Pricing Government Credit: A New Method for Determining Government Credit Risk Exposure

Brent W. Ambrose and Zhongyi Yuan

OVERVIEW

• The rapid growth of Fannie Mae, Freddie Mac, and other federal loan guarantee programs has prompted a debate on how to price government interventions in the credit markets.

• This study deploys a method for estimating the value of the government's loan guarantees that recognizes the interdependence of losses across programs.

• The estimation method reduces the total risk capital reserve that the government must hold to a level below the reserve required when all of the programs are considered in isolation; the extent of the reduction hinges on the strength of tail dependence among the program losses.

• Applying the method to the residential mortgage portfolio of Fannie Mae and Freddie Mac, the authors show that the strong tail dependence of the two programs permits only a modest reduction in the government's capital allocation and hence its guarantee fees.

In recent years, the pace of growth in government credit programs has accelerated markedly (U.S. Government Accountability Office 2016). Direct loans and loan guarantees by the federal government increased from $1.1 trillion in 2001 to almost $3 trillion by 2014 (Chart 1). Off-balance-sheet federal credit enhancements have reached still higher levels (Hanson, Scharfstein, and Sunderam 2016), with the combined value of Fannie Mae's and Freddie Mac's mortgage portfolio and mortgage-backed securities (MBS) guarantee exceeding $5.1 trillion in 2016.[1] One study (Lucas 2014b) puts the combined value of the federal credit programs and government-sponsored enterprises (Fannie Mae, Freddie Mac, the Federal Home Loan Banks, deposit insurance programs, and the Pension Benefit Guarantee Corporation) at more than $18 trillion in 2013. Underscoring the sheer magnitude of the federal government as a financial institution, the study notes that the federal government administers more than one hundred direct loan and guarantee programs—a figure that does not include programs such as Social Security or federal civilian and military pension benefits.

Not surprisingly, the rapid growth in government credit programs has prompted a debate on how to explicitly

Brent W. Ambrose is the Smeal Professor of Risk Management and Zhongyi Yuan an assistant professor of risk management at the Smeal College of Business, Pennsylvania State University. Email: bwa10@psu.edu; zuy11@psu.edu.

CHART 1
Federal Direct Loans and Loan Guarantees Outstanding

■ Loan guarantees ■ Direct loans
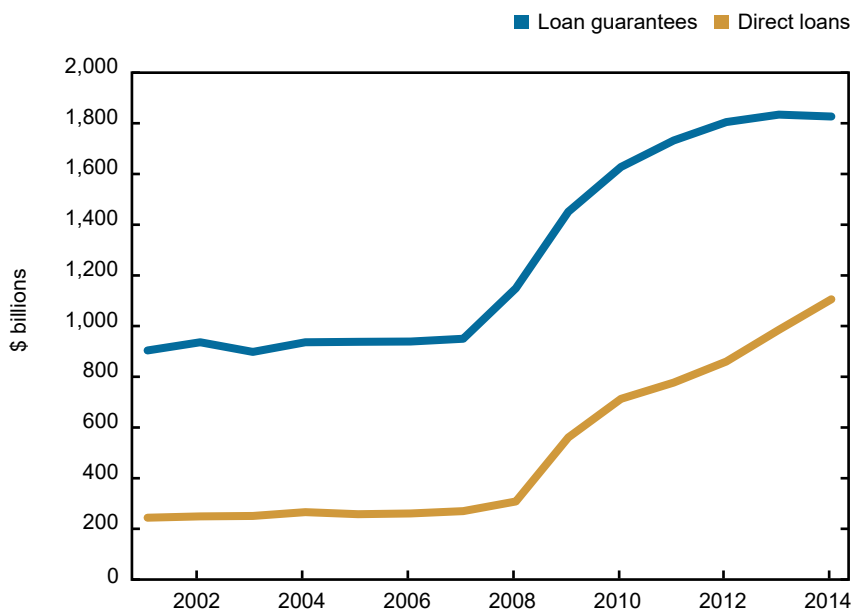
Source: U. S. Government Accountability Office (2016).

recognize—and possibly price—government interventions in the capital markets. At issue are the methods used to represent government loan and guarantee programs in the federal budget and balance sheet. Beginning in 1992, the Federal Credit Reform Act of 1990 (FCRA) mandated that the federal budget reflect the estimated net long-term cost (or subsidy cost) of credit programs in the year that the loan or guarantee was made.[2] Furthermore, these estimates were to be updated annually to reflect information on the actual and expected program performance. However, a growing literature contends that the FCRA underestimates the true costs associated with these programs and that a "fair value approach" offers a better metric for evaluating the costs that federal guarantees impose on taxpayers (see, for example, Congressional Budget Office [2012]; Lucas [2014a,b]; Hanson, Scharfstein, and Sunderam [2016]; Kelly, Lustig, and Van Nieuwerburgh [2016]; Lucas, [2012, 2016]; and Heaton, Lucas, and McDonald [2010]).[3]

In general, the distinction between the FCRA and fair value methods rests on the choice of the discount rate used in determining the present values of the programs. For example, under FCRA rules, future expected cash flows are discounted at the prevailing Treasury rates to reflect the capital cost ascribed to the federal government. In contrast, the fair value approach relies on market-based discount rates that reflect the risk premium associated with the programs' cash flows.[4]

While the choice of discount rate is undeniably important, generally missing from the debate is the recognition that a complete valuation method of the government's risk exposure should include the potential for loss dependence across various government credit programs.

One exception is the recent study by Hanson, Scharfstein, and Sunderam (2016), which develops a novel framework for analyzing government programs that explicitly recognizes the government's social and fiscal risk management objectives. One of the key insights from their model is that when the government manages a portfolio of programs (for example, direct loans and guarantee programs), each program's social and fiscal risks depend on how that program co-varies with others. Building on this insight, we apply a method that explicitly recognizes the interdependence associated with the likelihood of joint risk across loan and guarantee programs in order to estimate more precisely the government's risk exposure to each program through the program's required capital allocation.

The federal government's loan and guarantee portfolio contains a number of programs that share similar fundamental drivers of potential losses. The government loan guarantee programs in housing are an example of how default losses across programs may co-vary. The federal government operates a number of direct and indirect programs that provide insurance to mortgage investors in the event of borrower default.[5] With respect to mortgage credit risk, it is well known that the three primary factors driving borrowers' decisions to default are uncertainty over collateral value relative to mortgage value, future interest rates, and income.[6] Thus, since mortgage default rates rise during periods of economic stress when house prices fall and unemployment rises, losses across the loan products of the Federal Housing Administration (FHA) and the government-sponsored enterprises (GSEs) are expected to have strong tail dependence. As a result, one would expect that losses in the government's mortgage programs will co-vary.

In addition, our theoretical analysis demonstrates that to correctly assess the economic value of a government insurance program, one should consider not only the expected value of the program's losses but also a risk margin that reflects possible tail losses over and above the expected loss as well as uncertainties in the estimate of the expected loss. Including the risk margin allows us to value the federal programs in a manner consistent with insurance firm regulators' valuation methods for insurance liabilities. Put in the context of loan guarantee pricing, our price of the guarantee is the present value of the sum of the expected loan loss and a risk margin. Various methods for determining the risk margin have been discussed in the literature, and our method (with or without recognizing the interdependence of losses across programs) builds on the cost of capital approach (see, for example, Sandström [2016] for related discussions).

