and the Brookings Institution – has conducted an independent analysis of the PBGC models. (For more information about members of the research team, please see Appendix A). This paper represents the culmination of that analysis.

Although the analysis was conducted independently, the research team did have three formal meetings with PBGC staff (coordinated through SSA and including observers from SSA and the Departments of Commerce, Labor, and Treasury). The team also interacted with PBGC staff at the April 2013 Wharton conference on this topic. The research team relied on technical model documentation provided by PBGC as the authoritative source on assumptions and methods used in these models as well as using a few other key sources (see Appendix B for a discussion of these sources). In addition to consulting these materials, the team asked clarifying questions in the above noted conversations with PBGC staff. The team then compared the methods and assumptions used in PIMS to the existing academic literature on the relevant subject matter.

As noted on the NBER’s public website, the NBER is a “private, nonprofit, nonpartisan research organization ... committed to undertaking and disseminating unbiased economic research in a scientific manner, and **without policy recommendations**, among public policymakers, business professionals, and the academic community,” (emphasis added). Rather than making recommendations, our analysis is intended to provide an objective, independent, and research-based assessment of the PBGC single-employer and multiemployer models. We will point out areas where the modeling can be brought into closer alignment with the existing state of knowledge in the academic literature and research community. We will also highlight areas where particular modeling decisions, parameter assumptions or data choices may lead to biases in the outcomes projected by the model.

Given our areas of expertise, we have chosen to focus our analysis on the following areas suggested in the SSA call for proposals:

- Economy-wide economic modeling, primarily interest rates and rates of return
- Firm-level modeling of the financial health of firms that sponsor the pension plans
- Modeling of plan transitions, including the probability of bankruptcy and the probability that a plan or sponsor bankruptcy will generate a claim on PBGC plus the distribution of amounts recovered

As requested by the original request for proposals, our review provides an assessment of the strengths and weakness of modeling choices, including but not limited to the:

- Economic data and assumptions
- Actuarial data and assumptions
- Financial data and assumptions
- Modeling of plan requirements under the Employee Retirement Income Security Act of 1974 (ERISA)
- Modeling processes, calibration, and output
- Quality control and validation processes

In addition to these areas, there are many other topics related to PBGC that are of intellectual and/or policy interest. For example, there has long been a robust discussion in the literature on how to place a market value on the insurance that PBGC provides. Despite the importance of these questions, we focus our review on the PIMS model rather than on these broader economic or policy questions. Our omission of discussion of such issues should be viewed in the context of the purpose of this report – a review of the PIMS model –rather than as an implicit criticism or endorsement of any issue or policy.

Pension Benefit Guaranty Corporation

The Pension Benefit Guaranty Corporation, or PBGC, is an independent government corporation responsible for insuring the pension plans of about 43 million American workers and retirees. PBGC protects “defined benefit” pension plans offered by private sector firms. Traditional defined benefit plans offer a monthly pension payment for the life of a retiree, with the amount determined based on his or her past salary and work tenure. Many such plans now allow an employee to take a lump sum payment at retirement or provide payouts after retirement that depend on future investment returns. The PBGC insurance programs do not provide any coverage of 401(k) or other “defined contribution” plans.

PBGC was established in 1974 as part of the Employee Retirement Income Security Act (ERISA). Before ERISA, companies were generally allowed to choose whether, and to what extent, they prefunded their pension obligations or contributed to pension trusts. ERISA created minimum funding requirements, which have been periodically revised over time. However, many pension trusts continue to have assets that are considerably below the value of their obligations. These unfunded liabilities create a risk for the employee that the sponsoring employer may go bankrupt and be unable to fill the funding gap. PBGC was created to protect employees and retirees from the risk of losing pension benefits, up to certain specified limits, as a result of such bankruptcies. PBGC pension guarantee caps are set at a level high enough to ensure that a substantial majority of covered employees have received their full pensions even in cases where major firms have gone bankrupt and terminated their pension plans.

When a firm goes bankrupt and cannot meet its future pension obligations, PBGC assumes control of the pension trust's investment assets and takes over responsibility for the delivery of promised benefits. PBGC currently insures over 25,000 different pension plans. In addition, it pays benefits, or is committed to pay benefits in the future, to over a million individuals whose plans it has taken over as a result of bankruptcies. PBGC is charged with three purposes by ERISA, according to Section 4002 (a):¹

- To encourage the continuation and maintenance of voluntary private pension plans for the benefit of their participants
- To provide for the timely and uninterrupted payment of pension benefits to participants and beneficiaries under covered plans
- To maintain premiums established by the corporation under Section 4006 at the lowest level consistent with carrying out its obligations

As an insurer, PBGC relies on premiums and on investment income to meet the net costs of claims from bankrupt companies. Unlike most insurers, however, PBGC does not set its own premium rates, although it has frequently asked Congress for the right to do so. Instead, Congress has retained the power to set premium rates by amending ERISA from time to time, as it chooses. Also unlike most insurers, PBGC does not have the option of refusing to cover excessively risky plans; by law, PBGC must cover all private sector defined benefit plans that meet certain minimal criteria.

In the early 2000's, PBGC encountered severe financial difficulties, prompting Congress to pass the Pension Protection Act of 2006. The Act strengthened funding requirements for most existing pension plans and also raised premiums for plan sponsors with underfunded or newly terminated plans.² Pension rules were further modified by Congress after the recent financial crisis.

¹ These purposes are slightly paraphrased to be more understandable out of the original context.

² It simultaneously loosened funding requirements for the airline industry, given its financial troubles.

Currently, just under a quarter of all employees, current and retired, are covered by a private sector employer sponsored defined benefit plan that is insured by PBGC. PBGC itself encompasses two programs, the single-employer program and the multiemployer program. Each program insures different types of pensions. It is important to note that the single-employer and the multiemployer programs operate with separate accounts and according to distinct procedures.

The single-employer program guarantees benefits for 32.5 million individuals among 24,200 pension plans. It delivers benefits to 887,000 individuals and has commitments to deliver benefits to approximately 614,000 additional workers upon retirement. PBGC's single-employer program not only insures failing pension programs due to firm bankruptcy, known as "distress" terminations; it also initiates "involuntary" terminations of plans in cases where the plan meets one of four statutory tests, including the failure of the plan to meet minimum funding, the expected losses in the long run from continuation of the plan exceeding the costs of termination, or the plan lacking sufficient assets to pay benefits currently due.

Part of PBGC's role in guaranteeing retirement security thus involves monitoring the financial status of existing plans. Firms also have the option of undergoing a "standard" or "voluntary" termination where the employer chooses to terminate a pension plan and has sufficient assets to cover all of its existing liabilities by buying a qualified annuity from an insurer.

In addition to terminations, a plan sponsor can also amend its plan to freeze the accrual of new benefits. In a "hard freeze," no participant receives benefits for new service from the date of the freeze onward. In a "soft freeze," existing employees as of the date of the freeze can continue to accrue benefits for additional service, but the compensation levels used for benefit calculations are frozen. Plans may also be closed to new entrants, without affecting the benefit accruals of existing employees.³

When PBGC terminates a single-employer plan, it assumes control over both the investment assets and the future obligations of that pension. Under the PBGC single-employer program, the maximum annual benefit a 65 year old retiree can receive is \$57,477 per year for a participant in a plan terminating in 2013. The cap is lower for those who retire before that age. Because some covered pensioners are relatively well compensated, the PBGC cap can mean a reduction in benefits, especially for those who have taken early retirement.

In addition to absorbing assets from terminated pension plans, the single-employer program collects a per participant premium from each plan. However, the premium level has been insufficient to prevent the PBGC single-employer program from amassing a \$29 billion dollar deficit as of September 2012.

³ The terminology here is based on PBGC's usage. Some others in the pension sector use "soft freeze" in different ways.

The multiemployer program is smaller than the single-employer program, with only 10.4 million employees covered, representing 1,450 plans. Of these individuals, 39% were active employees and 61% were retired or vested participants separated from employment. In 2012, PBGC's multiemployer program provided financial assistance to plans covering some 51,000 individuals plus 21,000 terminated vested employees.

Multiemployer plans typically provide benefit levels bargained between one or a few unions and a number of different employers, generally in the same industry. Because these plans are supported by many firms rather than a single company, PBGC's multiemployer program operates differently in the case of financial stress than does the single-employer program. Rather than assuming control of the investment assets of the plan, PBGC most often serves as a source for emergency loans for multiemployer plans. These loans ensure that distressed plans are able to pay at least the insured portion of their promised benefits. In the past, multiemployer programs were more resilient than single-employer plans. However, in recent years the multiemployer program has faced significant financial difficulties, and more challenges are anticipated in the near future.

In 2011, over 20% of protected multiemployer plans were assessed as having significant and immediate funding problems. Often, distressed multiemployer plans do not recover and end up paying significantly less in benefits than promised. PBGC's maximum guarantee in these cases is considerably lower than for the single-employer program. For example, at 30 years of employment, the maximum annual guarantee PBGC can extend to multiemployer plans is \$12,870 per participant.

Like the single-employer program, the multiemployer program also charges a premium per plan participant. However, these premiums are quite small. Over the next decade PBGC's multiemployer program anticipates collecting \$1.3 billion in premiums, whereas potential obligations are projected to rise by \$37.6 billion. Because of this mismatch, PBGC estimates that absent premium increases or Congressional action, the multiemployer program has a 36% chance of insolvency by 2022, and a 91% chance of insolvency by 2032. In 2012, this program had a deficit of \$5.2 billion, an increase from \$2.8 billion in the previous year (PBGC 2012).

PBGC's financial condition at any point in time depends on a number of key variables. At the most basic level, PBGC's revenues and expenses consist of:

PBGC Premiums: Congress requires that all covered single-employer pension plans pay an annual premium of \$42 per participant (participants include employees, former employees with a vested pension, and retirees). In addition, plans must pay a variable premium if they have assets less than the value of their pension promises. The current level of variable premiums is 0.9% of the underfunding of vested benefits and that rate will rise over time according to a formula set by Congress. For multiemployer plans, the premium rate for 2013 is \$12

per participant with no variable premium. The premiums for both programs are set to increase each year beginning in 2014, in line with inflation. Thus, PBGC's total premium levels depend on the number of participants and the level of underfunding of the pension plans that it insures, in conjunction with premium rates set by Congress.

Investment income on PBGC's assets: PBGC owns investments purchased with premium revenues in addition to assets received through the take-over of the pension assets of plans with insolvent sponsors. These investments, primarily in bonds and stocks, earn income. PBGC's investment income depends on the total size and composition of its investments and the rate of return on those specific investments each year.

PBGC's own expenses: PBGC is an organization of several thousand employees and contractors and has the operating expenses that one would associate with such an entity, including compensation. PBGC's expenses depend primarily on its volume of activity, in terms of new claims, the number of participants in payout mode, and so forth, plus a certain level of fixed or quasi-fixed expenses such as real estate.

Insurance claims on PBGC: As an insurer, PBGC ends up paying claims. When a single employer plan is taken over by PBGC, the insurance program takes over the assets and liabilities of pension plans where the liabilities for PBGC-provided benefits exceed the assets. Thus, there is the expectation that each claim will cost PBGC money overall, over time, even though initially the claims bring an influx of assets from the plan that is taken over that more than cover the initial payouts. The level of claims on PBGC in a given year is determined by a complex set of factors, including:

Frequency and severity of bankruptcies: PBGC primarily covers pension funding deficiencies of single employer plans whose sponsors are bankrupt or insolvent. Therefore, critical determinants of claims on the insurer are the number and size of bankruptcies in a given year. Bankruptcies, in turn, reflect a combination of general economic and financial market conditions as well as idiosyncratic factors.

Limits to PBGC guarantees: The insurance provided by PBGC contains several limits, including a cap on the amount to be paid to any participant. The effect of these limits depends heavily on the benefit design of the plan, particularly its generosity, as well as the average compensation level and tenure of employees and certain technical factors.

Initial funding levels at each covered plan: At the beginning of each year, every plan has some level of existing assets and liabilities that determine how well funded it is.

Contributions by plan sponsors: Employers make contributions into their pension trusts each year. Federal law sets formulas for minimum contributions, but sponsors (especially single employer plan sponsors) often contribute more than those minimums.

Investment income in pension trusts: Pension assets are invested, principally in stocks and bonds, and these securities and other investments earn income each year. The amount is determined by the size and composition of the investments and the rate of return on the specific investments.

Pension payouts: Each year, retirees receive benefits and end the year with one less year of expected future payouts. Payouts reduce assets and the reduction in remaining life expectancy of the starting retirees reduces liabilities. Payouts are determined by benefit formulas, salary levels, rates of hiring, firing, and retirements, and other factors.

Increase or decrease in pension liabilities: Each year, pension liabilities go up or down as the number of remaining participants changes and as a variety of actuarial factors, such as mortality rates and interest rates, vary.

Additional factors for Multiemployer pension plans: Unlike single employer claims, multiemployer insurance program claims do not generally result from employer bankruptcies or insolvencies. The employers contributing to a multiemployer plan have joint and several liability as long as they remain in the plan, meaning that each is wholly responsible for providing any funds the plan needs to pay its full liabilities. However, there are provisions in the law that permit employers to withdraw from the plan and related requirements for them to pay a withdrawal liability to cover their share of the existing underfunding plus a penalty to cover the costs of performing a valuation of the plan's total assets and liabilities. However, there are also caps on the payments that affect these calculations and often leave the withdrawal payments well below the expected value of the employer's share of the liabilities. Thus, the estimation of multiemployer plan claims requires an examination of the potential for withdrawals by employers, whether individually or in a mass withdrawal of all employers.

Pension Insurance Modeling System

To assess its future obligations and financial position each year, PBGC employs a detailed microsimulation model known as the Pension Insurance Modeling System (PIMS). There are two versions of PIMS – one for the single-employer program and one for the multiemployer program.

The PIMS model allows PBGC to forecast a distribution of potential economic scenarios and identify the level of exposure these scenarios would produce for PBGC.

Determining the future financial position of PBGC is a complex process, as this position depends on the state of the economy, the behavior of individual firms, the status of specific pension plans, the starting position of PBGC's finances, and the policy parameters set by relevant laws and regulations.⁴

PIMS uses a probabilistic (or "stochastic") modeling procedure designed to assist PBGC in quantifying its level of risk by computing the probabilities associated with various economic futures. In this task, the PIMS model relies on historical data on a variety of factors, such as the incidence of bankruptcy and the effects of key economic variables on the level of funding in the universe of pension plans. This historical analysis rests upon the assumption that the future volatility of these variables is governed by the same dynamics underpinning their past performance.

The historical data includes bankruptcy data from 1980 onwards, mortality tables constructed from demographic data, as well as data on stock returns and interest rates from 1926 onwards. These data then inform the choice of parameters used for the stochastic model's future projections. For example, two key variables in the PIMS model are the (real) rate of return on common stock investments and the interest rate. Although in any particular PIMS simulation the values for these key variables are determined by a random walk process, the overall probability distributions for these variables across all simulations are derived from their long-run historical behavior. However, the stochastic nature of the model allows PBGC to assess their financial position in the context of a range of potential economic futures, including those without historical precedent.

Before PIMS, PBGC relied on a deterministic forecasting model. The pre-PIMS model involved three separate estimations of risk: a forecast based on the entire history of PBGC claims (forecast A), a forecast using only the most recent 14 years (forecast B), and an estimation of the total claims on PBGC in the event that over the next decade all plans with a reasonable chance of placing a claim on PBGC did so (forecast C). There were a number of problems with this approach, most importantly the absence of calculated probabilities and thus little capacity for assessing the relative likelihood of the three scenarios.

The first use of PIMS was in 1998, when PBGC ran it alongside its former model. In contrast to the previous method, PIMS was designed to calculate a broad distribution of possible scenarios, each with its own specific probability. The added complexity of PIMS was made possible by more sophisticated computational technologies, and the initial deterministic approach was superseded by the combination of random elements with deterministic processes in the new stochastic simulations. While PIMS still does not enable PBGC to predict individual future claims, its projections do offer a range of potential levels of exposure that enable

⁴ See PBGC, (1999), "Pension insurance data book 1998," specifically pages 10-17, (subtitled: "Pension Insurance Modeling System: An Introduction"), for an excellent overview of the PIMS model. The background given here is a summary of the description found in that document.

PBGC to better understand its level of overall risk moving forward. Following the single-employer version of PIMS, PBGC developed a multiemployer version of PIMS as well, the results from which were first released in their 2009 Annual Report.

The PIMS model involves the interaction of five distinct modules, each modeled according to its own particular characteristics. These modules are as follows:⁵

Economy: The general macro-economic climate, crucially interest rates and the prevailing rate of return in the stock market. These are modeled using stochastic processes.

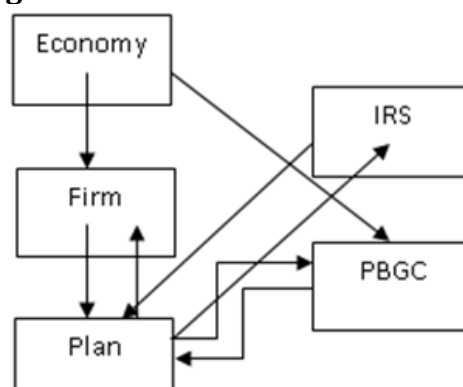
Firms: The financial and employment conditions specific to a firm, with special attention to how these variables impact the level of contribution to that firm's pension plan as well as their impact on its probability of bankruptcy.

Plan: The funding level, demographics, assets, and liabilities of specific pension plans.

IRS: The rules set by the Internal Revenue Code and ERISA, such as funding rules and the minimum level of contribution, which establish the policy parameters for the simulation as a whole.

PBGC: The initial financial assets and liabilities of PBGC, the expected value of premium collections, as well as an estimate of investment returns for PBGC's asset portfolio.

