How to Make the Global Market Shock Coherent

Greg Hopper | May 5, 2025

Introduction

This note examines how the Global Market Shock (GMS), a key component of the Federal Reserve's stress test and resulting stress capital charge, can be made coherent. As currently operated, the Global Market Shock is incoherent in its design, and its results thus likely misstate significantly the capital depletion a bank would experience under a real-world stress.

A scenario is coherent if it is internally consistent and targets the set of risks the scenario designer intends to capture. Internal consistency implies that the shock sizes and the correlations between shocks make sense in the context of the risks being targeted.

Scenario coherence is necessary because it makes the GMS results interpretable: scenario coherence makes it clear what risks are being targeted, how likely the scenario's risks are to materialize in the context of the current market environment and exactly what assumptions are being made about each risk.

The GMS as currently conducted is not coherent. In the GMS, for example, the calibration of illiquidity risk is completely opaque, making it impossible to tell how likely the assumed illiquidity risk is or whether it makes sense in the context of the entire scenario. In the GMS, risk factors, which are financial prices such as equity prices, bond spreads, and interest rates, are assumed to be illiquid over liquidity horizons – the assumed periods over which no trades can be completed. However, the GMS documentation, such as the 2025 CCAR documentation, leaves unexplained the motivation for and choice of the illiquidity horizons by asset class, providing only vague guidance:

> Risk factor shocks are calibrated based on assumed time horizons. The calibration horizons reflect several considerations related to the scenario being modeled. One important consideration is the liquidity characteristics of different risk factors. These characteristics may vary depending on the specified market shock narrative. More specifically, the calibration horizons reflect the variation in the speed at which banks could reasonably close out, or

> effectively hedge, risk exposures in the event of market stress. The calibration horizons are generally longer than the typical times needed to liquidate exposures under normal conditions because they are designed to capture the unpredictable liquidity conditions that prevail in times of stress. In addition, shocks to risk factors in more liquid markets, such as those for government securities, foreign exchange, or public equities, are calibrated to shorter horizons (such as three months), while shocks to risk factors in less liquid markets, such as those for non-agency securitized products, have longer calibration horizons (such as 12 months).[1]

Similarly, the correlations between GMS shocks across asset classes, such as between equities and credit, or within asset classes are unexplained; so too are the correlation of shocks to different rates in the term structure of interest rates.

Thus, the Federal Reserve's scenario designers are subject to no objective coherency standard when they design the GMS, and banks, financial markets, policymakers and the public can draw no meaningful conclusions about the results of the test. Are the risks tested in the GMS plausibly what banks might actually experience in a stressed environment, or are they essentially so rare as to be irrelevant, with more relevant, material risks left unanalyzed?

One way to make the GMS coherent is by following a methodical process, consisting of the following elements, as suggested in an earlier article[2]:

- Conduct a risk identification process that identifies the most important risks in the current market environment
- Write a GMS scenario narrative that describes the scenario in detail, showing how the narrative captures the identified material risks
- Develop econometric models for the risk factors to make explicit the assumptions being made

In this note, we show how this process would work by designing a simple coherent scenario that features

1) illiquidity in the credit and equity markets
2) correlations between equity and credit market shocks that are significantly higher than historically observed
3) an unusual movement of the U.S. term structure

To keep this example simple, we are not considering other aspects of the capital regime that may already be in force. Most notably, the Federal Reserve is considering changes to the GMS at the same time that it and other banking agencies are considering how to implement the market risk component of the Basel Accord, known as the Fundamental Review of the Trading Book, or FRTB. As described in earlier notes[3], the purpose of the GMS appears very similar to the purpose of the FRTB in terms of the risks each seeks to capture. Thus, designing each independently would produce an incoherent capital requirement for market risk, and almost certainly dramatically overstate the capital needs of banks

[1] Board of Governors of the Federal Reserve System, "2025 Stress Test Scenarios," February 2025, pg. 9, available at https://www.federalreserve.gov/publications/files/2024-stress-test-scenarios-20240215.pdf

[2] Hopper, G, "Designing Coherent Scenarios: A Practitioner's Perspective," Chapter 7 in *Handbook of Financial Stress Testing,* edited by Farmer, D, Kleinnijenhuis, A, Schuermann, T, and Wetzer, T, Cambridge University Press, 2022

[3] See for example Hopper, G, "How Can the Global Market Shock More Effectively Complement the Fundamental Review of the Trading Book?", Bank Policy Institute, 2023, available at https://bpi.com/how-can-the-global-market-shock-more-effectively-complement-the-fundamental-review-of-the-trading-book/

making markets. The GMS scenario design process should take into account the risks captured in the entire regulatory capital regime when using the techniques suggested in this note.

Designing a Coherent GMS

Risk Identification

The most important first step in designing any stress test is to perform a thorough analysis of the risks the test is designed to capture. The procedure for a regulatory stress test will be a little different from a bank stress test. A bank designs its stress tests to target idiosyncratic risks it may face. A regulatory stress test, on the other hand, is designed to capture risks that any bank could face under unusual or extreme market conditions.

Let us suppose for the purposes of this example that our risk identification process identified illiquidity of the credit and equity markets as a key risk. As an additional risk, we might expect historical correlations to become much higher than historically observed in the credit and equity markets. Along with that risk, current market conditions also suggest that although U.S. Treasury rates would be expected to fall across the board, the term structure curve also could steepen.

The output of the risk identification exercise should be a narrative that discusses each risk that has been identified and how they all fit together, i.e., the story line for how the realization of risks and market events results in the GMS. A written narrative helps to keep the stress test design coherent and plausible. The narrative also helps independent model validators as well as commenters understand the motivation for the stress test design and what risks the stress test intends to capture.