Lastly, we present an empirical analysis using the residential loan portfolios of Fannie Mae and Freddie Mac to illustrate the concept of tail dependence in pricing government credit guarantees. Assuming that Fannie Mae and Freddie Mac represent two insurance "lines" of the federal government, we demonstrate that they exhibit strong tail dependence. The tail dependence leads to values for Fannie Mae's and Freddie Mac's credit enhancement fees (also known as the guarantee fees) that are higher than those that would obtain if the two lines were independent. We also show that the current practice of setting the guarantee fees in isolation overprices the government's overall risk exposure to borrower default.

Our paper proceeds as follows: In Section 1, we present a model and simulation illustrating how capital allocation and program pricing can change as a result of recognizing the dependence between programs. Section 2 illustrates these concepts with an empirical analysis using mortgage data from Fannie Mae and Freddie Mac. Finally, Section 3 concludes the analysis with a look at some potential extensions of our framework.

1. Pricing: Theory and Simulation

1.1 Model

We develop a model for a portfolio of simple loan guarantee programs run by the government and study the government's risk exposure from the pricing perspective (that is, by pricing the guarantees). Our model can be extended to the reporting of economic value and the estimation of subsidy costs. For example, the subsidy cost of running a loan guarantee program can be estimated by the price of the guarantee minus the actual premiums collected.

Mimicking current government accounting conventions, we first price each loan guarantee as a stand-alone program. Then we consider the loan guarantees as a government portfolio and price the guarantees using a top-down capital allocation approach—first without, and then with, the default dependence recognized. As we will see, the price crucially depends on the diversification effects among the programs during normal times and the contagion effects during stressed times.

We begin by supposing that the government has a portfolio of n loan guarantee programs. Each of the guarantee programs may cover a different segment of the economy (for example, student loans, small business loans, or the housing market through mortgage insurance). For the sake of simplicity, we focus on a risk management period of one year and assume that the i^{th} program ($i = 1, \dots, n$) has a nonnegative random loss of X_i at the end of the period. We model the programs such that the government pays X_i to the investors that originate loans in the i^{th} program. Thus, the total payment by the government (that is, the guarantee payment to loan originators) is

$$S = X_1 + \cdots + X_n. \tag{1}$$

Note that these guarantees are not traded in the financial market, cannot be replicated by existing assets, and hence do not form a complete market to allow for risk-neutral pricing. As a result, we employ an actuarial pricing approach. Actuarially fair pricing of insurance policies is extensively discussed in the actuarial and insurance literature. For example, the premium on an insurance policy issued by a limited-liability insurance company is typically the present value of the sum of the expected claim cost, administrative expenses, and the cost of the capital needed to support the policy, minus the value of the put option of default (see, for example, Zanjani [2002] and Ibragimov, Jaffee, and Walden [2010]). A critical difference in the case of government guarantees is that default is no longer an option. Therefore, the price of the government loan guarantee mainly contains two components: the expected loss and the cost of the capital needed to support the guarantee.[7] This is equivalent to pricing the guarantee by using a cost of capital approach to determine the risk margin and then adding it to the expected loss.

We start with the case in which each program is priced as a stand-alone program. In order to limit the taxpayers' exposure, the government requires each program to hold some amount of risk capital, which is determined through the tail value at risk (TVaR) of the program loss at the z percent level. The TVaR at level z percent for program i is defined as

$$E\big[X_i \big| X_i \geq VaR_{z\%}[X_i]\big], \tag{2}$$

where $VaR_{z\%}[X_i] = \inf\{s \in (-\infty, \infty) : \Pr(X_i \leq s) \geq z\%\}$ is the value at risk (VaR) of X_i at level z percent.[8] Unlike the VaR risk measure, which concerns only the frequency of a capital

shortfall, TVaR, a more conservative risk measure, also concerns the amount of the shortfall when it occurs.[9] The choice of risk measure in the following demonstration is TVaR at the 99 percent level, but our analysis is also valid for other choices. The risk capital needed for assuming program i's risk is then given by[10]

$$TVaR_{99\%}[X_i] - E[X_i]. \tag{3}$$

If the annual risk-free interest rate is r_f and the annual cost of capital rate is δ, both positive, then the premium charged for the guarantee on program i is

$$prem_i = \frac{E[X_i] + \delta(TVaR_{99\%}[X_i] - E[X_i])}{1 + r_f}. \tag{4}$$

Here, the second term in the numerator represents the time-1 cost of holding the risk capital. The entire numerator is the time-1 "value" of the guarantee and is discounted at the risk-free rate for one period to get the time-0 premium. The risk-free rate used here simply reflects the time value of money and should not be confused with the debate over the choice of discount rate used in financial reporting. Of course, an important factor in the premium calculation is the cost of capital rate δ. Since the government's cost of capital rate is close to the risk-free rate, we shall use $\delta = r_f$ hereafter.[11]

Next, we suppose that the guarantee programs do not need to hold risk capital individually; instead, the government holds a certain amount of risk capital for the programs as a group, also based on the 99 percent level TVaR. Therefore, the total amount of capital is

$$E\big[S|S \geq VaR_{99\%}[S]\big] - E[S]. \tag{5}$$

This construct allows us to first use a top-down approach to allocate risk capital to individual programs and then to price the guarantees.

If the risk capital is allocated to each program based on its marginal contribution of risk (that is, based on Euler's principle), then the amount of capital allocated to program i is[12]

$$CA_i = E\big[X_i|S \geq VaR_{99\%}[S]\big] - E[X_i], \tag{6}$$

which, in spirit, is consistent with the marginal expected shortfall used in the literature for measuring the contribution of systemic risk from an individual program (see, for example, Acharya, Pedersen, and Philippon [2017]).[13]

In summary, our cost of capital approach produces a premium of

$$prem_i = \frac{E[X_i] + \delta \times \big(E[X_i|S \geq VaR_{99\%}[S]\big] - E[X_i]\big)}{1 + r_f} \tag{7}$$

for program i. It is clear from the pricing formula that the dependence—and, in particular, the tail dependence—across the program losses should have an impact on the capital allocation and hence on the guarantee premium. We now briefly discuss the concept of tail dependence and then use a few examples to demonstrate its impact.

Occurrences of Simultaneous Extreme Losses across Programs
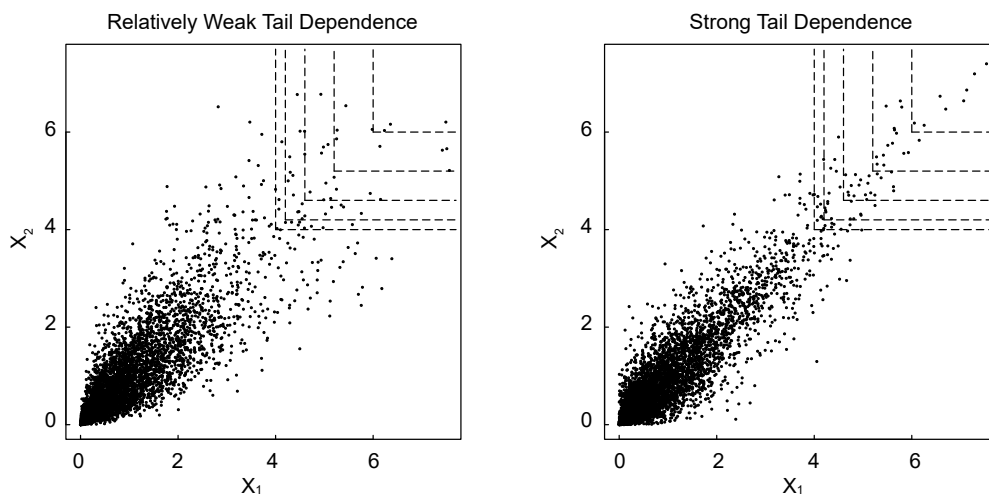
Source: Authors' calculations.