Figure 1: The Five Modules of PIMS



Source: Anderson, 1999, pages 11-12.

Figure 1 shows how PIMS modules interact with and inform each other.⁶ Simulations begin with known initial values, which include: macroeconomic

⁵ The following discussion of the modules and simulation procedures used in PIMS draws heavily on Anderson, (1999), "Models for retirement policy analysis."

variables, firm or plan sponsor variables (cash flow, employment levels, debt to equity ratios), pension plan variables (total liabilities and assets, salary information on participants), and the initial conditions of PBGC itself (such as assets, liabilities, current benefit payments).

The actual procedure used by PIMS is as follows. First, the simulation parameters are set. These parameters specify the number of scenarios to run, number of years per scenario, number of firms involved in a simulation, and other macro-level specifications. Second, the initial values for all modules are entered, including the policy parameters governing contributions and benefits. Then, the temporal parameters are set to describe interest rate behavior, bankruptcy incidence, portfolio performance, and demographic trends (such as mortality rates, retirement behavior, and employment levels.)

PIMS then runs a large number of macro-scenarios (500), which estimate a potential macro-economic climate across the specified years of the simulation using the above specified stochastic modeling procedure. Each scenario consists of a time path of up to 20 years, with PBGC paying special attention to and reporting the forecasts for 10 years into the future.

For each macro-scenario, PIMS runs at least 10 “cycles” for each individual firm in the single-employer model – these are stochastic microsimulations specific to individual firms which allow for a greater degree of accuracy in the model overall. Theoretically, this cycling procedure allows PBGC to better understand its exposure to risk from firms that represent a relatively large share of the overall system. However, because these cycles are firm specific, they may miss potential scenarios of intra-industry “linked risk” that may precipitate a large increase in liabilities. We return below to a discussion of contagion and correlated bankruptcy risk.

It is important to note that the firms and plans included in the PIMS model do not encompass the entire population of plans that are insured by PBGC. Instead, the single-employer PIMS model includes over 300 firms sponsoring about 400 plans, and the multiemployer program PIMS includes fewer than 200 plans. However, firms and plans are chosen for their relatively large share of PBGC’s overall liabilities (within each program). Thus, the firms and plans included in each version of PIMS comprise about half of PBGC’s total liabilities.

While the number of scenarios run can vary, PBGC has found that at least 300 macro-scenarios, with at least 10 “cycles” per plan sponsor, are required for PIMS to produce stable results. Balancing computational limitations and statistical accuracy, PBGC typically runs around 500 macro-scenarios, yielding 5,000 simulation outcomes after including each of the cycles of plan sponsor behavior. The outputs of

⁶ Although not displayed in the figure from Anderson (1999), we note that one may also add an arrow flowing directly from “Economy” to “Plan” to reflect the manner in which asset returns and discount rate assumptions flow directly.

these thousands of simulations are then used to populate a probability density function of the potential financial positions of PBGC in 10 years. PBGC reports this table, with specific attention to the mean expected financial position, the median expected position, and the worst expected case with at least a 15 percent projected probability of occurring.

General Modeling Considerations

In this section, we consider the overall modeling approach, techniques, and procedures used for PIMS, and discuss potential variants or extensions that may warrant further examination. Topics include simulation methodology, sensitivity analysis and modeling testing, communication of results and stakeholder engagement, and computational efficiency.

Modeling Approach and Methodologies

The literature on modeling complex stochastic dynamics offers a wide range of methodologies. These fall roughly along a continuum – from compact, top-down, mean-field analytical approaches (that is, ones that ignore geography or networks) at one end to bottom-up, simulation approaches that explicitly consider not only individual dynamic trajectories but also interaction effects and adaptive behavior at the other. The current implementation of PIMS as a microsimulation with non-interacting, non-adaptive actors is an approximate mid-point in this continuum. This approach offers a nice balance of detail and parsimony, and is well suited to the specific datasets it uses. However, the overall PBGC modeling effort would likely benefit from comparative analysis using other “points” on the continuum of modeling approaches as well.

For example, comparison of PIMS to a much simplified, equation-based analytical model—both at the level of assumptions and at the level of dynamic output—would likely be revealing. Such a comparison could help PBGC and stakeholders to “unpack” and explicitly discuss the impact on results created by each of the more complex or uncertain assumptions in the model (also see *Sensitivity Analysis* below). Where a simpler model or assumption produces comparable results to a more complex one, the efficiency of computation, interpretation, and communication can be improved by simplifying (see *Computational Demands* and *Communication* below). Where the two differ, the comparison is likely to produce insight into the specific role each assumption plays. We note that the Congressional Budget Office has an options-based model of PIMS that also provides another useful basis for comparison, although in a very different and less detailed fashion.⁷

⁷ See CBO, (2005a), “The risk exposure of the Pension Benefit Guaranty Corporation,” for a detailed comparison between these two models. For a more general analysis of the dynamics underlying

Another possibility that PBGC has considered is to develop a PIMS-light, a model that uses a much less detailed actuarial component in order to allow users to concentrate on simulating the other key variables, such as bankruptcy rates, overall economic factors, and financial markets factors. Such a complementary model would have many advantages in considering non-actuarial factors in a less complex structure.

On the other end of the “continuum of complexity,” the current analysis might benefit from enriching the model in key places where it may be leaving out important dynamics. These dynamic interactions might include:

Consideration of contagion (as in bankruptcy or plan freezing):

Simulation modeling in epidemiology (Epstein 2009; Eubank et al. 2004) has demonstrated that accurate modeling of contagion often requires considering the spatial structure of links between actors or networks. In this case, it would likely add to the potential accuracy of PIMS forecasts if the model could take into account the fact that any given firm's bankruptcy probability at any time (t) may be affected by previous ($t-1$) state of firms to which the focal firm is linked (through business transactions or industry networks such as airlines, steel, and so forth). Capturing this dynamic would involve two changes: first, moving from a model in which individual firms are passed through the model in sequence to one in which they go through in parallel; second, developing an empirically-grounded way to assess connections between firms. The latter is certainly challenging, and actual network information may not be readily available. However, it might be possible to use a statistical approximation and historical data to estimate this kind of contagion.

Consideration of adaptive behavior by firms: Firms in the real world are likely to adjust their behavior adaptively based on their own financial state, the financial health of close competitors in the same sectors, and the overall state of and trend direction in the economy. For example, a firm's choice of when to freeze or discontinue a pension policy in a multiemployer setting is likely to depend in part on observation of the financial health of peers. Similarly, under (or over) funding of pension plans by firms will depend in part on poor (or exceptionally good) economic performance. The decision by firms of whether to offer a pension plan in the first place may also be adaptive—potentially skewing the sample of who is in the PBGC coverage pool to begin with (sufficient membership by “healthy” individuals is of high concern to health insurers, for example).

PBGC's finances and prompting the development of the CBO model, consult CBO, (2005b), “A guide to understanding the Pension Benefit Guaranty Corporation.”

Capturing these types of dynamics would likely involve moving from a non-interaction, non-adaptive microsimulation (as in current PIMS) to something more like an agent-based microsimulation (Epstein 2007, Hammond 2009, Tesfatsion and Judd 2006).

Best Practices for Modeling

As the field of computational modeling has matured, a set of best practices has been well-established to help ensure that the results from modeling are generated in a rigorous way and interpreted appropriately (Tefatsion and Judd 2006; Saltelli, Tarantola and Campolongo 2000). Several of these practices are especially important for stochastic modeling, and for modeling that may inform policy decision-making. PIMS fits into both of these categories. While PBGC is following many important best practices, there are others that do not appear in the documentation of PIMS that we reviewed and are thus either not implemented or were not included in the documentation we received.

Sensitivity Analysis

Complex models generally involve a number of important conceptual or logical assumptions and draw on empirical estimates for multiple quantitative parameters. Understanding the sensitivity of model results to variation in both conceptual assumptions and quantitative parameter values is critical for appropriate interpretation of results. This can be especially important when models are stochastic, involve multiple dynamic mechanisms, or rely on limited data samples or uncertain empirical estimates. While it is important to use the best available evidence to inform baseline assumptions, variation around these estimates—both individually, and in various combinations—is critical for understanding the robustness or brittleness of the findings and for providing guidance to decision-makers. In recent years, sophisticated computational algorithms have been developed for use with microsimulations to allow “multi-dimensional sweeps” in which key parameters are co-varied to explore wide ranges of parameter space. Equally important can be confronting the model with extreme or rare, but potentially highly significant, events—draws from the “tail” of a distribution—to assess its ability to appropriately handle projections under such circumstances.

PIMS fits several of the criteria that make application of broad sensitivity analysis particularly important. It relies on a fairly small sample of historical data, involves estimates of key parameters that are based on data but are necessarily of limited accuracy, and makes several important conceptual assumptions designed to facilitate the analysis. Capturing and communicating the impact on results of variations in these assumptions and parameter values is important for developing confidence in the model’s findings, both internally and externally. If the model’s expected accuracy is highest within clearly defined bounds, this is important to state clearly.

Model Testing

Another important best practice designed to improve confidence in model findings and predictive or explanatory power is retrospective (or *ex post*) performance testing. This process involves initializing the current model (time t) with data from an earlier period of time (for example, $t-2$), and comparing the results it produces in the interim (time $t-1$) to those observed in actuality. If the model is able to “retrodict” (or reproduce an existing historical case or data set with high accuracy), this improves confidence in its projections into future (as yet unobserved) time periods. For example, the liabilities incurred by PBGC since the 2008 financial crisis offer a rare opportunity to test the robustness of the model's results during a period of high volatility (and refine its assumptions as necessary). Parameterizing the model as in 2003 allows calculation of the probability PIMS *would* have estimated for the observed ten-year PBGC liabilities—including during the period of volatility. If the actual values lie within the projected distribution, this would considerably strengthen confidence in the model's current ten-year projections. If not, it might be important to understand exactly why not—which sets of assumptions or parameters drive divergence from the observed outcome. We recognize, of course, that any testing of this type must account for other changes over the time period in question – including legislative changes – that might complicate the comparisons.

Sub-modules of the full model can also benefit from more direct testing. For example, the population of “virtual firms” created by PIMS (a key input into the microsimulation process) could be more fully compared to the full universe of real-world firms along important dimensions.

Communication and Stakeholder Engagement

Communicating both the results and the design of complex models to stakeholders without technical backgrounds is often challenging—but can be critical to gaining acceptance and maximizing the utility and impact of modeling. One strategy for increasing accessibility that has been widely adopted in fields such as public health (Epstein 2009, Hammond 2009) and business (Sterman 2006) is to create *interactive simulation interfaces*. These interfaces allow non-technical audiences to interact with and generate results from a model without requiring a technical background or detailed engagement in the complexities of the model's design. In these other contexts, stakeholders are able to vary the stochastic assumptions and view firsthand the predictions that each set of assumptions produces. Highly visual model interfaces help to facilitate this type of interaction. By communicating clearly the key input assumptions, the key outcomes, and the dynamics that produce one from the other, this approach has proved highly effective in situations where decision-makers from varied perspective and backgrounds wish to make use of complex models to inform important policy choices (see, for example, the work by the National Institute of Health Models of Disease Agent Study network).

Expanding Application and Utility of PIMS

As currently designed and deployed, PIMS includes consideration of multiple macroeconomic scenarios, of multiple firm and plan universes, and of multiple investment return and plan default scenarios. However, the model itself does not incorporate policy uncertainty, that is, the impact of variation in decision-making by regulatory authorities.⁸ In other areas where modeling is used to inform policy (such as retirement, public health, etc.), models offer great utility as “virtual laboratories” for exploring policy alternatives in a simulation and understanding the potential impact different choices might have. PIMS, or an extension of it, could also be used in this way. With complex stochastic dynamics, what appear to be simple dependencies (for instance, the impact of a premium rate change today on PBGC’s financial position in the future) may actually be far from straightforward. Modeling, and especially full sensitivity analysis (see below), can help inform such decisions.

Computational Demands and Resources

As with all computational models, the number and precision of analyses that can feasibly be conducted using PIMS may be determined in part by computational speed and resources. The initial deployment of PIMS and its design as a microsimulation were made possible in part by the existence of sufficient computation power. However, many of the possible extensions of the model discussed elsewhere in this report (and particularly those concerning multi-dimensional sensitivity analysis in the section above) would, if implemented, entail increases in computational processing demands.

Fortunately, in the last decade, there have been enormous advances in both computational hardware speeds and the efficiency with which microsimulation models (including agent-based models) make use of computation. For example, it is now possible to conduct agent-based simulations of spatial contagion across large geographies and with billions of individuals without specialist hardware (Epstein 2009). The ability to implement some of these model enhancements would require that PBGC has sufficiently modern and powerful computational hardware at their disposal (whether in-house or outsourced on an as-needed basis), and that PIMS code is optimized for maximum computational efficiency.

⁸ Although PIMS does not incorporate policy uncertainty, users of the model (including PBGC staff) can run the model under alternative modeling assumptions. In fact, PBGC staff tells us that the model has been used extensively in this manner in recent years to inform legislative and regulatory policy discussions.

Financial Market Assumptions

As noted elsewhere, the insurance program provided by PBGC is insuring underfunded pension benefits provided by a defined benefit plan sponsor that is experiencing financial distress. Both aspects of this risk – the probability of bankruptcy and the funding status of the bankrupt firm’s plan – are directly affected by realized returns in financial markets. Further, PBGC’s own asset portfolio is also invested in financial assets, providing a further avenue through which asset returns matter for the long-term health of the insurance program. Indeed, macroeconomic factors and the accompanying financial market returns are a source of correlated risk across these various elements. When the economy is strong and asset prices are rising, fewer firms may experience financial distress, plan funding ratios rise, and PBGC’s own asset portfolio may grow. Of course, the reverse is also true in states of declining asset valuations.

Given this economic foundation, financial market assumptions play an important role in assessing the distribution of possible future PBGC financial outcomes. The PIMS model simulates a range of financial market outcomes, including nominal and real interest rates, corporate yield spreads, and inflation. The model simulates these values according to a set of reduced form equations, parameterized based on historical data, and then uses them to compute items such as the growth in plan sponsor pension assets. In this section, we discuss some of the basic features of the PIMS models, and highlight a few areas where PBGC’s assumptions do not fully reflect the existing body of knowledge. Although PBGC has continuously improved the PIMS model since its creation, there are a number of elements where the model does not fully reflect the considerable advances in the literature that have occurred in recent years.

Overall, it appears that the existing model underestimates the likelihood of extreme or “tail” events. This observation is especially important given that these are the states of the world that are likely to have the most significant impact on long-term PBGC finances. Further, there is an asymmetric effect, in which poor financial market returns generally hurt PBGC as an insurer substantially more than good financial returns aid PBGC. Insurers essentially offer a “put” option in which they absorb the losses from bad events without benefitting directly from good events. PBGC would have lower losses if financial market returns are unexpectedly good, but the great bulk of that benefit would go to plan sponsors. PBGC absorbs a greater proportion of the losses when returns are sharply lower.

Real Interest Rates

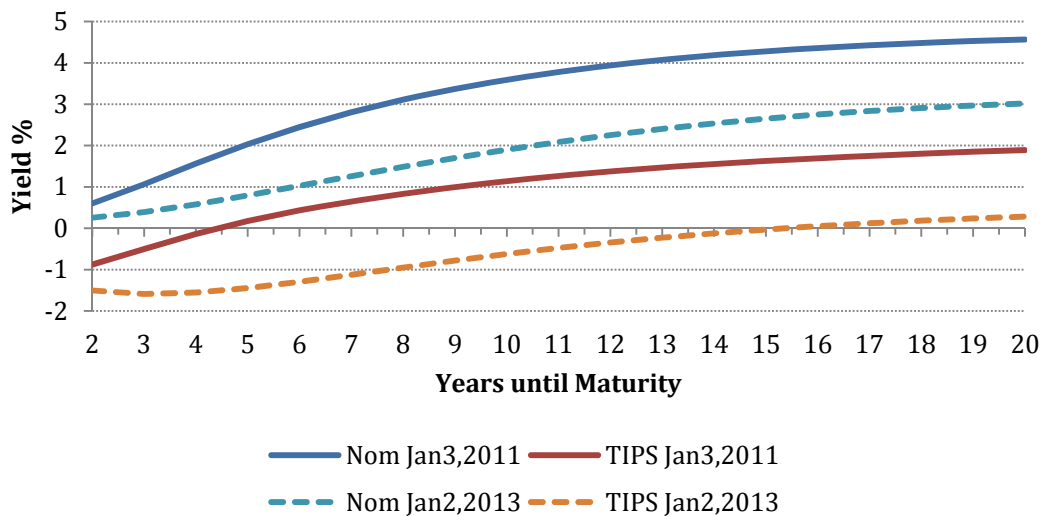
It appears from the PIMS documentation that the real interest rate in PIMS is fixed by user input, with a default value of 1.64%. This treatment of interest rates is clearly unrealistic. We note that, in the early 1990s, the U.S. did not yet issue TIPS (Treasury Inflation Protected Securities), and thus the modelers had less

information available for modeling real returns than we have today. The academic literature on how to model real term structures has also advanced significantly since that time. Given these advances, we make three overall observations about the treatment of interest rates in PIMS.

First, from a “model governance” perspective, it is unclear how PBGC goes about choosing the value of the fixed real interest rate parameter, or by what process changes to this parameter are considered. It may be that PBGC has a robust internal process for choosing how this is to be determined (as well as who has authority to make changes to the assumption), but this was not easily discernible from the documentation provided.

Second, the PIMS model assumes a flat term structure for both real and nominal yields, that is, there is no distinction between short-term and long-term interest rates. Yet even a casual observation of yields, as illustrated in Figure 2, clearly indicates that real interest rates are not independent of a bond’s maturity. For illustration, this Figure shows the yield curve for both nominal Treasuries and TIPS for January 2011 and January 2013.

Figure 2: Yield Curves



The upward slopes of the yield curves in this Figure are not unique to the specific dates shown: Haubrich, Pennacchi and Ritchken (2012) show that real interest rate risk premia increase with maturity, and thus it is typical for rates to be higher at longer time horizons. We note that in recent years, short-maturity real rates have been negative and have become positive only at horizons greater than four years or more.

In periods of a steep yield curve, the PIMS model’s assumption of a flat yield curve has the effect of under-stating the importance of nearer term liabilities relative to longer-term liabilities, for example. It should be relatively straightforward for PBGC

to incorporate the full term structure. At a minimum, it ought to be able to break it into several discrete components (for example, short, medium and long-term).

Third, the PIMS assumption of a fixed real interest rate across periods is in sharp contrast to the empirical distribution of real interest rates, which do, in fact, vary considerably over time. This result can also be seen from a comparison of the solid lines (January 2011) to the dashed lines (January 2013) in Figure 2 above, which clearly shows that the entire yield curve shifts over time. There is an extensive literature regarding interest rate processes, and discrete versions of Vasicek (1977), Cox, Ingersoll and Ross (1985) or Longstaff and Schwartz (1992) would provide a starting point if PBGC wishes to incorporate changes along these lines.

Nominal Rates and Inflation

According to PIMS documentation, the nominal yield for 30-year Treasuries follows a random walk, with the model parameters estimated from historical Ibbotson data. The empirical foundation for the random walk assumption is unclear; Haubrich, Pennacchi and Ritchken (2012) model nominal and real yield curves and find strong evidence that interest rates display mean-reversion.

Further, the random walk assumption, when combined with the assumption of a fixed real interest rate, leads to an assumption that the entire variation in nominal bonds is driven by inflation. This relationship is surely a clear mis-specification of the inflation process. It is extremely difficult for outside observers of PIMS, however, to assess whether it is quantitatively important for PIMS to more accurately measure the inflation process. One could easily imagine estimating the entire PIMS model on a real basis by inflation-adjusting all of the model parameters on the front end. It is possible, however, that inflation does matter in the model, as would be the case if some parameters of the PIMS model (for example, parameters imposed by regulation) are fixed in nominal terms (and therefore vary in real terms). If so, then model accuracy would be improved by modeling inflation directly, rather than backing it out from an arbitrary assumption about the real interest rate.

There are many ways of modeling inflation and inflation expectations directly, including using data on inflation swaps, surveys of professional forecasters, and the spread between nominal Treasuries and TIPS yields. Although each approach has relative advantages and disadvantages, any one of them would provide a more meaningful inflation process than attributing the entirety of nominal bond yield fluctuations to changes in consumer prices. However, we stress that the importance of modeling time-variation in real interest rates may be even more critical than modeling time-variation in inflation.

Corporate Bond Yields

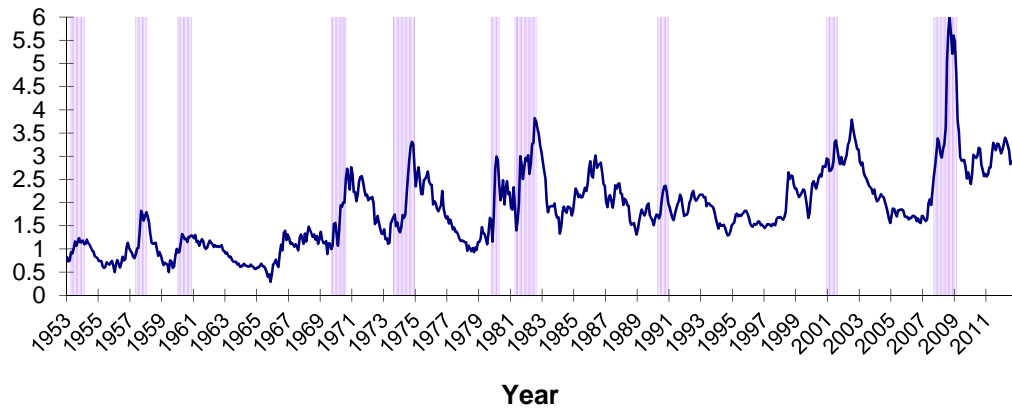
The PIMS model assumes that corporate bond yields are a linear function of the yield on Treasuries of the same maturity. This approach ignores the view in the

finance literature that the credit spread between corporate bonds and Treasuries varies with the state of the economy. For example, Fama and French (1993) and Bhamra et al. (2010) are two examples of papers providing empirical evidence that there are common risk factors in the returns to stock and bonds.

One can see graphically that the credit spread is typically larger during recessions (and, indeed, increased dramatically during the 2008-2009 financial crisis). In Figure 3 below, one can see that the credit spread between BAA corporate yields and Treasuries is typically larger during recessions (as indicated by the purple shaded areas).

Figure 3: Corporate Yield Spread
(Purple areas indicate NBER recessions)

BAA - Treasury Bond Spread



If PIMS moves in the direction of incorporating time-varying bankruptcy probabilities and simulating firm pension portfolios based on actual holdings (both of which are discussed below), then allowing yield spreads to vary with the underlying state of the economy would do a better job of capturing the dynamic aspects of plan funding.

Stock Returns

In PIMS, it appears that equities are modeled using a mean-reverting process, in which the mean is based on a long-run average of historical returns (since 1926).

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

where s_t denotes stock price at time t , s_{t-1} denotes the price at time $t-1$, and $\varepsilon_{s,t}$ is an innovation term. The terms α_s and β_s are parameters.

Our reading of the literature, however, suggests that although there is some evidence of *positive* momentum at a monthly level (for example, Jegadeesh and Titman 1993), there is no evidence that the *annual* stock market return is negatively related to the prior year's realization of stock market return. Given this evidence, a more natural starting point for modeling stocks is to maintain the random walk hypothesis. According to this hypothesis, stock prices move randomly, and consequently, stock returns are not predictable.⁹ Using PIMS notation, this would reduce the equation to:

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

Although there is no evidence of mean reversion at annual time horizons, there is a vast literature in finance exploring the predictability of aggregate stock returns over the long-run. Campbell (2000) discusses some of the “econometric pitfalls” that are relevant for evaluating this issue. A key one is that as the forecast horizon lengthens, the number of non-overlapping observations decreases rapidly. With a limited historical sample upon which to draw, there are fundamental limits to the efficacy of statistical controls to account for this. Despite these pitfalls, Campbell (2000, p. 1523) concludes that “the evidence for predictability survives at reasonable if not overwhelming levels of statistical significance. Most financial economists appear to have accepted that aggregate returns do contain an important predictable component.”

Whether to account for some form of long-run mean reversion is relevant, because mean-reversion serves to dampen the dispersion of long-run returns, which is obviously relevant for estimating PBGC's exposure to extreme events. Although the random walk assumption is a reasonable starting point, if PBGC wants to take advantage of long-run equity return predictability, then the predictive equation might relate the log stock return (or the excess log stock return) to the price-to-earnings ratio.

Another possible concern about PIMS' treatment of stock returns is the treatment of idiosyncratic risk. Specifically, the $\varepsilon_{s,t}$ term – which captures the annual, unpredictable variation in stock returns – is assumed to be “independent and identically distributed” (i.i.d.) over time. Conversations with PBGC staff indicate that the distribution of stock returns is assumed to be log normal (that is, $\varepsilon_{s,t}$ is drawn from a normal distribution), an assumption that represents another avenue through which the PIMS model may be underestimating the probability of extreme events. Research dating back at least 50 years (Fama 1963; Fama 1965) has found that stock returns do not follow a normal distribution. Rather, the distribution of returns has “fat tails,” that is, the probability of very large or very small return

⁹ Although not clear from the model's documentation, conversations with PIMS staff suggests that – in practice – the beta parameter is set equal to zero, which thus reduces the equation to the random walk equation noted.

shocks (very good or very bad returns) is much greater than that implied by a standard log normal distribution.

Further increasing the concern that PIMS may be underestimating the probability of “bad tail events,” research by Chang et al. (2013) indicates that the return on the market portfolio of stocks is negatively skewed. In other words, extreme declines in the stock market are more likely than extreme increases. In Figure 1 of Chang et al.’s paper, they graphically show evidence of changing volatility, negative skewness and fat tails (technically known as kurtosis).

Given the important role that equity returns play in determining plan funding status within the PIMS model, future modeling exercises might explore the robustness of the models to the assumptions that may bias downward the probability of extreme events. This can be done by separately and jointly re-estimating the model by setting the mean-reversion parameter to zero and allowing the innovation term to be drawn from a fatter tail distribution.

A “Systems” Approach to Estimation

Each of the above equations – those for treasuries, corporate bond yields, and stocks – would ideally be *jointly* drawn from historical estimates of the appropriate variances and covariances. It appears that PIMS uses covariances of stock returns and bond yields from a relatively short time period (1973 – 2007), although it is unclear from the documentation why this decision was made.

To further complicate matters, a concern in the literature is that the covariances may not be constant. For example, Pollet and Wilson (2010) discuss how an increase in systematic risk also increases the covariance between the stock market and the bond market. This concern seems an especially important to PBGC for the reasons outlined at the start of this section: if asset correlations rise at times of greater systemic risk – which is also likely to correspond with a greater probability of financial distress by plan sponsors – then this means that PBGC has much greater tail risk exposure than PIMS currently captures. For example, during the 2007-2009 financial crisis, stock returns declined dramatically while nominal interest rates declined sharply. Consequently, plan asset values plunged at the same time that the present value of plan liabilities increased.

PBGC may wish to explore the use of Vector Autoregression (VAR) models to estimate the joint behavior of log stock returns, changes in log of US Treasury yields, credit spreads, GDP growth and inflation. Notably, this approach “nests” the random walk hypothesis. In other words, it is possible to set coefficients to zero for specific elements and estimate a restricted model then test whether the restricted model performs better than the unrestricted model. Campbell, Lo and MacKinley (1997) discuss this approach in sufficient detail to guide future work in this area.

Plan Portfolio Allocations

When a plan sponsor experiences bankruptcy (both in the real world and in the PIMS model), the size of the liability inherited by PBGC depends on the funding status of the plan. Funding, in turn, depends on, among other things, the rates of return earned by the pension fund assets. Above, we discussed the PIMS process for projecting equity and bond market returns. How that translates into firm level funding, however, also depends substantially on the asset allocation decisions made by the plans.

Currently, the PIMS model assumes that all plan sponsors invest pension funds in a portfolio that is 60% in stocks and 40% in bonds. This approach is simplistic and can be misleading on several levels. First, the portfolios of the typical pension fund today allocate funds across a much wider range of assets, including international stocks and bonds, real estate, commodities, agriculture, private equity and hedge funds, among others. Each of these asset classes have very different risk and return profiles, and the variation in allocations across these classes can lead to very different risk / return profiles for the insured firms.

Second, although some of this variation in portfolio holdings will be random, much of it is not. Indeed, we know that asset allocations are correlated with other plan sponsor and fund characteristics, a fact that is especially important to PBGC. For example, Adonov, Bauer, and Cremers (2013) show that pension funds allocate more to riskier asset classes when they have a smaller fraction of pension plan participants still actively working, when the fund is larger, and when the defined benefit plan offers some form of inflation protection. The fact that larger funds with more generous benefits hold a higher share in risky assets is an especially important factor to consider when modeling future PBGC liabilities.

We suspect, on net, that the 60/40 portfolio assumption biases the PIMS model against finding the fat tails that are likely to exist in the real world. Incorporating dispersion in asset allocations will lead to great dispersion in funding levels. This dispersion will likely be further extended by the factors mentioned above related to greater risk taking by firms with more generous benefits. However, without a robust system in place to test alternative assumptions in the PIMS model, we cannot be certain of the direction or the magnitude of any impact of changing this assumption on the distribution of outcomes.

We also note that the shortcomings of the 60/40 approach have been highlighted in prior reviews (for instance, Buck Consultants) and, even more importantly, have been acknowledged by PBGC staff. Indeed, PBGC staff has indicated that they are currently working on a plan to use the detailed asset allocation data available in the Form 5500 data to do additional work in this area.

Bankruptcy Frequencies and Severities

PBGC insures against the joint occurrence of two events: (i) a plan sponsor declares bankruptcy and (ii) the plan sponsor has an under-funded pension plan. As with any other insurance company or financial institution, an assessment of the sustainability of PBGC's finances depends on a reasonable estimation of future claims. Given that PBGC liabilities are triggered by firm bankruptcy, the modeling of the bankruptcy process is one of the most important inputs into any model of PBGC claims.

As noted elsewhere, PBGC's PIMS-SE model was ahead of its time when it was first developed in the early 1990s. It uses a dynamic modeling technique that has only become common in the bankruptcy literature in economics and finance more recently. PBGC chooses a sample of firms with initial financial and pension plan characteristics intended to represent actual firms covered by PBGC. These sample firms are then modeled to evolve through simulation based on a set of equations representing the underlying economic and financial processes.

In each period of the model, once firm characteristics have been updated based on the simulations, a bankruptcy probability is calculated for each firm based on a logit model calibrated using historical data. Firms that are deemed in the model to experience bankruptcy are then evaluated based on their funding status, and sufficiently underfunded liabilities are then added to PBGC's simulated balance sheet. This is a reasonably sophisticated approach to modeling, particularly relative to the state of modeling at the time it was developed.

Despite this strong foundation, the bankruptcy modeling process in PIMS exhibits two features that could materially affect the outcomes of the simulations. First, there is a range of data concerns. For example, the process for updating parameter estimates based on new data has been inconsistent and inadequate. Conversations with PBGC staff confirm that the model has not been re-estimated in many years and thus, in all likelihood, the model parameters likely differ from today's best available estimates. There are also potential concerns about the treatment of missing data, a subject to which we will return below.

The second important feature of the model is that it does not allow for macroeconomic factors to lead directly to correlated defaults among firms. The bankruptcy model relies primarily – although not exclusively - on firm-specific accounting measures, and thus does not model the potentially important linkages with overall economic activity that market-based measures may be more apt to capture.

Given the importance and the complexity of the bankruptcy modeling process, and the substantial developments in the field since this model was created, PBGC may find it particularly valuable to invest additional resources in evaluation and validation of this model beyond what this report is able to do.

Data Concerns

There are two data concerns with the current incarnation of PIMS, both of which we believe can be addressed relatively easily.

First, the model's predictive power would be enhanced if PBGC would regularly update the data set used to estimate the bankruptcy equation. The number of public firms that have filed for bankruptcy in the U.S. is relatively small as a fraction of the total universe of firms. Of those that do file, many of them are small enough or well-enough funded that they do not represent a substantial part of PBGC's balance sheet. Rather, PBGC's balance sheet is dominated by a very small number of extremely large bankruptcies: As of 2011, pension plans from only 10 firms accounted for over 60% of PBGC liabilities. Given that PBGC is especially sensitive to "tail events," accurately modeling these extreme outcomes is important for assessing the insurance program's financial status.

From a statistical perspective, the very low ratio of bankrupt-to-total firms presents econometric challenges. With so few firms coded as a "1" (indicating bankruptcy) in a model with a binary dependent variable, even small changes in the composition of firms that have entered bankruptcy can have a significant impact on parameter estimation. Additionally, econometric models – including the logit model used by PBGC – impose distributional assumptions on the error term. Although these distributional assumptions often make little difference in practice, they are much more likely to matter when there are so few observations taking on one of the two possible values of the dependent variable.

It is somewhat surprising that PBGC does not place a higher priority on regularly updating the parameters in the bankruptcy model as new data becomes available. Conversations with PBGC staff confirm that the data used in this model is "extremely" outdated. Although we are sympathetic to the significant resource constraints under which PBGC is operating, the amount of effort required to annually update the data file and re-run a logit model is very small relative to the importance of the bankruptcy model to the PIMS model.

The second data concern is that "missing data" can be especially acute when studying bankruptcy because firms often do not report data during or immediately preceding bankruptcy. It is unclear from the PIMS documentation how, if at all, PBGC addresses this problem. Procedures along the lines of Shumway (2001), Chava and Jarrow (2004), or Campbell et al. (2008), which substitute averages based on existing data for missing variables, is one potentially useful approach to address this problem. If this is already done, then it would be helpful to consumers of the model to be more explicit about this in the documentation.

Structural Form of the Bankruptcy Model

PBGC uses what might be called a “dynamic” logit model to estimate bankruptcy probabilities. This choice placed the model somewhat ahead of its time by academic standards in the early 1990s. In the 1980s, logit models of bankruptcy were static in the sense that each firm was assigned a predicted probability of bankruptcy, and this was assumed to be constant over time (or, if multiple observations of a firm were available, they were treated as separate firms) (See, for example, Ohlson 1980; Mensah 1984).

In contrast, dynamic logit models place structure on the errors to recognize the information contained in a series of observations of the same firm. Although PBGC has used this approach since the early 1990s, it was not until after Shumway’s (2001) contribution – in which he showed, among other things, that static logit models are biased – that this type of model became broadly used in economics and finance. Shumway showed that a discrete-time hazard model with time-dependent covariates is computationally equivalent to a dynamic logit model with a specified error structure. Shumway recognized PBGC explicitly when he noted “estimating hazard models with a logit program is so simple and intuitive that it has been done by academics and regulators without a hazard model justification ... The Pension Benefit Guaranty Corporation forecasts bankruptcies by estimating a logit model by firm year.”

Some studies have recognized that using a Cox proportional hazard model rather than dynamic logit allows the model to capture macroeconomic trends, without having to split the sample or arbitrarily assume that the effects of the macroeconomic variables are linear within the exponent of the logit. For example, Hillegeist, et al. (2004) use the number of bankruptcies within the past year (relative to the number of firms in the sample) as the baseline hazard, and show that this increases the predictive power of the model. Nam et al. (2008) use change in interest rates and exchange rate volatility to model the baseline hazard and find that these are important for modeling bankruptcy in Korean data. These papers suggest a way that PBGC may be able to capture macroeconomic trends and/or bankruptcy contagion effects – by using macroeconomic variables to establish the baseline hazard rate. This is potentially important for modeling “waves” of bankruptcy. We will return to a discussion of macroeconomic factors below.