Notes: The chart shows 5,000 samples of two losses, X_1 and X_2, both following *Exp(mean = 1)*. The two losses have a Gaussian copula with parameter $\rho = 0.87$ (left panel) or a Gumbel copula with parameter $\theta = 3$ (right panel).

1.2 Tail Dependence

Tail dependence captures the tendency for losses to be simultaneous and extremely large (that is, appearing in the tails of their respective distributions). While the government can use premiums collected from, or capital held for, other programs to cover an extreme loss, it may run short of capital when all programs experience simultaneous extreme losses. The tendency of programs to have simultaneous extreme losses is referred to as strong tail dependence, and it is important to model the likelihood of such scenarios correctly. We use an example to convey the difference between weak and strong tail dependence.

As an illustration, Chart 2 compares the number of occurrences of simultaneous extreme losses assuming weak tail dependence (left panel) and strong tail dependence (right panel).[14] For each dot in the chart, the coordinates represent values of the two risks, X_1 and X_2, both assumed to follow an exponential distribution with mean 1. The dots in the upper-right corner represent scenarios of simultaneous large losses. As we move the rectangle toward the upper right (that is, as we focus on larger losses), we observe more such scenarios for the strong tail dependence case than for the weak tail dependence case.

Examples

To illustrate the capital allocation and tail dependence concepts, we assume that the government sponsors three loan guarantee programs ($n = 3$) and examine the capital allocation for

each program under four scenarios that highlight various assumptions about the distribution of losses. The four scenarios are as follows:

1. Price the programs as stand-alone entities using each line's capital reserve (the base case).

2. Assume independence across programs.

3. Assume losses are dependent through a Gaussian copula.

4. Assume losses are dependent through a Gumbel copula.

For scenarios 2–4, we use the programs' marginal risk contributions to allocate risk capital and set premiums.

Under all four scenarios, we assume that risk capital is set at the 99 percent level TVaR, and that both the risk-free rate and the cost of capital rate are 2 percent per year. According to equation (4), for scenario 1, we assume that the price for the program i^{th} is

$$prem_i = \frac{E[X_i] + 2\% \times (TVaR_{99\%}[X_i] - E[X_i])}{1.02}, \qquad (8)$$

whereas for scenarios 2 through 4, the price for the program is

$$prem_i = \frac{E[X_i] + 2\% \times CA_i}{1.02}. \qquad (9)$$

In all cases, we assume that program losses follow an exponential distribution, which has a "light" tail and implies that the probability of observing a catastrophic loan loss is relatively small. For the stand-alone pricing and independent loss scenarios, we assume that the program losses (X_1, X_2, and X_3) follow $Exp(mean = 1)$ and are mutually independent. In contrast, to introduce dependence for scenarios 3 and 4, we assume that the losses follow $Exp(mean = 1)$ and have either a Gaussian copula (weak tail dependence) or a Gumbel copula (strong tail dependence).[15]

The top panel of Table 1 reports the simulated capital and premium for each program and the total TVaR. The stand-alone scenario best matches current FCRA methods whereby each program is priced independently. Since we assume that the programs have the same loss distribution, the premiums associated with each program are equal. Furthermore, since the government does not consider potential diversification in default losses across programs, the government's required capital reserve (total TVaR) to limit losses at the 99 percent level is large.

The remaining three scenarios in the top panel consider the pricing and capital allocations assuming that premiums are set based on the programs' marginal risk contributions (MRCs) to the overall portfolio. In the case of independent losses, we assume that the government can shift capital from one program to another in the event of default, but we maintain the assumption of independence across programs in determining the capital allocation. As a result, the premiums charged for each program are lower than in the stand-alone case. In addition, the overall required capital reserve (TVaR) in this scenario is lower than that in the other scenarios as a result of the government's ability to shift capital across programs.

Table 1
Simulated Premiums and Capital Allocation

Panel A: Losses Follow an Exponential Distribution with Equal Means

	Program 1	Program 2	Program 3	Total TVaR
Stand-alone pricing	1.071	1.071	1.071	16.816
Marginal risk contribution (MRC) pricing				
Independent losses	1.024	1.024	1.024	9.638
Gaussian dependence	1.068	1.052	1.025	13.425
Gumbel dependence	1.062	1.062	1.062	15.465

Panel B: Losses Follow an Exponential Distribution with Unequal Means

	Program 1	Program 2	Program 3	Total TVaR
Stand-alone pricing	1.071	2.141	3.212	33.631
Marginal risk contribution (MRC) pricing				
Independent losses	0.990	2.023	3.168	21.235
Gaussian dependence	1.058	2.065	3.138	25.406
Gumbel dependence	1.058	2.122	3.200	31.115

Source: Authors' calculations.

Notes: The table shows the premiums and total risk capital for three government-sponsored loan guarantee programs under four scenarios. Under stand-alone pricing, the programs are priced independently; under marginal risk contribution (MRC) pricing, program premiums and capital allocation are determined on the basis of each program's relative contribution to the overall portfolio risk. Tail dependence is introduced through the assumption that losses have a Gaussian copula (weak tail dependence) or a Gumbel copula (strong tail dependence).

Next, we allow for correlated default risks across programs assuming weak (Gaussian) or strong (Gumbel) tail dependence.[16] In the case of weak tail dependence, we see that premiums vary across programs, with program 3 having the lowest premium owing to its assumed low correlation with other programs. However, the government's overall capital position (TVaR) is higher than in the independent losses scenario because of the recognition of potential correlated default losses. Finally, owing to the sub-additivity of TVaR as a risk measure, the total TVaR for scenarios 3 and 4 can never exceed that for the stand-alone scenario. Nevertheless, we note that the total TVaR for the strong tail dependence case approaches that of the stand-alone case.

We next modify the assumption regarding the distribution of losses to allow for different expected losses across programs. To accomplish this, we assume that X_1, X_2, and X_3 are distributed as $Exp(mean = 1)$, $Exp(mean = 2)$, and $Exp(mean = 3)$, respectively. The assumptions on the dependence among X_1, X_2, and X_3 are the same as in the previous example. The bottom panel of Table 1 reports the simulated guarantee premiums and total risk capital. Here, we note that the higher expected default losses for programs 2 and 3 result in higher premiums. However, the total TVaR follows the same pattern, with the government able to lower total

capital costs by recognizing the diversification effect. We again see that the three programs are priced differently. However, the strong tail dependence causes the total cost to the government (TVaR) to be almost the same as in the stand-alone pricing case (when the lines are priced independently). As a result, recognizing the potential for program losses to have strong tail dependence pushes the government's capital allocation pricing back toward the level assumed when the lines are priced in isolation.