Although this review is focused on possible changes that can easily be implemented in the existing PIMS infrastructure, we should acknowledge that there exists a range of other more flexible approaches. For example, Cheng et al. (2010) use semi-parametric estimation techniques to overcome the restrictive linearity assumption between predictive variables and default probabilities. Using local likelihood methods to specify the baseline hazard function, they find their model has higher predictive accuracy than other models. However, this method is likely too computationally burdensome for PIMS, can be difficult to interpret, and is not yet widely used in the economics and finance literatures. These reasons also limit the

attractiveness of using non-parametric techniques, which are more widely used in other literatures.

There has also emerged a small literature focusing on multinomial logit models which allow for more than two outcomes: for example, they allow modeling of financial distress through forms other than bankruptcy. These analyses (for example, Lau 1987, Astebro and Winter 2012, and Jones and Hensher 2004) find that there can be advantages to this type of model. Because this literature lacks the depth of the existing literature on binary bankruptcy outcomes, PBGC may wish to simply monitor this literature going forward to assess whether such an approach would have a meaningful impact on PIMS.

As an alternative to the reduced form approach used by PIMS, there is also a literature focused on structural modeling of bankruptcy. The most common structural model is based on Black-Scholes-Merton (BSM) option pricing model, which uses techniques from continuous time finance. This approach recognizes that a firm near distress has a payoff structure similar to that of a call option on the value of the firm, with a strike price equal to the value of the firm's debt. One can then use BSM to value this option-like firm. Moody's KMV proposes a method for handling the fact that neither the firm's true value or the volatility of its assets are known, calculating a "distance-to-default" (DD) measure (see Bharath and Shumway 2008 for a description of this approach). Empirical DD measures have been widely used in finance, a topic to which we will return below.

In terms of practical implementation, Duffie et al. (2007) and Duffie et al. (2009) provide the best references for using structural models that incorporate parameters based on historical data using a statistical bankruptcy equation. The data generating processes for the macroeconomic variables are specified in a manner that is reasonably close to methods used by PIMS in other parts of the model, suggesting that such an approach may be feasible to implement by PBGC. Of course, a more structural PIMS model would also have to incorporate a dynamic process for plan underfunding in addition to firm asset values. We recognize – as noted in PIMS documentation – that current funding rules lead to levels of volatility in minimum and maximum funding contributions that may render options-based models intractable. We are unsure whether this is actually the case, however, given that most of these rules are deterministic within PIMS. Ultimately, however, we simply have too little information to assess the costs versus the benefits of moving toward a structural model.

Choice of Covariates in PIMS Bankruptcy Model

The parameters used in PBGC's bankruptcy model appear to be based on the empirical models that were common in the early 1990s when the PIMS model was developed. These include proxies for liquidity, leverage, firm size, and industry controls.

Specifically, PIMS proxies for liquidity using two lags of cash flow minus pension contributions, divided by total assets. It proxies for leverage using:

- Two lags of the log of equity-to-debt ratio
- Lag of log of firm net funding ratio across all sponsored DB plans
- An indicator of missing funding ratio

PIMS measures firm size using:

- Lagged log of employment
- Lag of change in log of employment

Finally, PIMS uses indicator variables for:

- Transportation, communications and utilities
- Financials and insurance

We also note that the PIMS model includes interactions of the equity-to-debt ratios with the firm size proxies.

Like the PIMS model, most bankruptcy studies also include controls for liquidity, leverage, and firm size, although the specific proxies vary. For example, in contrast to the above choices, Campbell et al. (2008) measure liquidity using the ratio of cash and short term investments to the market value of total assets, whereas Shumway (2001) uses the ratio of working capital to the book value of total assets. For leverage, Campbell et al. use total liabilities as a fraction of the market value of total assets whereas Shumway uses retained earnings-to-book value of assets as well as the ratio of market equity-to-book value of liabilities. For firm size, Campbell et al. use the log of the firm's market equity relative to the total value of the S&P 500, whereas Shumway uses the log of the ratio of the firm's market value of equity to the total valuation of the New York Stock Exchange and American Stock Exchange.

Table 1: Parameter estimates from major logit bankruptcy studies and PIMS

Results from Campbell et al. (2008) are reproduced from Table 3, Logit Regression of Bankruptcy/Failure Indicator on Predictor Variables, Bankruptcy Model 2. The first set of results from Shumway (2001) are reproduced from Table 2, Forecasting Bankruptcy with Altman's Variables, Panel B: Hazard Model Estimates. The second set of results from Shumway (2001) are reproduced from Table 6, Forecasting Bankruptcy with Market-Driven Variables, Panel B: Market and Accounting Variables. Results from Chava and Jarrow (2004) are reproduced from Table 5, Panel A: Estimation using 1962-99 data for NYSE-AMEX-NASDAQ companies including financials (column 2, Public firm model). Results from Ohlson (1980) are reproduced from Table 4: Prediction results, Model 1. Finally, results from PIMS 1 are reproduced from Table 6-4 and PIMS 2 are reproduced from Table 6-6. Bold numbers indicate significance at the 5% level or below, except for PIMS 2 which did not provide t-statistics or p-values. Cash flow variables from PIMS 2 are assumed to be net of pension contributions, as in PIMS 1, although this was not specified. Dates for PIMS 1 and 2 are best estimates based on descriptive evidence.

Variable definition	Name of variable	Campbell et al. 1963-1998	Shumway 1 1962-1992	Shumway 2 1962-1992	Chava & Jarrow 1962-1999	Ohlson 1970-1976	PIMS 1 1980-1997	PIMS 2 1980-2009
Intercept	Intercept or Constant	-7.65	-3.226	-13.303	-14.886	-1.32	-4.725	4.662
<i>Profitability variables</i>								
Net income/Book value of total assets	NITA			-1.982	-1.9236	-2.37		
Net income/Market value of total assets (with lags, geometrically weighted)	NIMTAAVG	-32.52						
EBIT/Total Assets	EBIT/TA		-8.946					
Sales/Total assets	S/TA		0.158					
First difference in net income/Sum of abs. value of net incomes in numerator	CHIN					-0.521		
<i>Liquidity Variables</i>								
Stock of cash and short term investments/market value of total assets	CASHMTA	-4.89						
Working capital/Book value of total assets	WC/TA		-0.732			-1.43		
Current liabilities/Current assets	CLCA					0.757		
Funding provided by operations/Total Liabilities	FUTL					-1.83		
One if net income was negative for last 2 years	INTWO					0.285		
Cash flow net of pension contributions/Total assets (1 lag)**	$(CF_{t-1} - Cont_{t-1})/A_{t-1}$						-3.7866	3.8125
Cash flow net of pension contributions/Total assets (2 lags)**	$(CF_{t-2} - Cont_{t-2})/A_{t-2}$						-1.4679	2.0676
<i>Leverage Variables</i>								
Total liabilities/Book value of total assets	TLTA			3.593	4.0338	6.03		
Total liabilities/Market value of total assets	TLMTA	4.32						
Retained earnings/Book value of total assets	RE/TA		-0.818					
One if total liabilities exceeds total assets	OENEG					-1.72		
Market equity/Book value of total liabilities	ME/TL		-1.712					
Log Market equity/Book value of total debt (1 lag)	$\ln(ED_{t-1})$						-0.8336	0.9197
Log Market equity/Book value of total debt (2 lags)	$\ln(ED_{t-2})$						0.781	-0.2406
Log Pension funding ratio	$\ln(FundingRatio_{t-1})$						-0.0902	0.1302
<i>Size variables</i>								
Log of firm's market equity/Total valuation of S&P 500	RSIZE	0.246						
Log of firm's market equity/Total valuation of NYSE and AMEX	Relative Size or RSIZ			-0.467	-0.3475			
Size	SIZE					-0.407		
Log Employment (1 lag)	$\ln(N_t - 1)$						-0.2925	0.2503
Log Change in Employment (1 lag)	$\Delta \ln(N_{t-1})$						-0.7154	0.6676

Table 1, continued: Parameter estimates from major logit bankruptcy studies and PIMS

Variable definition	Name of variable	Campbell et al. 1963-1998	Shumway 1 1962-1992	Shumway 2 1962-1992	Chava & Jarrow 1962-1999	Ohlson 1970-1976	PIMS 1 1980-1997	PIMS 2 1980-2009
<i>Market variables</i>								
Log of gross excess return/Value-weighted S&P 500 total return (with lags, geometrically weighted)	EXRETAVG	-0.51						
Gross excess return/Value weighted NYSE and AMEX return	EXRET (2)			-1.809	-2.662			
Square root of the sum of squared stock returns over 3-month period (annualized)	SIGMA	0.92			0.8312			
Idiosyncratic standard deviation of stock returns	SIGMA2			5.791				
Market to book ratio	MB	0.099						
Log price per share	PRICE	-0.882						
<i>Other firm variables</i>								
Log age	Ln(age)		0.015					
Missing pension funding ratio	Missing FR_{t-1} dummy						-0.0275	
Manufacturing and minerals	Ind2				-0.4597			
Transportation, communications, and utilities	Ind3 or UT				-0.0178		-1.642	1.6721
Financials and insurance	Ind4 or F				-1.226		-3.0279	2.7136
Manufacturing and minerals (interaction with NITA)	NITA*IND2				0.3414			
Transportation, communications, and utilities (interaction with NITA)	NITA*IND3				-2.5921			
Financials and insurance (interaction with NITA)	NITA*IND4				-3.4877			
Manufacturing and minerals (interaction with TLTA)	TLTA*IND2				0.3547			
Transportation, communications, and utilities (interaction with TLTA)	TLTA*IND3				-0.3423			
Financials and insurance (interaction with TLTA)	TLTA*IND4				-0.2175			
Interaction of ED and N	$\ln(ED_{t-1})\ln(N_{t-1})$						-0.0485	0.0373

In Table 1, we report the coefficients of the variables used in these models and other models. The coefficient estimates themselves are of little comparative value because the distribution of the underlying covariates is quite different. However, we view this table as valuable primarily for showing the diversity in approaches used.

What is perhaps more instructive than a comparison of the coefficients is an observation about measures that are excluded from PIMS, but included in most other models. First, all of the other studies reported in Table 1 include measures of firm profitability. These measures include the ratio of net income to book value of total assets (used in three studies), net income divided by market value of assets (used in Campbell et al. 2008), the ratio of EBIT-to-total assets (Shumway 2001), the ratio of sales-to-total assets (Shumway), and related measures. Importantly, however, the PIMS model does include two lags of cash flows (minus pension contributions, divided by assets). Because cash flow and profitability are highly correlated, this may mean that the exclusion of other profitability measures may not be quantitatively important. It is, however, difficult to know this for sure without re-running the PIMS regressions.

More importantly, over the two decades since the PIMS-SE model was created, studies have increasingly incorporated market-based (rather than accounting-based) variables as predictors of bankruptcy. As noted in Table 1, this includes measures related to excess returns, market-to-book ratios, and the like. In the academic literature, market-based measures are often used to test various theories: according to the efficient markets hypothesis, market measures of distress should completely subsume the predictive power of contemporaneous accounting variables because market prices should already reflect the information contained in financial statements and also reflect additional information available to investors about the firm.

Of course, the point of the bankruptcy logits in PIMS is not to test market efficiency, but rather to parameterize the bankruptcy equations in the PIMS model. In this light, there does seem to be broad (although not universal) agreement that market-based measures contain more information than accounting measures alone. There is, however, less agreement on exactly what combination of market-based measures should be used.

One commonly-used variable – already alluded to above - is the Black-Scholes-Merton “Distance-to-Default” (DD) measure. As one of the few measures with a clear theoretical motivation, it is intended to reflect the value of the volatility of firm assets relative to liabilities. Unfortunately, the empirical success of this approach is mixed. For example, Foreman (2003) argues that accounting variables explain most of a firm’s bankruptcy probability (although this paper is limited to the telecommunications industry). In contrast, Hillegeist et al. (2004) show that an empirical measure of DD does a better job at predicting default than the accounting variables often used in prior studies (which would include those used in PBGC’s model). Beaver et al. (2005) use both accounting and market variables, and show that although accounting variables remain significant in models that also include market variables, the market variables have been increasing in explanatory power relative to accounting variables over time. More recently, Campbell et al. (2008) show that the empirical DD variable does not add significantly to the model’s explanatory power when other *market* variables are already present.

One potential concern about using market based measures in PIMS is that the PIMS model is applied to simulated data. This change would require a process for allowing the market-based measures to evolve in the simulation. If the process by which these are set to evolve is simply a function of the accounting data already used in the model, then the model will not be able to capture the most salient aspect of the market information. Of course, completely excluding market-based measures that have been shown to predict bankruptcy means that the existing PIMS model suffers from omitted variable bias.

Some scholars have suggested that bankruptcy models ought to include covariates that directly influence the difference between market-based and accounting measures, and this is an approach that PBGC may wish to explore. For example,

Grice and Dugan (2001) note that unions and lawsuits affect a firm's likelihood of filing for bankruptcy, yet these factors are rarely included in bankruptcy models (including PBGC's). Deterioration of worker morale or reduced work effort by employees may also contribute, although this is extremely difficult to quantify. It may be in PBGC's interest to support independent research that would move this literature forward. Differences in firm governance are another example that might be important determinants of the difference between market and accounting based measures.

Absent substantial empirical work that is beyond the scope of this paper, it is difficult to provide a sign or magnitude for any bias that may result from the exclusion of market-based information. However, given the central role of bankruptcy in triggering PBGC liabilities, studies of the robustness of PIMS modeling to alternative bankruptcy models would seem to be a high priority for assessing the overall robustness of the PIMS model.

More generally, PBGC may wish to undertake a systematic process for model validation. A commonly used – and relatively simple to implement – approach would be to divide up realized time periods into an earlier and a later “holdout” period. The bankruptcy model can be estimated on the earlier data, and then these coefficients applied to a simulation that represents the same length of time as the later period. One can then compare the simulated results to the actual realizations from the later holdout period. Indeed, given how long it has been since the bankruptcy model has been updated, PBGC could apply this validation methodology to the period that follows the estimation sample currently used.

Time-Varying Macroeconomic Factors

The PIMS-SE model does not allow for the bankruptcy parameters to vary over time or with the state of the economy. Rather, macroeconomic factors can only affect bankruptcy probabilities *indirectly* in the PIMS model via the impact of the macroeconomy on the value of the firm level characteristics that are included in the bankruptcy model.

In contrast, virtually every bankruptcy forecasting study that has considered the impact of the macroeconomy has found that it plays an important role in firms' distress risk. For example, Figlewski et al. (2012) investigates various macroeconomic variables that might cause firm bankruptcy for different magnitudes of credit downgrades. Although his model largely omits internal firm covariates, he finds that macroeconomic conditions have a stronger influence on downgrades from speculative classes to default than they have on downgrades from higher ratings to default. Hol (2007) tests a model that includes both macro and firm variables on a sample of Norwegian firms and finds that the GDP gap is a significant influence on default probabilities. Duffie et al. (2007) include 3-month Treasury bill rates and the trailing 1-year S&P 500 return as covariates in a regression of bankruptcy. This paper may be of particular interest to PBGC because, like PIMS, it

uses these coefficients as parameter estimates to predict default data on model-generated data.

Another way to frame this issue is to note that the relationship between firm-specific variables and bankruptcy may be time-varying (and may depend on the state of the economy). Statistically, models with time-varying coefficients may be too complex to be incorporated as a feature of the PIMS model. A relatively simple way to circumvent this problem is to divide the sample based on factors believed to drive time variation in the coefficients. For example, business cycles might be one natural source of non-stationarity in a bankruptcy model (Mensah 1984), as might different inflation or interest rate regimes. Further, there are reasons to think that PBGC might want to allow industry effects to operate differently in different economic environments. For example, industries with a high concentration of defined benefit plans may also be industries with high operating leverage (for example, wages fixed by long-term union contracts) and high financial leverage (due to being old-economy, large tangible-asset firms) that may do extremely poorly in downturns, whereas other industries may behave quite differently.

There is evidence that period and industry effects matter: Grice and Dugan (2001) show that the predictive accuracy of the Ohlson (1980) and Zmijewski (1984) models vary across economic regimes, as well as across industrial versus non-industrial firms. Anyane-Ntwo (2011) also provides evidence that manufacturing and service firms respond differentially to macroeconomic factors. Another possible approach would be to include some measure of firm sensitivity to macroeconomic shocks (such as an asset beta) into the model.

The actual history of PBGC's loss experience appears to be statistically close to impossible in models that do not allow for time-varying factors. Thus, it would seem that an exploration of alternative ways of incorporating macro factors in a tractable way is of first order importance to improving the PIMS model. In all likelihood, allowing for time-varying bankruptcy processes will lead to more correlated firm defaults. Understanding the extent to which there are correlated episodes – or “waves” – of bankruptcy is quite important to understanding the long-term viability of any credit insurance program, and PBGC is no exception.

Finally, we note that there is evidence that the costs of financial distress, in addition to the probability of it occurring, are also correlated with the macroeconomy. Specifically, bankruptcy is more costly when the economy is doing poorly (Almeida and Philippon, 2007; Chen, forthcoming). To the extent that these costs affect the firm's funding behavior (which we discuss in the next section) and/or PBGC's ability to partially recover value through the bankruptcy process, this serves as another source of temporal concentration of risk.

Firm Behavior In or Near Bankruptcy and Implications for Modeling

There is a lively debate among bankruptcy experts about the extent to which bankruptcy is a strategic choice on the part of a firm. For example, Delaney (1999) states: “bankruptcy no longer means ‘being broke’ ... Bankruptcy is sometimes an offensive weapon used by the rich and powerful rather than a refuge for the weak and ailing.”