In the appendix, we repeat the simulation analysis but modify the assumption regarding the distribution of losses such that X_1, X_2, and X_3 follow a Pareto distribution ($Pareto(2, 1)$). This modification recognizes the possibility that program losses may have "heavy" tails. We further explore the effect of strong tail dependence by maintaining the assumption that the program losses follow a Gumbel copula. The results continue to hold, with the diversification effect dissipating with the introduction of strong tail dependence.

2. Empirical Analysis

2.1 Data

We now turn to an empirical analysis of the residential loan portfolios of Fannie Mae and Freddie Mac to illustrate how the concept of tail dependence might be applied in pricing government credit guarantees. Fannie Mae and Freddie Mac are government-sponsored enterprises that were placed into conservatorship by the U.S. Treasury Department in September 2008 during the financial crisis. Given the nature of the GSEs' business activity and their relationship with the federal government, our analysis views Fannie Mae and Freddie Mac as representing two different insurance programs of the federal government. Thus, we test the tail dependence between the two programs and price the guarantees.

We use the GSE public data sets comprising thirty-year fixed-rate mortgages purchased or securitized by the GSEs and made available through the Federal Housing Finance Agency (FHFA). The data contain performance information on GSE loans originated from 2000 to 2016. The available loan characteristics include origination date, original unpaid principal, loan-to-value (LTV) ratio, zero-balance date, unpaid balance prior to foreclosure date, zip code, and source (Fannie Mae loan or Freddie Mac loan).[17]

We group the loans according to their foreclosure dates (month/year) and three-digit zip codes. For analysis of loan loss dependence, we use monthly three-digit-zip-code-level losses of the two programs projected on the foreclosure dates. The projected losses are equal to the product of the unpaid balance prior to the foreclosure date and the projected loan loss severities. We use the GSEs' loss severities reported in Table 9 of Goodman and Zhu (2015) to project the severities based on origination year and LTV.[18]

2.2 Analysis of Tail Dependence

In this subsection, we analyze the tail dependence between the losses of the two programs at the zip code level. As a preliminary approach to the issue, we create a scatterplot diagram of the relationship between the two programs' annual losses, with loans further

CHART 3
Annual Zip-Code-Level Losses of Fannie Mae and Freddie Mac
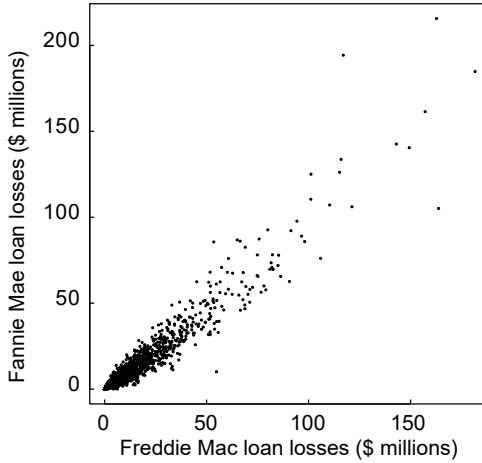
Source: Authors' calculations.

grouped by their years of foreclosure (Chart 3). The tendency for the GSEs' losses to be large simultaneously—evident in the chart—offers a hint of tail dependence.

Mathematically, we focus on the tail dependence coefficient defined as follows: for two losses Y_1 and Y_2 following distribution functions F_1 and F_2, respectively, their tail dependence coefficient is given by

$$\lambda = \lim_{u \to 0} \frac{\Pr(F_1(Y_1) > 1 - u, F_2(Y_2) > 1 - u)}{u}, \tag{10}$$

provided that the limit exists. A positive value of λ means strong tail dependence between the two losses. Specifically, we say that the two losses are tail dependent if $\lambda > 0$ and tail independent if $\lambda = 0$. We test the null hypothesis that the losses from the two lines are tail independent. Following the methods introduced by Hüsler and Li (2009), we assume that the data are summarized by n observations of loss vectors $(Y_{11}, Y_{21}), \ldots, (Y_{1n}, Y_{2n})$, with (Y_{1i}, Y_{2i}) denoting the vectors of monthly losses. The test requires an equal split of the data into two subgroups, each containing $m = n/2$ observations of loss vectors denoted by $\{(Y_{11}, Y_{21}), \ldots, (Y_{1m}, Y_{2m})\}$ and $\{(Y_{1(m+1)}, Y_{2(m+1)}), \ldots, (Y_{1n}, Y_{2n})\}$.[19]

For the i^{th} member of the first subgroup, R_{1i} denotes the rank of Y_{1i} in $\{Y_{1(m+1)}, \ldots, Y_{1n}\}$ and R_{2i} the rank of Y_{2i} in $\{Y_{2(m+1)}, \ldots, Y_{2n}\}$ (a larger value means a higher rank). The two test statistics (the integral test statistic and the supreme test statistic) introduced by Hüsler and Li (2009) are given, respectively, by
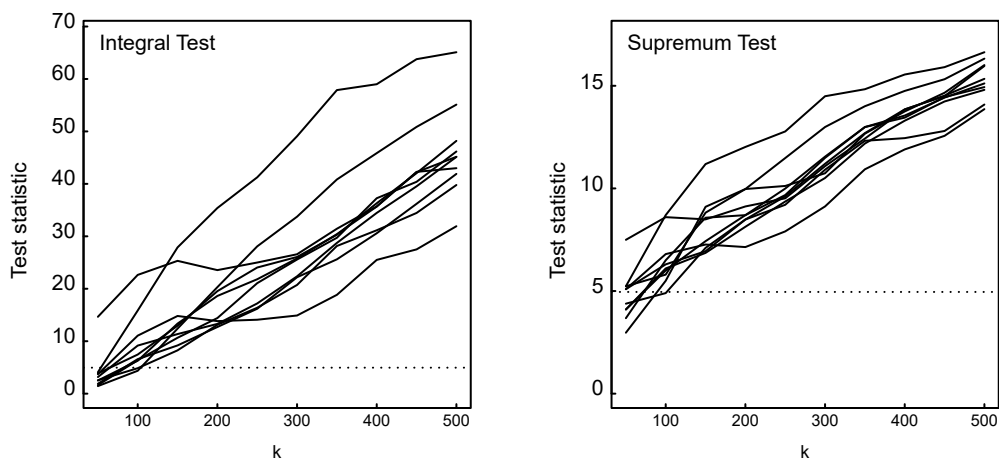
$$T_I = \int_0^1 \int_0^1 D(x, y)^2 \, dx \, dy \tag{11}$$

and

$$T_S = \sup_{x, y \in [0,1]} |D(x, y)|, \tag{12}$$

Testing the Null Hypothesis That the Losses of Fannie Mae and Freddie Mac
Are Tail Independent

Source: Authors' calculations.

Note: The dotted lines denote the levels of critical value.

where

$$D(x,y) = \sqrt{k}\left(\frac{1}{k}\sum_{i=1}^{m}\mathbb{I}\{R_{1i} > m + 1 - kx \ OR \ R_{2i} > m + 1 - ky\} - (x + y)\right). \qquad (13)$$

Here $\mathbb{I}\{\cdot\}$ denotes the indicator function, and k satisfies $k = k(n) \to \infty$ and $k(n)/n \to 0$
as $n \to \infty$. The choice of k for tail/extreme value analysis is subjective, and therefore, as a
robustness check, we perform the tests for a range of values of k, from 50 to 500.