Although his evidence is only anecdotal, there have indeed been cases of financially healthy companies that have chosen to undergo Chapter 11 reorganizations in order to protect themselves from future liabilities unrelated to their current financial position, such as potential liabilities resulting from asbestos litigation (Tweedale and Warren 2004). Lemmon et al. (2009) discuss a benefit of Chapter 11 as being the ability to exercise the option inherent in lease contracts.

In her provocatively titled book *Pension Dumping*, Hawthorne (2008) provides anecdotes of firms that filed for bankruptcy with the purpose of shedding legacy costs such as pension plans. Although one might think that this could or should result in fraudulent conveyance suits, Orr (1998) finds that these cases are very difficult to win. Ippolito and James (1992) also discuss strategic pension terminations.

Of course, strategic behavior is not only limited to the decision of whether to declare bankruptcy: it can also affect other relevant behaviors that affect the size of the liabilities inherited by PBGC. A number of studies, for example, have documented an increase in pension underfunding during financial distress. Duan et al. (2012) find that firms significantly lower contributions prior to default. Bean and Bernardi (2000) argue that pension underfunding is a strategic choice by management, and a transfer of risk from equity-holders to workers. An et al. (2013) find that the risk management strategies of defined benefit pension funds can be characterized by risk-shifting during periods of distress.

Bergstresser et al. (2006) show that managers change pension assumptions to manipulate earnings. Although their paper was focused on managerial incentives rather than financial distress, one could see how this could adversely affect PBGC on two margins: bankruptcy risk would be higher than estimated because earnings are artificially inflated, and pension underfunding would be greater than estimated because pension returns are inflated. Rauh (2009), however, finds that the incentive to avoid distress may be stronger than the incentive to shift risks.

Why does any of this matter? To the extent that bankruptcy is the outcome of a strategic choice, rather than a (conditionally) random outcome, it becomes much more difficult to model in a simple reduced form way based on historical data. Absent observable characteristics that are correlated with the strategic behavior, simple statistical models may provide misleading estimates (that are likely biased downward) for “strategic defaulters.”

There is no simple way to deal with this possibility. To the extent that PBGC has evidence that strategic behavior is associated with certain situations in large firms – such as union disputes, costly and unforeseen litigation, etc. – that might be more likely to lead to strategic bankruptcy behavior, PBGC may actually be better off taking these firms out of the PIMS model and assessing the risks of this small number of large firms more qualitatively.

Although this approach may seem incongruous with a simulation model, it is, in spirit, not all that different from the Early Warning program. Separating the firms into quantitatively and qualitatively modeled sub-segments is also sensible in light of the political nature of some plan terminations. For example, during the recent financial crisis, the U.S. government took the highly unusual step of directly involving itself in the restructuring of the automobile industry, a factor that undoubtedly affected the liability exposure of PBGC.

The possibility that larger firms are more likely to behave strategically also has implications for PBGC's sampling approach. PBGC documentation notes that “the sample currently available for PIMS simulations represents only about 4% of PBGC-insured plans with greater than 100 participants; these represent almost half of all insured plans' liabilities and underfunding. The weighting scheme effectively creates ‘partner firms’ for each plan sponsor in the PIMS sample.”

In essence, this choice means that PIMS is basing its stochastic bankruptcy risk model on plan sponsors for which plan termination is most likely to be strategic rather than stochastic in nature. Although we certainly understand the gain in computational speed that drove the decision to use a “partner firm” approach, the tremendous gains in computing speed that have occurred over the past two decades should make it much more feasible to reflect a more complete distribution of the PBGC's true universe of firms.

Other Modeling Choices Affecting Claims on PBGC

Issues Regarding the Composition of the Universe of Covered Plans

PIMS essentially assumes that the starting universe of pension plans covered by PBGC remains constant, except as a result of insolvencies. Thus, it does not include the effect of voluntary plan terminations nor of new freezes of benefit accruals within existing plans. In reality, there will be some voluntary plan terminations and potentially a quite significant number of new freezes.

For example, from 2008 to 2012 the proportion of PBGC-insured plans that had undergone a partial or a complete accrual freeze increased from 24.3% to 35%.

Most of this increase came in the form of hard freezes, which increased by over 2000 during this period. Freezes disproportionately occurred in smaller plans, so that only 29% of PBGC-insured participants in 2012 were affected by freezes, compared to 35% of PBGC-insured plans. It is worth noting that in 2003, only 2.5% of participants were enrolled in plans with a hard freeze.¹⁰

Ignoring voluntary plan terminations has the following direct effects:

PBGC expected premiums are overstated: A terminated plan ceases to be covered by PBGC and therefore no longer needs to pay premiums.

Potential claims may be overstated: Plans that are voluntarily terminated can no longer make a claim on PBGC. If a plan that should have been assumed to voluntarily terminate is instead assumed in a PIMS simulation run to present a claim on PBGC, then total claims will be overstated.

PBGC expenses are overstated: A modest portion of PBGC's expenses represents the marginal cost of dealing with the sponsor and participants in a healthy plan. The claims figure would also include an expense component for handling a distress termination.

It is difficult to quantify, but it appears likely that ignoring voluntary terminations to date has had little net effect on PBGC's projected financial condition. In principle, the effect could be either to overstate or understate the financial strength of PBGC. Voluntary terminations to date have almost all been of smaller plans for whom the administrative burdens appear to have been the major factor in prompting termination. (For example, around 99,000 plans entered into a standard termination between 1986 and 2004, with the average number of participants in these plans being about 18 times smaller than the average size of existing plans covered by PBGC currently).¹¹

There is a considerable disincentive for large plans to be voluntarily terminated, since it requires the sponsoring employer to pay a life insurer to take over the obligations. Insurers will generally charge significantly more than the cash contribution that would be required to bring a plan up to full funding. (Insurance regulation and prudent practice cause the investment portfolios of insurers to have quite low returns. In addition, insurers have to include an expense load, a risk load, and an expected profit in order for the transaction to make sense for them.) The price charged by an insurer creates not only a significant cash outlay for pension plan sponsors but also a hit to reported net income in at least the amount of the excess of the payment to the insurer over the contribution that would have been necessary to achieve full funding.

¹⁰ PBGC, (2013), "Single-Employer DB Plan Freezes."

¹¹ Data on terminations taken from: PBGC, (2005), "An Analysis of Frozen Defined Benefit Plans." Calculations are our own.

The same logic generally means that any large sponsor that is considering a voluntary termination likely either has a very well funded plan or is quite strong financially. Therefore, it presents little risk of making a claim on PBGC over the usual 10-year projection period. To the extent that voluntary terminations occur among sponsors that will not present a claim on PBGC, then PIMS' ignoring of future voluntary terminations causes a net overstatement of PBGC's financial position, since PIMS overstates premium revenues by more than it overstates the combination of expenses and the likely minimal volume of claims that are assumed from these plans.

Ignoring future plan freezes has similar types of direct effects, although the magnitudes differ.

PBGC premiums may be overstated over time: A plan that freezes its membership will see a decline over time in its participant count and therefore in its per capita PBGC premiums. Variable premiums may also be reduced by the freeze, since the plan effectively shrinks over time. However, this outcome will depend as well on contribution levels and any changes to investment strategy that may occur in conjunction with the decision to freeze the plan.

Claims are potentially overstated: Frozen plans shrink over time and therefore may have lower underfunding levels. Given enough time they will effectively be terminated as the number of participants declines to zero, although this would take many decades. Assuming no net effect from differences in contribution rates and investment strategies between frozen and active plans, claims on PBGC would be lower and even perhaps less frequent. (Small potential claims on PBGC are often ignored in bankruptcy in order to avoid the complications of PBGC involvement). In reality, firms may well change their strategy on contributions and investments in conjunction with a freeze, but it is not clear whether such changes would improve or harm PBGC's position on average.

PBGC expenses are overstated: A modest portion of PBGC's expenses is proportional to participant count and would be lower for frozen plans, over time.

Many more plans of substantial size initiate freezes rather than voluntary terminations, so this factor is potentially more important. Offsetting this factor, the effect of freezes is much more gradual than the effect of plan terminations. The impact during the typical 10-year projection period would not necessarily be large, particularly as there are offsetting factors in the overstatement of both costs to PBGC (claims and expenses) and revenue (premiums).

Overall, the question of how PIMS should deal with voluntarily terminations and freezes is complex. There are not good statistical models to project these activities, so PBGC would be required to exercise considerable subjective judgment. However, it would be quite helpful to users to have sensitivity analyses informed by historical

experience for different levels of terminations and freezes, so that policymakers could determine the extent to which PBGC's future financial condition will be affected by these activities. If those sensitivity analyses demonstrate a strong effect, then further consideration should be given to explicitly modeling voluntary terminations and/or new freezes. Such modeling would need to consider the relationship between factors driving such actions and other key variables in PIMS.

Contributions by plan sponsors

The level of contributions that plan sponsors make to their pension trusts is determined by two prime factors: (a) the legal requirements on minimum contributions and (b) the sponsor's willingness to contribute more than legally required or its inability or unwillingness to fund even the required amount. In general, PIMS assumes that employers make precisely the minimum legally required contributions, with the exception that once a firm is determined by the PIMS simulation to be in bankruptcy, there is an *ex post* assumption that the contribution for the previous 12 months was not in fact made. The exception in the model exists because it has been generally true in practice that the most recent contributions have been skipped as firms head into bankruptcy.

In practice, many employers do contribute more than the minimum that is legally required of them, often by substantial amounts. Sometimes these contributions are intended to raise funding levels sufficiently to reduce or eliminate the variable rate premium on underfunding. Other times, they stem from the overall cash and tax planning situation and objectives of the employer, including a sense of greater prudence, which of course is not directly observable to PBGC. Contributing earlier than required to a pension trust is analogous to an individual making more generous contributions to his or her 401(k) plan. In that case, with a few exceptions, the firm benefits from an immediate tax deduction and the benefit of earnings on the contribution that accrue free from corporate tax.

PIMS would be more accurate if it included an accurate estimate of the additional contributions that firms are likely to make above the legal minimums. Contributions above the legal minimum impact funding levels, and funding levels in turn affect ultimate claims on PBGC. In addition, as described below, the variable portion of the premiums collected by PBGC are determined by funding levels.

Although PIMS does make a rough adjustment to the premium calculations to reflect the role of voluntary contributions in reducing variable rate premiums, it has not modeled excess contributions explicitly. There appear to be two reasons for this omission. First, there seems to be an implicit view that either PBGC does not have a good way of predicting those voluntary contributions and/or that any such assumptions would be subject to a great deal of criticism since there is not a clear basis for choosing the parameters.

Second, and at least as important, PBGC explicitly argues that those firms that present claims on PBGC are unlikely to have made voluntary contributions for some years prior to bankruptcy and that history demonstrates this. Claims on PBGC have seldom come as a sudden shock, but rather have usually resulted from the demise of firms that have been heading towards insolvency for years and therefore have been acting under constraints that would lead them to conserve as much cash as possible, including by avoiding voluntary pension contributions above the minimum threshold. Nor would the existence of the variable rate premium generally deter a troubled firm from minimizing its contributions, since the variable premium is still far cheaper than the net cost of borrowing funds to make the contribution.¹²

Thus, inaccuracies in claims forecasts introduced by assuming only the minimum contributions are made would arise solely from those firms that eventually became insolvent but chose to make significant voluntary contributions over the course of the PIMS simulation prior to insolvency. This outcome would require either that a significant volume of claims came from employers who became insolvent because of a sudden shock or that, improbably, firms sliding into insolvency would choose to make voluntary pension contributions.

Overall, it seems likely that this omission does not create a major distortion in the estimates of claims on PBGC for the standard 10-year simulations. Longer simulations might produce a somewhat greater over-estimate of PBGC claims from this factor, although it is difficult to quantify this effect.

That said, explicitly modeling voluntary contributions might improve its claims projections and might also reduce the confusion among external users regarding the need to make a substantial adjustment to variable rate premium calculations, described below, to reflect voluntary funding. It would also be useful to do sensitivity analyses because PIMS could significantly overstate the loss from a major bankruptcy in the unusual case that a firm that has voluntarily funded more than the minimums goes bankrupt suddenly due to an external shock.

Pension liabilities

The level of pension liabilities assumed for each sponsor is a critical variable in determining PBGC's financial condition. For one thing, claims on PBGC are based on the difference between those liabilities and the related assets. Also, PBGC's variable rate premiums are determined based on the relationship between pension assets and pension liabilities. Further, the level of liabilities has an indirect effect on pension contributions and therefore on future funding levels and future PBGC premiums.

¹² Even if management of the firm were convinced that it would survive, and therefore took a longer view of their cash planning, the variable premium is likely to be considerably lower than the net cost of borrowing the contribution minus the investment return in the pension fund.

PIMS collects information about the starting level of pension liabilities at each plan from the Form 5500 filed by the plan, with some adjustments made for more recent developments. The Form 5500 information is somewhat dated, for a variety of reasons. Sponsors have a number of months to file after the end of the plan year, and there can be some further delays in putting the data into electronic form for PBGC. Depending on the timing of the PIMS run, as much as another 364 days might have passed since the most recent information became available. In total, most of the data from the Form 5500 appears to be about 14 months old at the time that PBGC runs PIMS to produce the figures it uses in its annual reporting, and some of it will be older.

PIMS is explicitly designed to project the effects of these time lags in order to use the best estimates for the condition of the plans at the start of the relevant simulation period. PBGC updates the assumed asset and liability values from the Form 5500 to reflect key features that are time-dependent. Asset values are assumed to have gone up or down at the same rate as broad indices of those categories of assets over the relevant period. Liabilities are assumed to have grown, or occasionally shrunk, based on historic patterns and the intervening movement in interest rates.

PIMS uses some crucial simplifications to model the liabilities, primarily out of a desire to hold down computational complexity.

First, PIMS is not currently able to directly model certain types of benefit plans, instead making a series of simplifying assumptions that PBGC itself recognizes as limitations of the model. For instance, career average pay benefit plans are treated in the PIMS model as final average pay plans.

Similarly, PIMS does not fully model the benefits of soft-frozen plans, although some progress has been made in this regard in recent years. PIMS also has difficulty modeling the benefits of cash balance plans, “hybrid” plans that incorporate aspects of defined contribution plans. While PBGC is currently working to improve the PIMS model regarding cash balance plans, they indicate that because these plans utilize many different benefit formulas, the same limitations in accounting for sub-groups will “prevent many converted cash balance plans from being adequately modeled.”

Second, PIMS uses a single mortality table for all plans; this will usually be different to the one used by the plans’ actuaries to calculate their liabilities on the Form 5500. For purposes of simplification, SE PIMS assumes that all plan participants are males.

In recognition of these limitations, the single employer version of PIMS uses a complex adjustment process to “true up” the pension liability estimates from its simplified approach with the more detailed analyses of the plan actuaries. In essence, the key actuarial information is fed into PIMS and an initial pension liability estimate is produced using the standard simplifying assumptions. The size of the assumed employee and retiree base is then adjusted up or down by whatever factor

is necessary to produce a starting liability consistent with the figures from the Form 5500, as initially adjusted for the passage of time since the filing of the form. Although these adjustments are intended to de-bias the results of the model, this approach does reduce the transparency of the model.

However, this approach could produce estimates of future changes in liabilities due to the passage of time that are inconsistent with the information from the plan actuaries. Thus, further adjustments are made to bring PIMS' assumed "normal cost," an estimate of the cost of the additional benefits earned from a year's service, in line with that shown by the plan actuaries for the recent past. Thus, the assumptions in PIMS about the plans are calibrated to provide estimates of the initial liability and normal cost that closely match the actual values given in the Form 5500, despite differing assumptions about the exact nature of the plan benefits.

PBGC believes that these adjustments eliminate most of the mis-estimation of liabilities that would otherwise occur as a result of the simplifications described above. It is impossible to accurately estimate the extent of these distortions, however. A more complex and accurate approach to estimating the liabilities and normal cost may create a more accurate estimate, but this would also create additional computational complexity. We also note that plan actuaries have far more granular information available to them than does PBGC, and thus any model will likely have some need for a "true-up" process.

Over the course of the simulation, pension liabilities grow or shrink each year based on a number of factors:

- Employees earn additional benefits for working another year
- Payments to retirees reduce the remaining liability, as well as assets
- The present value of the remaining liability rises by an interest factor
- Actuarial adjustments may be necessary to reflect other developments

Earned benefits: In an ideal world, PBGC would have detailed information about each employee and about assumed departures and hirings during the year, much as the plan's own actuaries have. In the absence of this information, PIMS assumes that the normal cost will vary from the initial estimate described above according to a straightforward pattern based on changes in the size, age, and general composition of the workforce. In addition to assumptions about age and normal attrition and hiring of employees, the economic model within PIMS also produces variations in the size of the workforce due to overall economic factors combined with random movements between firms. As with the earlier discussion on bankruptcy variables, it is not clear whether PIMS takes sufficient account of the correlation between variables affecting employment status and those affecting bankruptcy probabilities and asset returns.

Payments to retirees: PIMS calculates the amount to be paid out to retirees based on its assumptions on the benefit structure of the plans and the size of the retiree base, along with mortality rates. The size of the retiree base is as initially determined according to the translation process from the Form 5500 information described earlier. This estimate then changes based on movements in the workforce size and the effects of aging combined with retirement rules.

Interest factor: The beginning of year liability for future cash payments needs to be increased by a discount rate to reflect the time value of money and the passage of the year. That is, a dollar that was due in two years from the beginning of the year might be worth roughly 90 cents in beginning of year value (assuming a 5% discount rate), but would be worth about 95 cents in end of year value. (One year later, when due, it would be worth the full dollar.)