We use numerical integration to approximate the integral test statistic and use the
maximum over a 100×100 grid to approximate the supreme test statistic. In addition, as a
robustness check, we randomly shuffle and split the data, and then recalculate the test statis-
tics. We repeat this ten times and show the test statistics and critical values in Chart 4.

The critical value at the 5 percent significance level is known to be 6.237 for the integral test
and 4.956 for the supreme test. Chart 4 shows that for the majority of cases (89 percent for the
integral test and 91 percent for the supreme test), the test statistics are greater than the critical value,
suggesting that it is reasonable to reject the null hypothesis of tail independence at the 5 percent
significance level. Thus, we conclude that the two programs have strong tail dependence.

2.3 Pricing

We now illustrate the impact of tail dependence on the pricing of the GSEs' guarantee fee for
thirty-year fixed-rate mortgages. We first fit the distributions of the annual losses from the two

Chart 5
Distribution of Annual Zip-Code-Level Fannie Mae Loan Losses
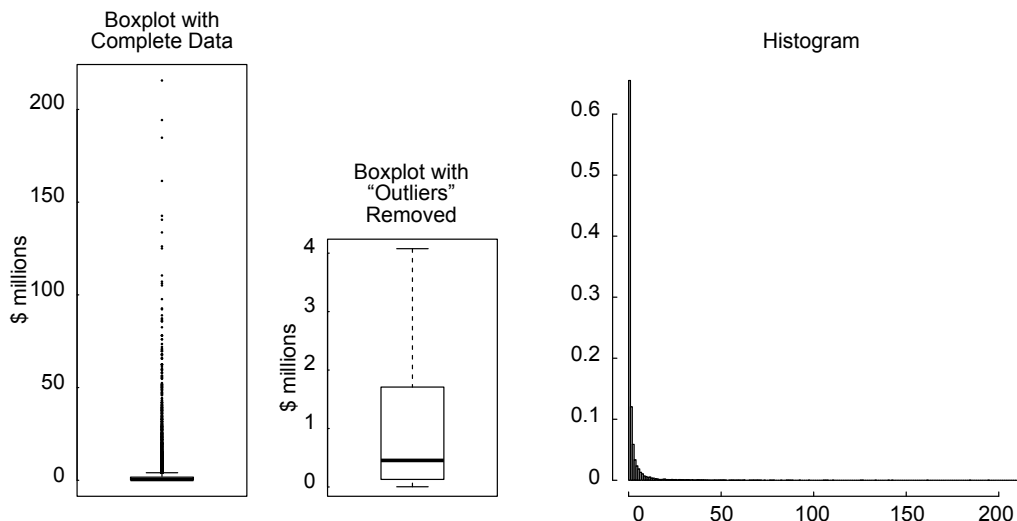
Source: Authors' calculations.

programs. Beginning with Fannie Mae, we chart the program's loan losses in two boxplots and a histogram, which together show that the losses are likely heavy tailed (Chart 5). We consider exponential, Gamma, log-normal, and Pareto distributions as candidates for the distribution of Fannie Mae's losses. The first two columns in Table 2 list the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) values for these distributions. According to the AIC and BIC criteria, the best fit among these candidates is log-normal. While the pp and qq plots in Chart 6 confirm that log-normal fits better than the other candidates and fits well for data in the normal range, it does not fit the tail well. Therefore, motivated by the Pickands-Balkema-de Haan theorem, which states that losses that exceed a high threshold roughly follow a generalized Pareto distribution, we use a two-component spliced distribution with the log-normal distribution fitting the data in the normal range and a generalized Pareto distribution in the extreme range. As shown by the qq plot in Chart 6, this leads to a much-improved fit. The spliced distribution density function is given by

$$f_{FN}(x) = \begin{cases} g_{FN}(x), & 0 < x < x_{FN}, \\ (1 - p_{FN}) \times h_{FN}(x - x_{FN}), & x \geq x_{FN}, \end{cases} \qquad (14)$$

where $x_{FN} = 10.658$, $p_{FN} = \int_0^{x_{FN}} g_{FN}(x)dx = 0.963$, $g_{FN}(\cdot)$ is the density function of a log-normal distribution with parameters $\mu_{FN} = -0.772$ and $\sigma_{FN} = 1.751$, and $h_{FN}(\cdot)$ is the density function of a generalized Pareto distribution with shape parameter $\xi_{FN} = 0.404$ and scale parameter $\theta_{FN} = 10.616$.

Similarly, Charts 7 and 8 show the box plots and histogram of Freddie Mac loan losses and the pp and qq plots of the corresponding fitted distributions, respectively. The values of the

TABLE 2
Model Selection

Distribution	Fannie Mae		Freddie Mac	
	AIC	BIC	AIC	BIC
Exponential	48, 154	48, 161	53, 966	53, 973
Gamma	37, 613	37, 628	45, 281	45, 296
Log-normal	32, 300	32, 315	40, 870	40, 885
Pareto	32, 777	32, 791	41, 313	41, 328

Note: The table shows the AIC (Akaike information criterion) and BIC (Bayesian information criterion) values for fitting an exponential, a gamma, a log-normal, and a Pareto distribution to the annual zip-code-level Fannie Mae and Freddie Mac loan losses.
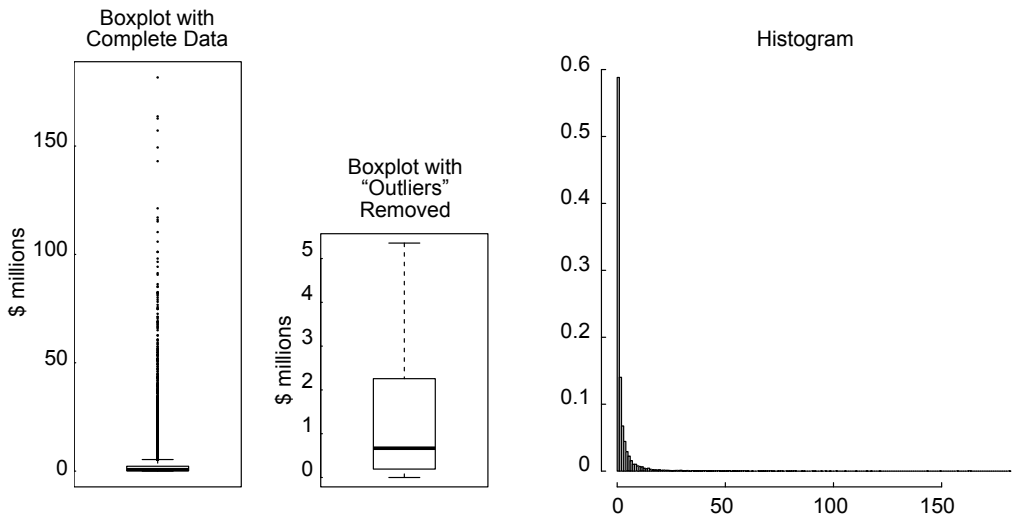
CHART 6
Comparisons of the Fitted and Empirical Distributions of Annual Zip-Code-Level Fannie Mae Loan Losses Using pp and qq Plots
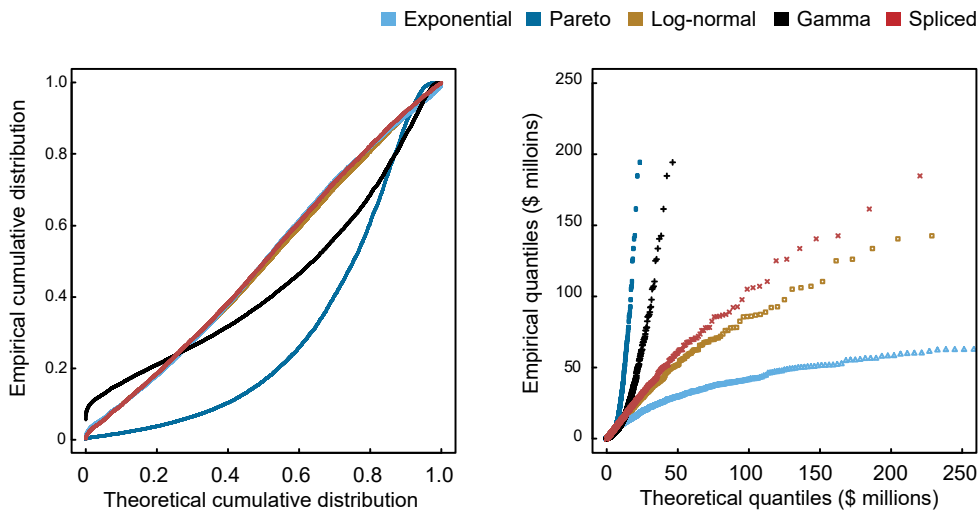
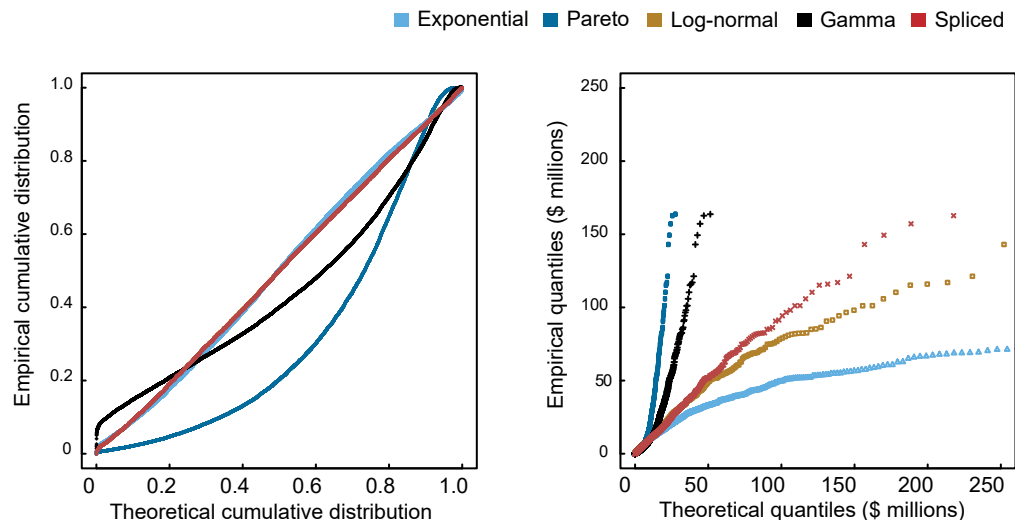Source: Authors' calculations.

CHART 7
Distribution of Annual Zip-Code-Level Freddie Mac Loan Losses

Source: Authors' calculations.

CHART 8
Comparisons of the Fitted and Empirical Distributions of Annual Zip-Code-Level
Freddie Mac Loan Losses Using pp and qq Plots

Source: Authors' calculations.

TABLE 3
Premiums and Required Capital
Billions of Dollars

Distribution	Fannie Mae	Freddie Mac	Total TVaR
Stand-alone	2.94	3.77	115.6
Independence	2.60	3.47	83.0
Fitted Gumbel	2.75	3.63	98.3

Note: The table shows the annual premiums and required capital for insuring against Fannie Mae loan and Freddie Mac loan losses nationwide. They are calculated based on the fitted marginal distributions of the loan losses under each of the following scenarios: 1) the two lines are priced as stand-alone lines; 2) the two lines are independent; 3) the two lines are dependent through the fitted Gumbel copula.

AIC and BIC for candidate distributions are listed in the third and fourth columns of Table 2. Our final choice of fitted distribution for Freddie Mac loan losses is also a two-component spliced distribution, with density given by

$$f_{FR}(x) = \begin{cases} g_{FR}(x), & 0 < x < x_{FR}, \\ (1 - p_{FR}) \times h_{FR}(x - x_{FR}), & x \geq x_{FR}, \end{cases} \tag{15}$$

where $x_{FR} = 12.131$, $p_{FR} = \int_0^{x_{FR}} g_{FR}(x)dx = 0.949$, $g_{FR}(\cdot)$ is the density function of a log-normal distribution with parameters $\mu_{FR} = -0.403$ and $\sigma_{FR} = 1.774$, and $h_{FR}(\cdot)$ is the density function of a generalized Pareto distribution with shape parameter $\xi_{FR} = 0.310$ and scale parameter $\theta_{FR} = 12.640$. Meanwhile, prompted by our conclusion that there is tail dependence between the two lines, we fit a Gumbel copula to the losses.[20]

We then use equation (9) to price the guarantees on the losses for thirty-year fixed-rate mortgages. While the FHFA charges guarantee fees that are calculated in terms of basis points, we report the total premium required. This choice reflects our holistic approach: the aim of our analysis is to summarize program losses rather to analyze program cash flows.

To estimate loan losses nationwide, we simply assume that they are equal to the number of zip codes producing losses times zip-code-level losses. Based on the fitted marginal distributions, the expected annual three-digit-zip-code-level losses for Fannie Mae and Freddie Mac thirty-year fixed-rate mortgages are $2.15 million and $2.99 million, respectively. Fannie Mae and Freddie Mac losses on thirty-year fixed-rate mortgages are calculated from 915 and 891 three-digit zip codes, respectively. Therefore, the national expected annual losses for Fannie Mae are 915 × 2.15 million, or $1.97 billion; those for Freddie Mac are 891 × $2.99 million, or $2.66 billion.

The allocations of capital, according to equation (6), are estimated as the average program loss over the cases where the total losses of the two programs exceed the 99 percent quantile minus the expected program loss. To show the impact of tail dependence on the value of loan guarantees, we first estimate the allocations of capital assuming that the losses from the two programs are independent, and then report the losses estimated with tail dependence fitted using a Gumbel copula. The premiums and the TVaR of the program losses are reported in Table 3 (above). We observe that tail dependence leads to higher premiums and higher total

required risk capital, a finding that is consistent with the simulations in Section 1. Specifically, the annual total premium for insuring against Fannie Mae loan losses nationwide increases by about 5.8 percent because of tail dependence. Similarly, the total premium for Freddie Mac to insure its losses increases by about 4.6 percent. In addition, we find that the total required risk capital for the GSEs increases by 18.4 percent. Because of data limitations, our analysis focuses only on the thirty-year fixed-rate mortgage product line, and thus, our risk capital estimates are not directly comparable to the capital reserves reported by the GSEs in their annual reports.