Actuarial adjustments: In real life, many of the assumptions made in the Form 5500 have to be revised each year as the world changes and as better information becomes available. Interest rates move up and down, assumptions change about how long employees and retirees will live, assumptions change about retirement patterns and choices about taking lump sum payments at retirement or other variables under the employees' control, and so forth. Each of these variables has an effect on the best estimate of the pension liability. PIMS mimics this process for some variables, but not others. Interest rates do move in PIMS and therefore there are actuarial adjustments for changes in discount rates. However, most of these variables are assumed to be static in PIMS. This is one area where there is considerable room to add potential volatility, as discussed elsewhere in more detail. For example, the model could be extended to allow for parameter uncertainty in the mortality distribution, rather than just assuming draws from a known distribution.

Mortality risk: As noted above, PIMS uses a specified mortality table updated with information from its participants in 2011 to reflect changes in systematic mortality risk. However, it does not perform sensitivity analyses for factors such as selective participation in defined benefit plans.

Changes in benefit formulas: PIMS assumes that plans continue with the same benefit formulas throughout the projection period. The effects of ignoring future voluntary terminations and freezes were discussed above. Another impact of this approach is that no recognition is given to actions taken as firms near insolvency. In the past, some troubled firms have sweetened pension promises in lieu of wage increases or to mitigate the reaction to pay cuts. ERISA and the tax code substantially reduce the effects of these actions, but do not eliminate them. For example, PBGC's guarantees are calculated to exclude all or part of benefit increases put in place near the time of insolvency, with the exclusion phasing out over a five-year period. In addition, severely underfunded plans are forbidden to increase benefits, based on statutory provisions put in a few years ago.

Employees and retirees also make decisions, such as about early retirement or lump sum payments, which can increase pension obligations and/or reduce funding ratios by draining cash from the pension trust. These decisions are effectively assumed by PIMS to continue according to historical patterns, since they only show up indirectly in PIMS through their effects on the Form 5500's information on investment levels, the value of future pension payments and so forth, and these pieces of data are used to calibrate assumptions on key PIMS variables. This is another area where a better reflection of parameter uncertainty would increase the usefulness and accuracy of the PIMS outputs.

Historically, PIMS appears likely to have substantially under-estimated the impact of the largest actual individual claims by projecting a substantially higher funding ratio than was associated with the ultimate claim on PBGC, as discussed in detail above in the section on bankruptcy risk.¹³ There are a variety of reasons for this lower estimate, some of which would be of lesser importance going forward as the result of changes in pension law. The changes in pension liabilities discussed immediately above may have been exacerbated by, for example, riskier investment strategies by firms that were gambling for survival. PBGC might consider conducting an extensive *ex post* analysis on the largest claims and how they were modeled in PIMS preceding the presentation of the claim in order to determine whether there are systematic effects of firms approaching bankruptcy that could be modeled effectively in PIMS.

Discount rates

In the analysis of pension obligations, there are many circumstances in which it is desirable to condense a projected set of future pension payments into a single figure that represents an estimate of the value in today's dollars of those future payouts. This estimate is referred to as a "present value" or a "net present value." The calculation is performed by using one or a series of interest rates referred to as a "discount rate." Essentially, if the chosen discount rate is 5%, then a payment one

¹³ For an elaboration of this point, see CBO, (2005a), "The risk exposure of the Pension Benefit Guaranty Corporation." The CBO study notes that "Bankrupt firms with underfunded pension plans have historically imposed larger costs on PBGC than the level of underfunding they reported immediately prior to bankruptcy," (page 24). This study cites the two "illustrative cases" of Bethlehem Steel and the US Airways pilots' plan, "both of which reported high funding levels in the years immediately prior to their termination, [after which] their funding levels were determined to be less than 50 percent," (page 11). Also see GAO, (2012), "PBGC: Redesigned premium structure could better align rates with risk from plan sponsors," page 12. "PBGC has conducted analyses showing that measures of underfunding are poor predictors of plan termination. In these analyses, PBGC reviewed funding levels... and found the average termination funding level was about 54 percent on the date of termination. For the year previous to termination, the average funding level measurement on which the variable-rate premium was calculated for these plans was about 84 percent." None of these analyses directly show that PIMS itself underestimates these claims, since we do not have access to the historical estimates in the model for specific firms. However, they are suggestive when viewed in combination with the historical performance of PIMS viewed as a predictive model, discussed below.

year out is valued at 95% (100% minus 5%) of its nominal value. A payment two years out is approximately 90% of its nominal value, and so on. The choice of a discount rate is clearly material: small changes in discount rates can lead to large swings in the present value of obligations that are many years in the future.

PIMS uses discount rates for several important purposes:

Calculating the minimum legally required contributions: Simplifying considerably, sponsors are required to contribute at least enough to cover the new benefits earned over the course of the year plus one-seventh of any remaining under-funding. Both the cost in today's dollars of the newly earned benefits and the level of initial under-funding rely on the use of a discount rate to value future pension payouts. For these purposes, PIMS uses the discount rate procedure set in law, which is roughly based on market interest rates for high quality corporate bonds.

Other statutory restrictions: Similarly, heavily underfunded plans have limits on the accrual of new benefits even if they would otherwise be automatic under a pension plan's rules. For these purposes, PIMS uses the discount rate procedure set in law, which is generally the same as that used in calculating the minimum required contributions.

Claims on PBGC: In order for PIMS to report an estimate of PBGC's financial condition at the end of the simulation period, it is necessary to calculate a net present value of any new claims on PBGC, as well as the value of the payments remaining at the end of the simulation for existing claims. For these purposes, PIMS simulates the procedure used by PBGC for its annual financial statements, whereby PBGC attempts to ascertain the market price of its claim liabilities via a survey of prices provided by the major group annuity insurers. PIMS estimates this future market price by using a discount rate equal to the Treasury rate plus 30% of the projected spread of corporate bond rates over Treasury rates.

Recoveries in bankruptcy: PBGC generally receives at least some payout from the estate of a bankrupt plan sponsor to partially offset the losses it incurs by taking over a failed pension plan. (Historically this has averaged a few cents on the dollar, but it varies considerably.) PBGC's claim under bankruptcy law is based on a calculation of underfunding that uses a termination discount rate. PIMS estimates this rate in the future in the same way that it estimates the rate used for claim calculations.

Valuation in today's dollars of PBGC's financial condition at the end of the simulation period: PBGC has chosen to report the results from PIMS by presenting the present value, at the time of the report, of the projected financial condition at the end of the

simulation period, usually 10 years out. For this purpose, the chosen discount rate is the 30-year Treasury rate used in the final period of the simulation.¹⁴

Theoretical considerations

Financial economics theory is unambiguous that a discount rate should reflect the risk of the cash flows being discounted. To a first approximation, the benefit obligations for which PBGC is responsible are close to being risk-free. Within the narrow confines of PIMS, however, the discount rate for the purposes noted above are typically chosen to reflect statute or PBGC's financial accounting practices. To the extent that a discount rate is prescribed by law – such as in the case of calculating minimum required contributions – it is appropriate to use the statutorily required rate.

Generally, because our review is focused on the PIMS model, rather than on the broader structure of the pension insurance system, we take these constraints as given. We note in passing, however, that application of finance principles might suggest using rates other than those implied by these constraints.

We note, however, that for the last purpose noted above – valuation in today's dollars of PBGC's financial condition at the end of the simulation period – PBGC does have latitude in choosing which discount rate to use. We understand that different rates might be useful for different purposes (for example, an appropriately risk-adjusted rate would be better for understanding the true economic value of the liability, whereas the rate used for accounting purposes might be more appropriate for estimating the impact on PBGC's balance sheet). One approach to meeting these various needs would be to simply provide multiple calculations using the various discount rates that are most appropriate for a given use.

Projections of PBGC premiums

PIMS needs to project future premium income at PBGC in order to estimate its financial condition. A portion of PBGC's premium income is based on a per capita charge, which can be calculated easily from PIMS's estimates of the number of covered employees and retirees. The harder part is to project the revenue from variable premiums, which are based on the level of underfunding at each plan.

PIMS cannot simply use the same estimates of funding levels as it uses for the rest of the model's calculations, since it is assumed for those purposes that the minimum legally required contributions are made by each plan sponsor. As described above, that may be a reasonable shortcut for claim purposes, but it would lead to a considerable over-estimate of expected variable premiums if it were not adjusted.

¹⁴ See PBGC, (2012), "Pension Insurance Modeling System: PIMS system description," page 23 of section 2.

Therefore, for premium purposes, there is an essentially *ad hoc* adjustment based on historical experience, reducing the variable premiums by a specific factor that represents the difference between actual past variable premiums and what PIMS would have estimated without this adjustment.

In addition, there is a further discount for the impact of increasing variable premium *rates* going forward, as the result of recent changes in pension law. Higher premium rates make it more expensive to be underfunded and some sponsors will respond by increasing their contributions. PBGC uses a statistical analysis done by an expert at the Office of Management and Budget as a rough approximation of this unknown effect on future funding levels. There has not yet been enough experience to form an empirical conclusion.

It would avoid confusion and potential inconsistencies to explicitly project the level of contributions by each employer. It might also yield more accurate estimates of both future variable premiums and future claims on PBGC, both of which are dependent on contribution levels. The principal trade-off would be in increased computational complexity, although this seems unlikely to be a calculation that would add greatly to the overall work in each simulation run. It is also possible that a poor model to estimate future contributions would decrease accuracy, but it is difficult to see why a carefully thought through formula would be worse than the quite simple approaches used now, where claims use the minimum legal contribution and variable premiums uses a very rough adjustment.

Effects of guarantee limitations on claims on PBGC

PBGC's pension insurance is subject to caps specified in ERISA, particularly a cap on annual pension payments to each individual, currently \$57,477 for someone in a single-employer plan who retires at the normal full retirement age, with reductions in the cap for those who take early retirement. There are also limitations on the payment of special benefits, such as those triggered by plant shutdowns that are contained in some union contracts. PIMS calculates any claims on PBGC by reducing the pension obligations of each plan to reflect these caps and limitations, based on the detailed employee censuses that are derived from the Form 5500 data with the modifications described earlier to better match the liability and normal cost numbers.

Regardless of these caps, however, ERISA also provides that, for single-employer plans, PBGC shall cover up to the full amount of certain priority payment categories, to the extent that there are sufficient assets in the pension trust to cover these payments. PIMS does not attempt to calculate the extent to which insurance is provided above the normal caps and therefore understates the potential claims on PBGC. PIMS does not include this factor because PBGC views the effects as relatively small, in part because the large majority of single-employer pension obligations fall below the normal caps anyway, and the additional data requirements and

computational complexity as too great to merit an attempt at estimation. A compromise might be to include a standard increase in the assumed claims based on historical experience with these special priority categories.

Special factors in ME PIMS

The multiemployer version of PIMS builds on the original single-employer version by tailoring it in two broad ways. First, there are a series of parameters that differ between the two insurance programs, such as the PBGC premium rates and the guarantee caps. The customization to reflect these variations is relatively straightforward and does not merit extensive discussion here. Second, there are three crucial features of the multiemployer program that significantly alter how PIMS needs to work.¹⁵

Joint and several liability of the sponsoring employers: Multiemployer pension plans are the joint responsibility of all of the sponsoring employers, not simply the aggregation of a set of individual benefit programs. As a result, every one of the plan sponsors is jointly and severally responsible for the pension promises. Even if all of the sponsors but one were to default on their obligations, or use the withdrawal option described below, the last employer standing would remain responsible for the entire obligation. If employer withdrawal were not legal, then claims on PBGC would only occur when all of the sponsors became insolvent. In practice, a plan will experience a “mass withdrawal” of sponsors, described below, before all of the sponsors would become bankrupt.

Provisions for withdrawals by plan sponsors: A key complication for PIMS is the ability of employers to pull out of multiemployer plans as long as they follow specific procedures and make a withdrawal payment to cover their share of any underfunding. A “death spiral” can occur in which the impending withdrawal or bankruptcy of some of the sponsors prompts healthier employers to pull out of the plan in order to lock in their share of the underfunding at a level that still reflects the participation of all the other sponsors. When the weaker employers then withdraw or default, the proportionate impact on the remaining employers becomes that much higher because the loss is spread over fewer, and less healthy, firms. This dynamic can lead to another round of withdrawals and defaults. Since the healthy firms are likelier to be able to find the funds to pay their withdrawal liability, there is also an adverse selection problem whereby the average financial strength of the remaining sponsors declines.

Further exacerbating this withdrawal problem, there are caps on the maximum withdrawal liability that can be charged to a sponsor that result in the charge for

¹⁵ These differences are spelled out most clearly in PBGC, (2011), “Multiemployer PIMS: Key differences between SE-PIMS and ME-PIMS.”

withdrawal often being quite considerably less than the present value of future contributions would have been. This outcome provides a further incentive for firms to withdraw and leaves those firms that remain with greater and greater proportionate liabilities.

PIMS does not attempt to model withdrawals by individual sponsors. Instead, it assumes that either all sponsors withdraw simultaneously, a “mass withdrawal,” or none do. PIMS employs a formula to calculate the probability of a mass withdrawal in any given year. Estimating employer withdrawals is very difficult and it is not clear that the parameters in PIMS are accurate estimates of this. There is relatively little history of employer withdrawals from which to form statistical judgments and the bulk of these have come in recent years, making it harder to judge under what conditions they might occur in the future.

Estimating withdrawals is a major issue for the reliability of PIMS projections of multiemployer claims on PBGC since there has clearly been a major shift in the viability of multiemployer plans in recent years. PIMS could be made more accurate if PBGC improved the statistical basis for modeling plan withdrawals, probably including an investigation of the role of individual firm withdrawals and the death spiral described above. Sensitivity analysis would also be useful given the very significant uncertainty about this parameter.

Ability of multiemployer pension trustees to reduce the pension obligations:

Unlike single-employer plans, multiemployer plans have the legal ability, and in many cases legal requirement, to reduce their pension promises if maintaining the old levels becomes unviable. ERISA specifies a series of stages of peril that are determined by actuarial forecasts of when and whether the plan will run out of the cash necessary to pay the promised pension benefits. Plans that appear able to meet all of their obligations have neither the right nor the obligation to reduce benefits. Plans that are projected to become insolvent well into the future may have their benefits cut back by their trustees, if raising the contributions sufficiently to solve the problem appears infeasible, including for the reason that the contributions would be too high for the sponsoring firms to support financially. As the date of projected insolvency draws nearer, tougher benefit cuts become optional and then mandatory.

The ability to reduce benefits, which is not available to single-employer plans, can significantly change the ultimate claim on PBGC, either by avoiding the necessity of a claim in the first place or by reducing the eventual claim through reductions in the promises prior to a plan’s demise. PIMS employs a very simple formulaic approach to estimate the effect of benefit cutbacks. In particular, it assumes that plans that get into trouble initially cut back their benefits by the full 20% that is allowed under those circumstances. As plans enter still worse phases of peril, further cutbacks are assumed.

Peer review of Multiemployer PIMS

In 2012, Buck Consultants performed a peer review of ME PIMS.¹⁶ They highlighted four areas where they believed it was particularly important for PBGC to improve ME PIMS, along with a series of lesser recommendations and a number of other areas where they were comfortable with the model as it currently exists. The key areas for change were:

Number of covered participants and plans: Buck states that ME PIMS essentially assumes a static level of active population in multiemployer plans, but recommends that there be an assumed decline, based on history and current indications. There is likely to be an interaction between the steps necessary to restore the plans to sound funding, which will often be painful, and the number and type of plan withdrawals. Buck suggests that an assumed decline in participant count of 1.5% a year would be reasonable. We do not have any independent information with which to evaluate Buck's specific suggestion, but it does appear that sensitivity analysis with regard to population declines would further enhance the model. Ideally, there would be a correlation in the model between changes in the plan population and other factors that affect funding levels and the strength of firms.

PBGC has indicated that they concur in principle with this recommendation and are implementing changes for the coming year similar to what Buck recommended.

Contribution levels: Buck states "we think that the level of [contribution] increases factored into the programming are unrealistically high, especially given that, in many cases, while the contributions are increasing two to four fold over the next 10 years or so, the benefit accruals are being drastically reduced or even eliminated. Trustees of most ME plans will not approve Rehabilitation Plans with such extremes, on the basis that such draconian demands would result in mass exodus of employers from the plans. The disconnect stems largely [from] the fact that ME-PIMS explicitly does not account for the 'exhaustion of reasonable measures clause' [that allows trustees to avoid taking actions that would be severely detrimental to the plans]."

PBGC has indicated that they concur in principle with this recommendation and are implementing a version of it for this coming year. It could also be useful to have sensitivity analyses to indicate the level of importance of these assumptions and the effects on PBGC's financial condition of a range of reasonable assumptions.

Funding improvement plan/rehabilitation plan hierarchy: Buck extends the same logic to point out that they believe it to be unrealistic to assume, as ME PIMS does, that trustees will take the full range of steps that would theoretically be required under improvement/rehabilitation plans to remedy funding problems.

¹⁶ The remainder of this section draws on: Kai Petersen, Darren French, et al., (2012), "ME-PIMS peer review."

This recommendation is an extension of the point about the “exhaustion of reasonable measures clause.”

PBGC has indicated that they concur in principle with this recommendation as well, although they do note that their original assumption that the full required benefit cuts would be made was based on the recommendation of an external consultant. This recommendation is complicated to put in practice, in part because of data problems, so there will be an estimate used in this coming year with the intention to have a more detailed approach in future years.

Employer withdrawal assumptions: As noted above, employers have the ability to withdraw from multiemployer plans and healthy firms are increasingly doing so in order to limit their risk to the creditworthiness of other plan sponsors. Buck has a number of technical suggestions on how the ME PIMS calculations might be improved in estimating withdrawals. This is another area where sensitivity analyses and other ways of handling parameter uncertainty may strengthen the model. PBGC has indicated that they agree in principle and are making a series of changes to next year’s model to implement the Buck recommendations.