Contrary to the conclusion of the FHFA's recent Fannie Mae and Freddie Mac Guarantee Fee Review, which found "no compelling economic reason to change the general level of fees,"[21] we observe a substantial difference between the premiums for the two programs. The disparity arises from the difference in the mean expected losses of the enterprises. Our estimates indicate that Freddie Mac's premium should be 32.0 percent higher than the premium for Fannie Mae. In addition, the programs' marginal risk contributions, measured as a percentage of risk capital allocations, are estimated to be 44.6 percent for Fannie Mae as against 55.4 percent for Freddie Mac.

3. Conclusion

A growing debate centers on how best to recognize and price government interventions in the capital markets. We apply a method for estimating the government's exposure to credit risk through its loan and guarantee programs and for estimating the value of the guarantee. Our method builds on the insight in Hanson, Scharfstein, and Sunderam (2016) that the risks associated with government programs co-vary, and it explicitly recognizes the interdependence among the program losses. We demonstrate that the current practice of reporting only the economic value implies an underestimation of the value of the government guarantee. However, we also show that pricing the programs after properly accounting for their risk margin on a stand-alone basis will overestimate the guarantee's value. We propose that a "correct" method should base each program's risk margin on the interaction of all the programs, which determines their marginal risk contributions to the government's program "portfolio." Our method also suggests that such an accounting method would lead to a reduction in the overall capital reserve required to insulate taxpayers from losses. Moreover, we point out that the extent of the reduction hinges on the strength of (tail) dependence among the program losses. By extension, the government's overall risk exposure will hinge on the dependence of losses across programs.

The application of our approach relies on two important building blocks: the expected value and the dependence of program losses. Therefore, future work should focus on identifying the appropriate discount rate for determining the present value of expected losses and on estimating the tail dependence.

APPENDIX: SIMULATION RESULTS ASSUMING THAT LOSSES FOLLOW A PARETO DISTRIBUTION

Example 1:

We modify the assumption regarding the distribution of losses in the stand-alone and independent cases such that X_1, X_2, and X_3 follow a Pareto distribution (*Pareto* (2, 1)). This modification provides for "heavy" tails in the expected losses.[22] To incorporate tail dependence, we assume that the losses X_1, X_2, and X_3 have a Gaussian copula with correlation matrix

$$\begin{bmatrix} 1 & 0.9 & 0.5 \\ 0.9 & 1 & 0.1 \\ 0.5 & 0.1 & 1 \end{bmatrix}$$

	Program 1	Program 2	Program 3	Total TVaR
Scenario 1	1.333	1.333	1.333	57.000
Scenario 2	1.197	1.196	1.193	35.874
Scenario 3	1.301	1.265	1.202	45.135

Example 2:

For the second version of this example, we modify the assumption regarding tail dependence such that the losses X_1, X_2, and X_3 have a Gumbel copula

$$C(u_1, u_2, u_3) = \exp\{-[(-\log u_1)^\theta + (-\log u_2)^\theta + (-\log u_3)^\theta]^{1/\theta}\}$$

with $\theta = 5$. Again, this assumption produces strong tail dependence.

	Program 1	Program 2	Program 3	Total TVaR
Scenario 1	1.333	1.333	1.333	57.000
Scenario 2	1.197	1.196	1.193	35.874
Scenario 3	1.292	1.293	1.292	50.706

Appendix (Continued)

Example 3:

Our final version of this example assumes that in the stand-alone and independent cases, the X_1, X_2, and X_3 follow *Pareto*(1.5, 1), *Pareto*(2.0, 1), and *Pareto*(3.0, 1) respectively, and hence, the expected losses are 2, 1, and 0.5, respectively. Again, for the tail dependence case, the losses X_1, X_2, and X_3 have a Gumbel copula $C(u_1, u_2, u_3) = \exp\{-[(-\log u_1)^\theta + (-\log u_2)^\theta + (-\log u_3)^\theta]^{1/\theta}\}$ with $\theta = 5$.

	Program 1	Program 2	Program 3	Total TVaR
Scenario 1	3.169	1.333	0.597	88.595
Scenario 2	3.109	1.139	0.496	70.462
Scenario 3	3.161	1.271	0.569	83.518

NOTES

[1] Data on Freddie Mac's position are available at http://www.freddiemac.com/investors/pdffiles/investorpresentation .pdf. Data on Fannie Mae's position are available at http://www.fanniemae.com/resources/file/ir/pdf/monthly -summary/013117.pdf.

[2] See U. S. Government Accountability Office (2016) for a discussion of the FCRA's history and its implementation.

[3] Others have proposed using an option-pricing framework (Lucas and McDonald 2006).

[4] The differences between the FCRA and fair value methods are nontrivial. Lucas (2014a) provides an illustration of the difference using the Tennessee Valley Authority (TVA) as an example and notes that between 2008 and 2012, the TVA's "understatement of capital costs totaled about $3 billion." This difference can be ascribed to the value of the implicit guarantee backing TVA debt. Thus, the concepts of fair value reporting are rooted in attempts to quantify the value of the government's implicit guarantee that backs various government-sponsored enterprises (GSEs). For example, Ambrose and Warga (2002) and Passmore (2005) quantified the value of the implicit guarantee backing the housing GSEs (Fannie Mae and Freddie Mac). Given the public-private nature of the GSEs, the purpose of these measures is to ascertain whether the value of the implicit guarantee is being captured by the shareholders or, alternatively, passed through to consumers.

[5] For example, the FHA administers the Mutual Mortgage Insurance (MMI) program to insure mortgage default risk associated with lower-income and first-time homebuyers, while the government-sponsored enterprises Fannie Mae and Freddie Mac, which are currently in federal conservatorship, provide credit enhancements for prime or conforming mortgages.

[6] See, for example, IFE (2007) for a comprehensive discussion of the literature on mortgage default.

[7] For simplicity, our analysis ignores the administrative expenses and assumes without loss of generality that the loan guarantee reimburses the entire loan loss.

[8] Different names have been used in the literature for TVaR, such as expected shortfall, conditional tail expectation, and tail conditional expectation, although they could differ slightly for discrete risks.

[9] We note that $TVaR_{z\%}[X] \geq VaR_{z\%}[X]$ for any value of z percent and any risk variable X. The majority of European insurance companies, regulated by the Solvency II framework, calibrate their risk capital based on the VaR risk measure at the 99.5 percent level, but Switzerland is an exception and uses 99 percent level TVaR. Moreover, the Life and Health Actuarial Task Force of the U.S. National Association of Insurance Commissioners (NAIC) and the International Actuarial Association have endorsed the use of TVaR for setting the capital requirement. See Geneva Association, Modernising Insurance Solvency Regimes—Key Features of Selected Markets, https://www .genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public/20160830_solvencyregimes _complete_0.pdf, and NAIC, "Consultation Paper on Regulatory Capital Requirements and Overarching Accounting/ Valuation Issues for the Solvency Modernization Initiative."