The relatively crude assumptions in ME-PIMS in regard to the composition and size of the plan universe, the extent of sponsor withdrawals, and the extent of benefit cutbacks by troubled plans may well have been justified when ME-PIMS was first created. At that time, PBGC’s expected claims from this program were far smaller than were expected from the single-employer program. However, the rapid increase in expected claims on PBGC from multiemployer plans suggests that ME-PIMS did not initially reflect fully the risks in the multiemployer system. (The multiemployer program has gone from a net worth deficiency that had never reached as much as \$1 billion throughout its history until 2010, to a 2012 estimate of a deficiency of over \$5 billion, with a potential for continued rapid escalation.)

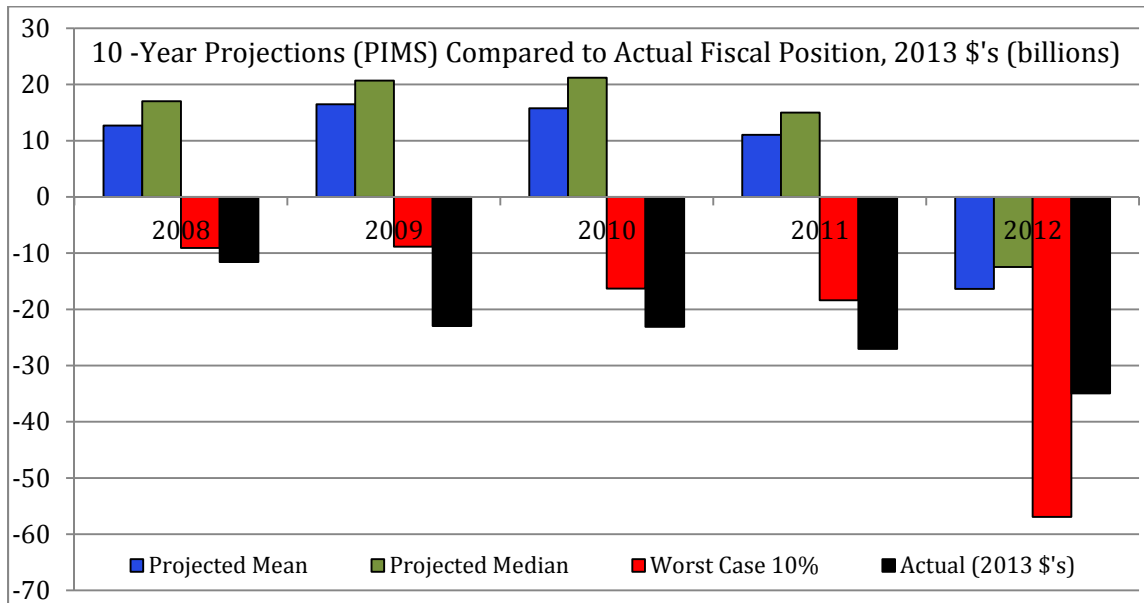
Historical predictive performance of PIMS

PBGC releases a comprehensive assessment of its financial standing in its annual reports, complete with a series of projections estimating its financial position in a decade. It relies on PIMS to arrive at these projections. PBGC runs 500 scenarios of PIMS and reports the distribution of its estimated financial condition 10 years out. As discussed above, PBGC clearly states every year in their annual reports that the “the PIMS model is not predictive” and that “PIMS provides a range of possible future outcomes.” However, it seems clear that many readers – both inside and outside government - are treating the projections as predictions, particularly in the absence of any other estimates that have truly detailed quantitative analysis underlying them.

Although there is potential value in examining the predictive accuracy of estimates, we underscore three important caveats.

First, we have only a very limited number of projections with which to assess the model's predictive accuracy. Second, most of the projections that can be used include the Great Recession. Recognizing that many of the macroeconomic features of the model experienced "tail draws," it should not be surprising if the median projections of PIMS turned out to be overly optimistic. Finally, we also note that any *ex post* evaluation of projections also needs to account for any legislative changes that altered plan rules as well as firm behavior (most notably with the Pension Protection Act in 2006). With these important caveats in mind, we turn to a review of the historical experience. There are only 5 projections from PBGC annual reports where the full simulation period has now ended, those from 10-year projections made in 1998 through 2002 (with projections for 2008 to 2012). Viewed as predictions, they consistently and significantly underestimated the financial risks facing PBGC. Actual deficits were in the worst 10% of simulation runs in the 1998-2001 projections. For the 2002 projection, the actual result was within one standard deviation of the mean, but was still close to \$20 billion worse than either the mean or the median of the projection. The chart below contrasts the PBGC's actual financial position against three outcomes projected by PIMS ten years prior: the mean fiscal position for all potential outcomes, the median fiscal position across all potential outcomes, and the worst case 15% fiscal position (that is, the projected financial position at the 15% mark on the probability distribution).

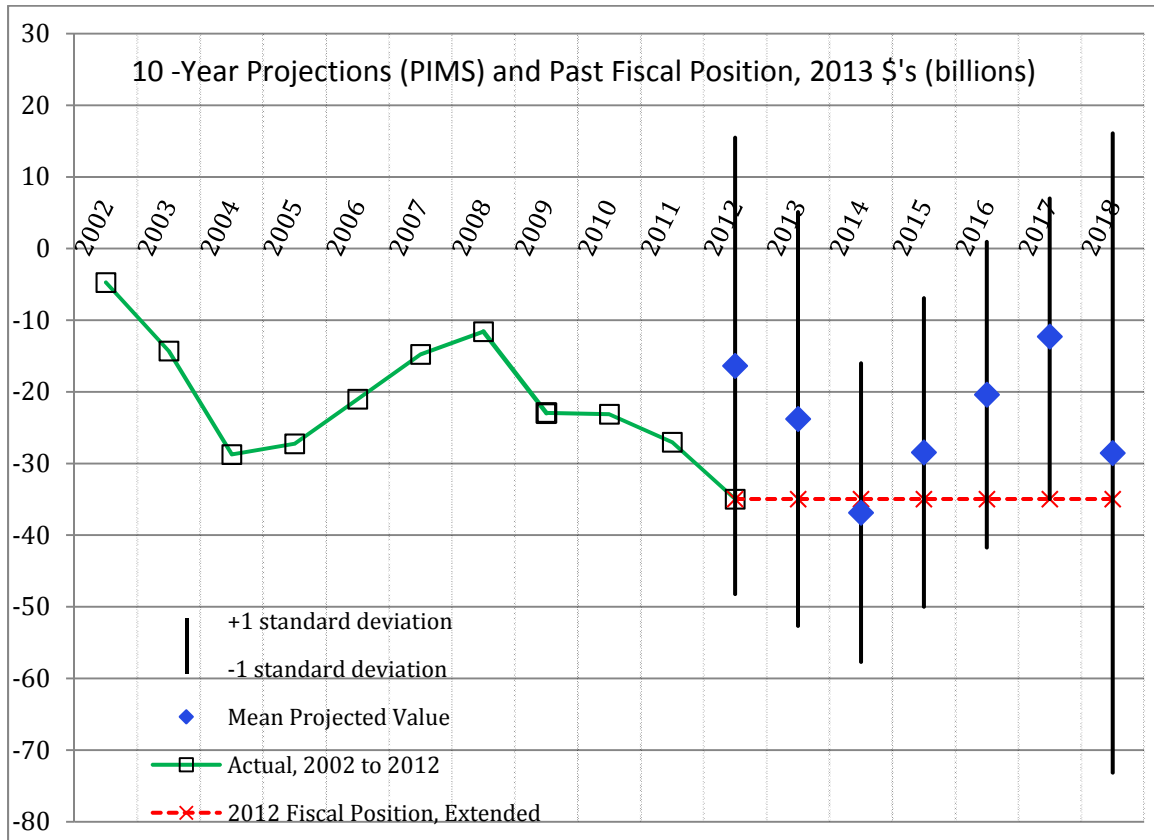
Figure 4: Evaluating PIMS Projections Against Recent Historical Experience



For each of the more recent projections, made between 2003 and 2008, PBGC's current deficit at the end of 2012 already exceeds the mean and median projections

made for 2013 to 2018. While PBGC's financial standing could improve in the interim, this would break the pattern of steady increases in the deficit in every year since 2008 and would be inconsistent with the historical insufficiency of PBGC's premiums to cover its new claims. The figure below shows PBGC's past financial position as well as the projections from PIMS for the near future. PIMS projections include the mean projected financial position, as well as the projected dispersion of outcomes, one standard deviation above and below the mean.

Figure 5: PIMS Projections, 2012 - 2018



We believe it is too early to draw any conclusions about the 2009 and subsequent projections.

There are two main classes of explanations for the substantial and consistent divergence between the median and mean projections by PIMS and subsequent actual PBGC deficits. First, it may be that the factors noted throughout this paper that bias estimates in PIMS towards excessive optimism significantly outweigh the factors that bias estimates towards pessimism. It is difficult to fully quantify this without creating a more refined version of PIMS that would eliminate these biases, but it certainly appears that PIMS is optimistically biased. Second, PIMS could be unbiased and PBGC simply experienced unusually bad luck, which is consistent with

the projection periods incorporating at least one, and sometimes two, of the major bear markets associated with the “dot-com bubble” and then the financial crisis and ensuing Great Recession. Of course, these projection periods have also included the major bull markets that preceded the downturns and the very sharp rebound after the both bear market, which has led stocks in the US to recent new highs in nominal terms.

As suggested above, it would be useful to have PBGC do a serious *ex post* analysis to assess, to the extent possible, which of the possibilities contains the most explanatory power, or to offer alternative explanations.

It should also be noted that PIMS appropriately assumes that pension laws remain static during the projection period, whereas the reality is that those laws change. In general, these changes are likely to pull the actual results back closer to the projections, since extreme positive or negative results for pension funds may produce offsetting legislative changes.

PIMS as a planning tool

In addition to the discussions of technical aspects of PIMS and potential improvements in that regard, it is worth focusing more on the purpose of PIMS, how its results are used and communicated, and how it might be supplemented. PBGC consistently argues that the model is intended to show the range of financial conditions that the agency might encounter going forward, but does *not* produce a forecast or a range of forecasts. However, this is at odds with how almost everyone external to PBGC appears to use the information, including members of Congress, Congressional staff, members of the Administration, and external parties, who commonly view PIMS results as a forecast.

We recognize that PBGC believes this interpretation may be inappropriate, but if so, then it is important to note that they have not sufficiently clarified their reasoning to external users. Further, it appears to us that PIMS *could* be used as a forecasting tool, although a number of improvements to the model would be required if the model were to be an effective tool for this purpose.

One obstacle to using PIMS as a forecasting tool may be the great caution PBGC exercises whenever judgment might be called for in the model. That is, there is clearly a strong concern about using any assumptions other than those directly obtainable by looking at history. This stance makes it difficult to incorporate reasonable estimates of how the future may differ from the past. In some cases, such estimates could be obtained from reasonable outside sources. In other cases, PBGC may be in the best position to set the estimate. It appears that this use of judgment could be appropriate, particularly if the process is open and transparent.

As discussed elsewhere in this paper, sensitivity analyses, including scenario analyses, represent one way to convey the uncertainties while also demonstrating which variables are most important in determining the financial condition of PBGC. However, one disadvantage of PIMS is that the model is so complicated that it is difficult for non-experts, and even many experts, to understand the main drivers of the projections. While sensitivity analyses would help with this, it may also be potentially useful to supplement this with a deterministic model that generated a base case and some alternative scenarios. Alternatively, PBGC could find a way to show the implied course of key variables, such as claims, expenses, and premium income over the projection period for the median case or for an average of a set of cases around the median. Ideally, either approach would allow an external party to view simplified income statements and balance sheets for PBGC on a year by year basis.

Using a simplified model, or a simplification of the outcome of the full PIMS model, has clear disadvantages. However, these must be weighed against the problems created by the current status of PIMS as essentially a “black box”. This supplemental approach could improve the credibility of the outcomes, through greater transparency, and allow key parties to more clearly understand the drivers of PBGC’s condition.

In the interest of accuracy, governance and transparency, PBGC might consider the lessons of the Technical Panel process used to review the methods and assumptions used in the Medicare and Social Security Trustees’ Report. The Social Security Advisory Board – an independent, bipartisan board with members appointed by the President and Congress – appoints an independent technical panel every three years to provide a rigorous external assessment of the actuarial models. This process has, over the years, greatly increased the transparency of the actuarial models and has also led to a number of recommendations for improvement that have been subsequently incorporated into the models.

Throughout our review process, it has also become clear that model transparency would be greatly enhanced if PBGC were to provide more accurate and up-to-date documentation. Although there is no shortage of documentation when measured by page count, there are internal inconsistencies in some places, whereas in other places the underlying assumptions are not specified. More generally, a more externally transparent process for the ongoing governance evaluation of the model would likely increase public understanding of and confidence in the model.

References

- Adler, B. E., Capkun, V., and Weiss, L. A. (2013). Value destruction in the new era of Chapter 11. *Journal of Law, Economics, and Organization*, 29(2), 461-483.
- Adonov, Aleksander, Rob Bauer, and Martijn Cremers. (2013). Pension fund asset allocation and liability discount rates: camouflage and reckless risk taking by U.S. public plans? http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2070054.
- Almeida, Heitor and Thomas Philippon (2007). The risk-adjusted cost of financial distress. *Journal of Finance*, 62(6), 2557-2586.
- An, H., Huang, Z., and Zhang, T. (2013). What determines corporate pension fund risk-taking strategy? *Journal of Banking & Finance*, 37(2), 597-613.
- Anderson, Joseph M. (1999). Models for retirement policy analysis: Report to the Society of Actuaries. Society of Actuaries. See chapter 11: PIMS: The Pension Insurance Modeling System. http://www.soa.org/files/research/projects/chapter_11.pdf.
- Anyane-Ntow, K. (2011). Accounting information and its relationship to corporate financial distress process. *Journal of Applied Business Research*, 7(3), 29-35.
- Astebro, T., and Winter, J. K. (2012). More than a dummy: The probability of failure, survival and acquisition of firms in financial distress. *European Management Review*, 9(1), 1-17.
- Bean, D. F., and Bernardi, R. A. (2000). Underfunding pension obligations while paying dividends: evidence of risk transfers. *Critical Perspectives on Accounting*, 11(5), 515-530.
- Beaver, W. H., M. F. McNichols, and J. Rhie (2005). Have financial statements become less informative? Evidence from the ability of financial ratios to predict bankruptcy, *Review of Accounting Studies*, 10(1), 93-122.
- Bergstresser, D., Desai, M., and Rauh, J. (2006). Earnings manipulation, pension assumptions, and managerial investment decisions. *The Quarterly Journal of Economics*, 121(1), 157-195.
- Bhamra, Harjoat S., Lars-Alexander Kuehn and Ilya A. Strebulaev. (2010). The levered equity risk premium and credit spreads: a unified framework. *Review of Financial Studies* 23 (2), 645-703.
- Bharath, S. T., and T. Shumway (2008). Forecasting default with the Merton distance to default model. *Review of Financial Studies*, 21(3), 1339-1369.

- Bodie, Zvi. (1996). What the Pension Benefit Guaranty Corporation can learn from the Federal Savings and Loan Insurance Corporation. *Journal of Financial Services Research*, 10(1), 83-100.
- Bodie, Zvi. (2006). On asset-liability matching and federal deposit and pension insurance. *Federal Reserve Bank of St. Louis Review*, 88(4), 323-29.
- Boyce, Steven, and Richard Ippolito. (2002). The cost of pension insurance. *The Journal of Risk and Insurance*, 69(2), 121-170.
- Brown, Jeffrey R. (2008). Guaranteed trouble: The economic effects of the Pension Benefit Guaranty Corporation. *Journal of Economic Perspectives*, 22(1), 177–198.
- Campbell, John Y. (2000). Asset pricing at the millenium. *Journal of Finance*, 55(4), 1515-1567.
- Campbell, J. Y., J. Hilscher, and J. Szilagyi (2008). In search of distress risk. *The Journal of Finance*, 63(6), 2899-2939.
- Campbell, John Y., Andrew W. Lo and A. Craig MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton University Press.
- CBO. (2005a). The risk exposure of the Pension Benefit Guaranty Corporation. Congressional Budget Office.
<http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/66xx/doc6646/09-15-pbgc.pdf>.
- CBO. (2005b). A guide to understanding the Pension Benefit Guaranty Corporation. Congressional Budget Office.
<http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/66xx/doc6657/09-23-guidetopbgc.pdf>.
- Chang, Bo Young, Peter Christoffersen and Kris Jacobs. (2013). Market skewness risk and the cross section of stock returns. *Journal of Financial Economics*, 107(1), 46-68.
- Chava, S., and R. A. Jarrow (2004). Bankruptcy prediction with industry effects. *Review of Finance*, 8(4), 537-569.
- Cheng, K. F., Chu, C. K., and Hwang, R. C. (2010). Predicting bankruptcy using the discrete-time semiparametric hazard model. *Quantitative Finance*, 10(9), 1055-1066.
- Chen, Hui. (Forthcoming). Macroeconomic conditions and the puzzles of credit spreads and capital structure. *Journal of Finance*.
- Cox, Ingersoll and Ross. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2), 385-408.

Delaney, K. J. (1999). *Strategic bankruptcy: How corporations and creditors use chapter 11 to their advantage*. Berkeley: University of California Press.

Duan, Y., Hotchkiss, E. S., and Jiao, Y. (2012). Corporate pensions and financial distress. Unpublished Manuscript.

Duffie, D., Eckner, A., Horel, G., and Saita, L. (2009). Frailty correlated default. *The Journal of Finance*, 64(5), 2089-2123.

Duffie, D., L. Saita, and K. Wang (2007). Multi-period corporate default prediction with stochastic covariate. *Journal of Financial Economics*, 83(3), 635-665.

Elliott, Douglas J. (2004a). PBGC: When will the cash run out? *The Center On Federal Financial Institutions*.
<http://coffi.org/pubs/PBGC%20When%20Will%20the%20Cash%20Run%20Out%20v8.pdf>.

Elliott, Douglas J. (2004b). PBGC: Updated cash flow model from COFFI. *The Center On Federal Financial Institutions*.
<http://coffi.org/pubs/PBGC%20Updated%20Cash%20Flow%20Model%20from%20COFFI.pdf>.

Elliott, Douglas J. (2005a). PBGC: Policy options. *The Center On Federal Financial Institutions*.
<http://coffi.org/pubs/PBGC%20Policy%20Options%20%20January%206%202005.pdf>.

Elliott, Douglas J. (2005b). PBGC: Legislation may not restore solvency. *The Center On Federal Financial Institutions*.
<http://coffi.org/pubs/PBGC%20Legislation%20May%20Not%20Restore%20Solvency.pdf>.

Epstein, Joshua M. (2007). *Generative Social Science*. Princeton: Princeton University Press.

Epstein, Joshua M. (2009). Modelling to contain pandemics. *Nature*. 460, 687.

Eubank, Stephen, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang. (2004). Modelling disease outbreaks in realistic urban social networks. *Nature*, 429, 180-184.

Fama, Eugene (1963). Mandelbrot and the stable Paretian hypothesis. *Journal of Business*, 36(4), 420-429.

Fama, Eugene (1965). The behavior of stock market prices." *Journal of Business*, 38(1), 34-105.

Fama, Eugene F., and Kenneth R. French, (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3-56.

Figlewski, S., Frydman, H., and Liang, W. (2012). Modeling the effect of macroeconomic factors on corporate default and credit rating transitions. *International Review of Economics and Finance*, 21(1), 87-105.

Foreman, R. D. (2003). A logistic analysis of bankruptcy within the US local telecommunications industry. *Journal of Economics and Business*, 55(2), 135-166.

GAO. (2012). PBGC: Redesigned premium structure could better align rates with risk from plan sponsors. United States Government Accountability Office. <http://www.gao.gov/assets/650/649838.pdf>.

Grice, J. S., and Dugan, M. T. (2001). The limitations of bankruptcy prediction models: Some cautions for the researcher. *Review of Quantitative Finance and Accounting*, 17(2), 151-166.

Hammond, Ross A. (2009). Complex systems modeling for obesity research. *Preventing Chronic Disease*, 6(3), A97. http://www.cdc.gov/pcd/issues/2009/jul/pdf/09_0017.pdf.

Haubrich, Joseph, George Pennacchi and Peter Ritchken. (2012). Inflation expectations, real rates, and risk premia: evidence from inflation swaps. *Review of Financial Studies*, 25(5), 1588 – 1629.

Hawthorne, F., 2008. *Pension Dumping: the Reasons, the Wreckage, the Stakes for Wall Street*. New York: Bloomberg Press.

Hillegeist, S. A., E. Keating, D. Cram, and K. Lunstedt. (2004). Assessing the probability of bankruptcy. *Review of Accounting Studies*, 9(1), 5-34.

Hol, S. (2007). The influence of the business cycle on bankruptcy probability. *International Transactions in Operational Research*, 14(1), 75-90.

Ippolito, R. A., and James, W. H. (1992). LBOs, reversions and implicit contracts. *The Journal of Finance*, 47(1), 139-167.

Jegadeesh, Narasimhan, and Sheridan Titman. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91.

Jones, S. and D. Hensher. (2004). Predicting firm financial distress: A mixed logit model. *The Accounting Review*, 79(4), 1011-1038.

Lau, A. (1987). A five-state financial distress prediction model. *Journal of Accounting Research*, 25(1), 127-138.

Lemmon, M., Ma, Y. Y., and Tashjian, E. (2009). Survival of the fittest? Financial and economic distress and restructuring outcomes in Chapter 11. Working Paper, University of Utah.

- Longstaff, F. A. and E.S. Schwartz, (1992). Interest rate volatility and the term structure: A two-factor general equilibrium model. *Journal of Finance*, 47(4), 1259–82.
- Marcus, A. (1987). Corporate pension policy and the value of PBGC insurance. Chapter 3 in *Issues in Pension Economics*, edited by Z.Bodie, J.Shoven, and D. Wise. Chicago: University of Chicago Press.
- Mensah, Y. M. (1984). An examination of the stationarity of multivariate bankruptcy prediction models: a methodological study. *Journal of Accounting Research*, 22(1), 380-395.
- Nam, C. W., Kim, T. S., Park, N. J., and Lee, H. K. (2008). Bankruptcy prediction using a discrete? Time duration model incorporating temporal and macroeconomic dependencies. *Journal of Forecasting*, 27(6), 493-506.
- Ohlson, J. A. (1980). Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research*, 18(1), 109-131.
- Orr, D. V. (1998). Strategic bankruptcy and private pension default. *Journal of Economic Issues*, 32(3), 669-687.
- Pennacchi, George and Christopher Lewis. (1994). The value of Pension Benefit Guaranty Corporation insurance. *Journal of Money, Credit, and Banking*, 26(3), 735-753.
- PBGC. (various years). Annual report. Pension Benefit Guaranty Corporation. <http://www.pbgc.gov/res/annual-reports.html>.
- PBGC. (various years). Pension insurance data book. Pension Benefit Guaranty Corporation. <http://www.pbgc.gov/res/data-books.html>.
- PBGC. (1999). Pension insurance data book 1998. Pension Benefit Guaranty Corporation, pages 10-17. <http://www.pbgc.gov/documents/1998databook.pdf>.
- PBGC. (2005). An Analysis of frozen defined benefit plans. Pension Benefit Guaranty Corporation. http://www.pbgc.gov/docs/frozen_plans_1205.pdf.
- PBGC. (2010). Pension Insurance Modeling System: User guide. Pension Benefit Guaranty Corporation. September 2010.
- PBGC. (2011). Multiemployer PIMS: Key differences between SE-PIMS and ME-PIMS. Pension Benefit Guaranty Corporation.
- PBGC. (2012). Pension Insurance Modeling System: PIMS system description. Pension Benefit Guaranty Corporation. Updated December 17, 2012.
- PBGC. (2013). Single-employer DB plan freezes. Pension Benefit Guaranty Corporation. <http://www.pbgc.gov/Documents/Plan-Freeze-2013-EA-meeting.pdf>.

- Petersen, Kai, and Darren French, et al. (2012). ME-PIMS peer review. Buck Consultants.
- Pollet, J., and M. Wilson. (2010). Average correlation and stock market returns. *Journal of Financial Economics*, 96(3), 364–380.
- Rauh, J. D. (2009). Risk shifting versus risk management: investment policy in corporate pension plans. *Review of Financial Studies*, 22(7), 2687-2733.
- Saltelli, A., S. Tarantola, and F. Campolongo. (2000). Sensitivity analysis as an ingredient of modeling. *Statistical Science*, 15(4), 377-395.
- Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *Journal of Business*, 74(1), 101-124.
- Sterman, J.D. (2006). Learning from evidence in a complex world. *American Journal of Public Health*, 96(3), 505-514.
- Tesfatsion, L. and K.L. Judd, editors. (2006). *Handbook of Computational Economics, Vol 2*. Amsterdam: Elsevier.
- Tweedale, G., and Warren, R. (2004). Chapter 11 and Asbestos: Encouraging Private Enterprise or Conspiring to Avoid Liability? *Journal of Business Ethics*, 55(1), 31-42.
- Vasicek, Oldrich (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2), 177-188.
- Wilcox, David W. (2006). Reforming the defined-benefit pension system. *Brookings Papers on Economic Activity*, 37(1), 235 – 304.
- Zmijewski, M. E. (1984). Methodological issues related to the estimation of financial distress prediction models. *Journal of Accounting Research*, 22, 59-82.

Appendix A

The Project Team

Jeffrey R. Brown is the William G. Karnes Professor of Finance and Director of the Center for Business and Public Policy at the University of Illinois' College of Business and a Research Associate of the National Bureau of Economic Research, where he also serves as Associate Director of the NBER Retirement Research Center and as Editor of the *Tax Policy and the Economy* series. Dr. Brown earned his BA from Miami University (Ohio), his MPP from Harvard University, and his PhD in economics from MIT. Dr. Brown has previously served on the faculty at Harvard's Kennedy School of Government, as a Senior Economist with the White House Council of Economic Advisers, and as an economist for the President's Commission to Strengthen Social Security. In 2006, he was nominated by the President and confirmed by the US Senate as a Member of the Social Security Advisory Board, a position in which he served through September 2008. He also serves as a Trustee for TIAA and on the Board for the American Risk and Insurance Association. Professor Brown has published extensively on topics related to retirement income security, including papers on Social Security, public and private pensions, PBGC, and markets for annuities, life insurance and long-term care insurance. In 2008, Dr. Brown received the Early Career Scholarly Achievement Award from the American Risk and Insurance Association, and the TIAA-CREF Paul A. Samuelson Award for Outstanding Scholarly Writing on Lifelong Financial Security. He is a founding co-editor of the *Journal of Pension Economics and Finance*.

Douglas Elliott is a Fellow at The Brookings Institution, where he analyzes the global financial sector, private sector pensions, state and local government pensions, federal financial institutions (including PBGC), and the Euro Crisis. He graduated from Harvard College *magna cum laude* with an AB in Sociology in 1981 and earned an MA in Computer Science from Duke University in 1984. Mr. Elliott is an expert on PBGC and appears to have written more papers on that institution than any other external analyst. The New York Times referred to his PBGC analyses as "refreshingly understandable" and "without a hint of dogma or advocacy." He appears to be the only external analyst to have created a detailed annual cash flow model projecting the PBGC's financial condition out for decades into the future. An investment banker to the insurance industry for two decades, principally at J.P. Morgan, he was the founder and principal researcher for the Center On Federal Financial Institutions (COFFI), a think tank devoted to the analysis of federal insurance and lending activities. He wrote the book, *Uncle Sam in Pinstripes: Evaluating the US Federal Credit Programs*, the only comprehensive review of the federal government's credit activities to be written in the last quarter century. Mr. Elliott has testified before both houses of Congress and participated in numerous speaking engagements, as well as appearing widely in the major media outlets.

Ross A. Hammond is Senior Fellow in Economic Studies at The Brookings Institution, where he is also Director of the Center on Social Dynamics & Policy. His primary area of expertise is modeling complex dynamics in economic, political, and public health systems. He has over 15 years of experience with mathematical and computational modeling techniques from complex systems science. His current research topics include: obesity etiology and prevention, behavioral epidemiology, food systems, tobacco control, crime, corruption, segregation, trust, and decision-making. Hammond received his B.A. from Williams College and his Ph.D. from University of Michigan. He has authored numerous scientific articles, and his work has been featured in *New Scientist*, *Salon*, *The Atlantic Monthly*, *Scientific American*, and major news media. Hammond was recently appointed to the Institute of Medicine and National Research Council committee on the Framework for Assessing the Health, Environmental, and Social Effects of the Food System. He also currently serves on the editorial board of the journal *Childhood Obesity*, on the steering committee for the NIH Comparative Modeling Network of the National Collaborative on Childhood Obesity Research, and as a member of the NIH MIDAS (Models of Infectious Disease Agent Study) and NICH (Network on Inequality, Complexity, and Health) networks. Hammond has been a consultant to the World Bank, the Asian Development Bank, and the Institute of Medicine, and he is currently a Public Health Advisor at the National Institutes of Health. He has taught computational modeling at Harvard School of Public Health, the University of Michigan, Washington University, and the NIH/CDC Institute on Systems Science and Health. Hammond has previously held positions as the Okun-Model Fellow in Economics, an NSF Fellow in the Center for the Study of Complex Systems at University of Michigan, a visiting scholar at The Santa Fe Institute, and a Consultant at PricewaterhouseCoopers LLP.

Tracy Gordon is a Fellow in Economic Studies at the Brookings Institution. She is also an affiliated scholar with the Urban Institute-Brookings Institution Tax Policy Center. Her research is in state and local budgeting and public finance, political economy, and urban economics. Before joining Brookings, Gordon was an Assistant Professor in the School of Public Policy at the University of Maryland, College Park. She was also a research fellow at the Public Policy Institute of California, where she is now an adjunct fellow. She holds a Ph.D. in public policy with a concurrent M.A. in economics from the University of California, Berkeley. Gordon has authored reports and journal articles on state and local budgeting, local property taxes, the local initiative process, and so-called “private governments” or common interest developments. Some recent publications include: “The Federal Stimulus Programs and Their Effects,” (with Gary Burtless) in *The Great Recession*, (Russell Sage Foundation), “State and Local Fiscal Institutions in Recession and Recovery,” in *Oxford Handbook on State & Local Government Finance*, “Addressing Local Fiscal Disparities,” in *Oxford Handbook of Urban Economics and Planning* (Oxford University Press), and “State and Local Government Finances: Where We’ve Been, Where We’re Going, and How to Get There,” 2010 National Tax Association Papers and Proceedings.

Appendix B

Key Sources and Background Literature

Unless otherwise noted, the basic information on PBGC and the PIMS model contained in this paper has been adapted from a few key sources. Information on PBGC's financial position as well as the predictions of the PIMS model can be found in the Annual Reports released by PBGC. Detailed data on PBGC finances and the behavior of covered pension plans is largely taken from the Pension Insurance Data Books released by PBGC annually. General information on the composition of the PIMS model can be found in PBGC's Pension Insurance Data Book (1998), which provides a concise introduction to the model in the first year of its use. A more detailed account of the simulation parameters and procedures utilized in PIMS is provided by Joseph M. Anderson (1999) in a report to the Society of Actuaries. More technical details are drawn from "Pension Insurance Modeling System: User's guide," (2010) and "Pension Insurance Modeling System: PIMS system description," (2012), which serve as instructional manuals accompanying the PIMS model.

Despite the important role of PBGC in the US economy, there is only a modest literature on the economic effects and future of PBGC. Although small, this literature spans a variety of topics, ranging from the implications of PBGC insurance on retirement security generally, to proposals to improve the long-term financial health of PBGC, to evaluations of the accuracy of PBGC premium pricing and modeling procedures.

Early literature on PBGC explored the costs, both private and social, of providing an insurance guarantee to pension plans. Pennacchi and Lewis (1994) introduced a model to estimate the cost of PBGC insurance as a put option with a random exercise date depending on firm bankruptcy. They also provided the tools to estimate this cost to PBGC as a lump-sum payment. Their approach built on and improved the work of Marcus (1987), which explored the same question of valuing PBGC's guarantee, but conceived of as a forward contract rather than as a (contingent) put option. Boyce and Ippolito (2002) also pursued this question, utilizing a Monte Carlo method for modeling the cost to PBGC of providing insurance. Comparing their results to the level of PBGC premiums at the time, Boyce and Ippolito found that PBGC's premiums significantly understated the costs of insurance provision compared to a theoretical "true market cost," suggesting that this creates moral hazard. Bodie (1996, 2006) also explored problems related to asymmetric information. He compared PBGC to the Federal Savings and Loan Insurance Corporation, specifically noting that both institutions faced a fundamental disjoint between assets and liabilities, and arguing that, absent proper policy reform, the former federal corporation may go the way of the latter. Drawing and expanding upon Bodie's work, Brown (2008) noted that PBGC's mounting deficits and the general decline in defined benefit pensions are related to problems with PBGC's pricing of premiums, anemic funding requirements, and the lack of transparency regarding plan funding status.

Other scholars have focused on the pension funding and investment decisions made by plan sponsors. For example, Bean and Bernardi (2000) studied the relations between

unfunded pension liabilities and earnings, dividends, and cash flows. They provide evidence that some pension plan managers were able to divert risks from shareholders onto the pension participants, and potentially, via PBGC insurance, onto society at large. Wilcox (2006) and Brown (2008) discuss a wider range of features of the PBGC insurance program that have contributed to significant underfunding. Although some of the root causes were addressed in the Pension Protection Act of 2006, not all the incentive problems were fixed. Adonov, Bauer and Cremers (2013) contend that the US pension structure allows plan managers, particularly of larger and older plans, to conceal their true level of underfunding by shifting their investments toward riskier assets and thus pushing up their discount rates. They identify this behavior as posing an undue risk for future workers, and potentially for society at large.

Aside from the scholarly sources noted above, there is a substantial literature concerning PBGC among governmental and actuarial organizations. For example, the Congressional Budget Office (2005a) issued a detailed report on PBGC's premium structure, financial projections, and ability to properly monitor plan underfunding. The study introduced its own model for estimating the probable losses facing PBGC, and explored a variety of potential reforms to PBGC's premiums to help reduce its deficit. The CBO report investigated the cases of Bethlehem Steel and US Airways as examples of how certain PBGC procedures hindered the proper assessment of plan liabilities. Finally, we should note the PBGC-focused modeling work of Elliott (2004a, 2004b), who created a cash-flow model for predicting PBGC's financial position that has informed and served as a foundation for a number of his subsequent works on PBGC finances. Elliott's cash flow model was an early, independent attempt to evaluate the health of PBGC, and was notable for paying special attention to some of the major risk factors – particularly airline pension plans – facing PBGC's balance sheet. In focusing on these large loss scenarios, scenarios which PIMS is admittedly insensitive towards, Elliott's model suggested that PBGC's financial condition was far worse than PIMS was predicting at the time, results which he used to inform the discussion of a series of reform proposals (Elliott 2005a), as well as a discussion of the structural forces within the PBGC economic model that limit the viability of potential reforms (2005b).¹⁷

¹⁷ We note that our current review is restricted to the PIMS model: we made no attempt to validate any other models, including that of Elliott.