[10] Note that the exact amount of risk capital should be $TVaR_{99\%}[X_i] - prem_i$, but $\text{TVaR}_{99\%}[X_i] - E[X_i]$ is widely used in practice as an approximation.

[11] Whether the government's cost of capital rate (δ) should be the risk-free rate is still an open question. However, since the government borrows at the risk-free rate and if the risk capital is allowed to be held in the form of Treasury bonds, which earn the risk-free rate, then the cost of capital largely reduces to agency costs and monitoring costs. Other factors, such as political constraints, may affect the government's real borrowing cost, particularly as these considerations often involve federal budget discussions and have implications for the federal debt limit. In contrast, for insurance companies, the Swiss Solvency Test framework and EU Solvency II both suggest that a cost of capital rate of 6 percent per year above the risk-free rate be used for risk margin calculation. See http://www.naic .org/documents/committees_smi_int_solvency_switzerland.pdf and Sandström (2016).

[12] See the discussion in Section 8.5.2 of McNeil, Frey, and Embrechts (2015).

Notes (Continued)

[13] Note that although the total risk capital is nonnegative, in theory, the allocation to a certain program could be negative. This happens mostly when the program hedges the risk of the government's portfolio, and in this case, the premium is determined to be less than its expected loss.

[14] The Gaussian copula producing the weak tail dependence has a correlation coefficient parameter of $\rho = 0.87$, and the Gumbel copula producing strong tail dependence has a parameter of $\theta = 3$. The parameters are chosen to make the comparison fair, in the sense that they produce the same level of dependence as gauged by a third measure, Kendall's tau.

[15] Simulation is not needed for the stand-alone pricing case, since it does not require capital allocation and the TVaR of an exponential random variable is explicitly known. However, for the other scenarios, we must rely on simulation to estimate the capital allocations and TVaRs. Note that because a loss exceeding the 99th percentile is a rare event, only a very small fraction of the simulation samples will have an observed loss. As a result, the simulations require an extremely large number of trials to numerically estimate the TVaR at the 99 percent level. We use 2.5×10^7 samples for each loss variable to reduce simulation errors.

[16] The dependence structure of X_1, X_2, and X_3 is given by the Gaussian copula with correlation matrix

$$\begin{bmatrix} 1 & 0.9 & 0.5 \\ 0.9 & 1 & 0.1 \\ 0.5 & 0.1 & 1 \end{bmatrix}$$

or the Gumbel copula

$$C(u_1, u_2, u_3) = \exp\left\{-\left[(-\log u_1)^\theta + (-\log u_2)^\theta + (-\log u_3)^\theta\right]^{1/\theta}\right\}$$

with $\theta = 1.5$.

[17] We delete loans with missing characteristics.

[18] For the loans originated after 2013, we use the 2011-13 severities.

[19] If n is an odd number, one may randomly delete an observation to make the total number of observations even.

[20] Using the inversion of Kendall's tau method implemented in the R package "copula," we estimate the Gumbel copula parameter as 1.245 with a standard error of 0.012. Thus, the fitted copula is given by

$$C(u_1, u_2) = \exp\{-[(-\log u_1)^{1.245} + (-\log u_2)^{1.245}]^{1/1.245}\}.$$

[21] See https://www.fhfa.gov/Media/PublicAffairs/Pages/Results-of-Fannie-Mae-and-Freddie-Mac-Guarantee-Fee-Review.aspx.

[22] For a Pareto random variable with shape parameter $\alpha > 1$ and scale parameter θ, and hence, density function

$$f(x) = \frac{\alpha \theta^\alpha}{(x + \theta)^\alpha}, \qquad x > 0,$$

its TVaR at level $q \in (0, 1)$ is

$$\theta \left[\frac{\alpha(1 - q)^{-1/\alpha}}{\alpha - 1} - 1\right].$$

References

Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson. 2017. "Measuring Systemic Risk." Review of Financial Studies 30, no. 1: 2-47.

Ambrose, B., and A. Warga. 2002. "Measuring Potential GSE Funding Advantages." Journal of Real Estate Finance and Economics 25, no. 2-3: 129-50.

Goodman, L., and J. Zhu. 2015. "Loss Severity on Residential Mortgages." Journal of Fixed Income 25, no. 2 (Fall): 48-57.

Hanson, S. G., D. S. Scharfstein, and A. Sunderam. 2016. "Fiscal Risk and the Portfolio of Government Programs." National Bureau of Economic Research Working Paper no. 22763.

Heaton, J. C., D. Lucas, and R. L. McDonald. 2010. "Is Mark-to-Market Accounting Destabilizing? Analysis and Implications for Policy." Journal of Monetary Economics 57, no. 1: 64-75.

Hüsler, J., and D. Li. 2009. "Testing Asymptotic Independence in Bivariate Extremes." Journal of Statistical Planning and Inference 139, no. 3: 990-8.

Ibragimov, R., D. Jaffee, and J. Walden. 2010. "Pricing and Capital Allocation for Multiline Insurance Firms." Journal of Risk and Insurance 77, no. 3: 551-78.

IFE. 2007. "The Measurement and Management of Mortgage Credit Risk in the United States: Implications for Emerging Mortgage Markets." U.S. Department of Housing and Urban Development, Office of Policy Development and Research.

Kelly, B., H. Lustig, and S. Van Nieuwerburgh. 2016. "Too-Systemic-to-Fail: What Option Markets Imply about Sector-Wide Government Guarantees." American Economic Review 106, no. 6: 1278-319.

Lucas, D. 2012. "Valuation of Government Policies and Projects." In A. W. Lo and R. C. Merton, eds., Annual Review of Financial Economics 4, October: 39-58.

———. 2014a. "Evaluating the Cost of Government Credit Support: The OECD Context." Economic Policy 29, no. 79: 553-97.

———. 2014b. "Evaluating the Government as a Source of Systemic Risk." Journal of Financial Perspectives 2, no. 3: 45-58.

———. 2016. "Credit Policy as Fiscal Policy." Brookings Papers on Economic Activity, Spring: 1-57.

Lucas, D., and R. McDonald. 2006. "An Options-Based Approach to Evaluating the Risk of Fannie Mae and Freddie Mac." Journal of Monetary Economics 53, no. 1: 155-76.

References (Continued)

McNeil, A. J., R. Frey, and P. Embrechts. 2015. Quantitative Risk Management. Princeton, N.J.: Princeton University Press.

Passmore, W. 2005. "The GSE Implicit Subsidy and the Value of Government Ambiguity." Real Estate Economics 33, no. 3: 465-86.

Sandström, A. 2016. Handbook of Solvency for Actuaries and Risk Managers: Theory and Practice. Boca Raton, Fla.: CRC Press.

U.S. Congressional Budget Office. 2012. "Fair-Value Accounting for Federal Credit Programs." Issue Brief, March.

U. S. Government Accountability Office. 2016. "Credit Reform: Current Method to Estimate Credit Subsidy Costs Is More Appropriate for Budget Estimates Than a Fair Value Approach." Report to the Ranking Member, Subcommittee on Financial Services and General Government, Committee on Appropriations, U.S. Senate.

Zanjani, G. 2002. "Pricing and Capital Allocation in Catastrophe Insurance." Journal of Financial Economics 65, no. 2: 283-305.