Incorporating the Three Identified Risks in the GMS

Risk 1: Illiquidity in the Credit and Equity Markets

A global market shock will always require a period of illiquidity as a common risk for at least some risk factors. Since an assumed financial panic will be precipitated by generalized concerns about bank counterparty risk, the likely length of the period of illiquidity can be derived from looking at a metric that measures bank counterparty risk. Chart 1 plots the Libor-OIS spread from 2007-2010. As is apparent, the spread shot up rapidly over a two-week period starting in September 2008, and then came rapidly back down over the next two weeks, where it remained unusually elevated until May 2009.

Chart 1

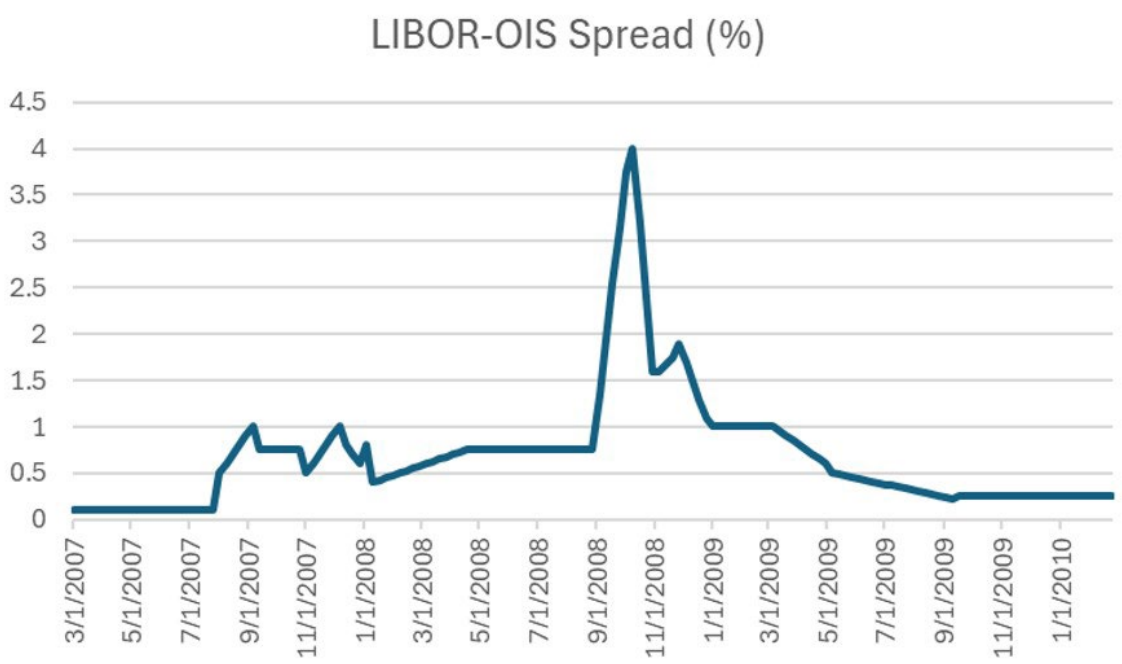

This analysis suggests that liquidity risk could be divided into two periods. In the first period, illiquidity for most assets would be extreme, and liquidity could be assumed to be non-existent. In the second period, illiquidity could be assumed to be partial, i.e., only trades of a certain size could be done without moving the market. The length of that second period would depend on the results of the risk identification performed for each asset class. To be able to generate the stress test shocks, we need to develop two complementary models:

- A simulation model for the risk factors—the S&P 500, the investment-grade (IG) index, and high-yield (HY) CDS index—in this example
- An illiquidity model for trades that depend on the risk factors over the two periods

Simulation Model for Equities and CDS Spreads

To simulate equities and CDS, it is important to capture their empirical features. Both risk factors have time-varying volatility, leverage effects (i.e., volatility rises when equity prices fall or CDS spreads widen) and fat-tailed distributions[4], which implies that large daily percentage changes in the risk factors occur more frequently than the standard normal distribution would imply.

To simulate these risk factors, we estimate an EGARCH model.[5] The EGARCH model will allow us to simulate daily returns of each risk factor to any horizon that we need. Because we want to simulate

[4] The conditional innovation in the model should be a fat-tailed distribution. We use a generalized error distribution in the analysis.
[5] Denoting the daily observation of an equity price or spread as s_t, the model below is estimated

$$r_t = 100.0 * \left(\frac{S_t - S_{t-1}}{S_{t-1}}\right)$$

$$r_t = \mu + u_t$$

$$u_t = \sqrt{h_t}\epsilon_t$$

$$\ln(h_t) = \omega + \alpha[|\epsilon_{t-1}| - E(\epsilon_{t-1})] + \gamma\epsilon_{t-1},\ \ \epsilon_t \sim GED(v)$$

paths that potentially have never been observed but that would also be consistent with a financial crisis, we sample the estimated residuals for each day's simulation from June 2008 to June 2009 only, ignoring other more benign periods. For the same reason, we also elevate the starting volatility in the simulations by assuming it is the average daily volatility observed over June 2008 to June 2009.

Chart 2 illustrates a simulation of the S&P 500 over a 20-day period. Chart 3 depicts how the simulated annualized equity volatility varies over the 20-day period. Chart 4 shows a representative simulation of the IG CDS index over 30 days. Chart 5 shows how the simulated IG CDS annualized volatility varies over that period. In the simulation models, we perform a high number of simulations of the risk factors so that we have many simulated paths over which to run the illiquidity model we will develop. To illustrate the equity results, for the S&P 500, we show the path that leads to the 1st percentile 20-day simulated level. To illustrate the CDS results, for the CDS IG index, we show the path that leads to the 99th percentile 30-day spread level.

Chart 2

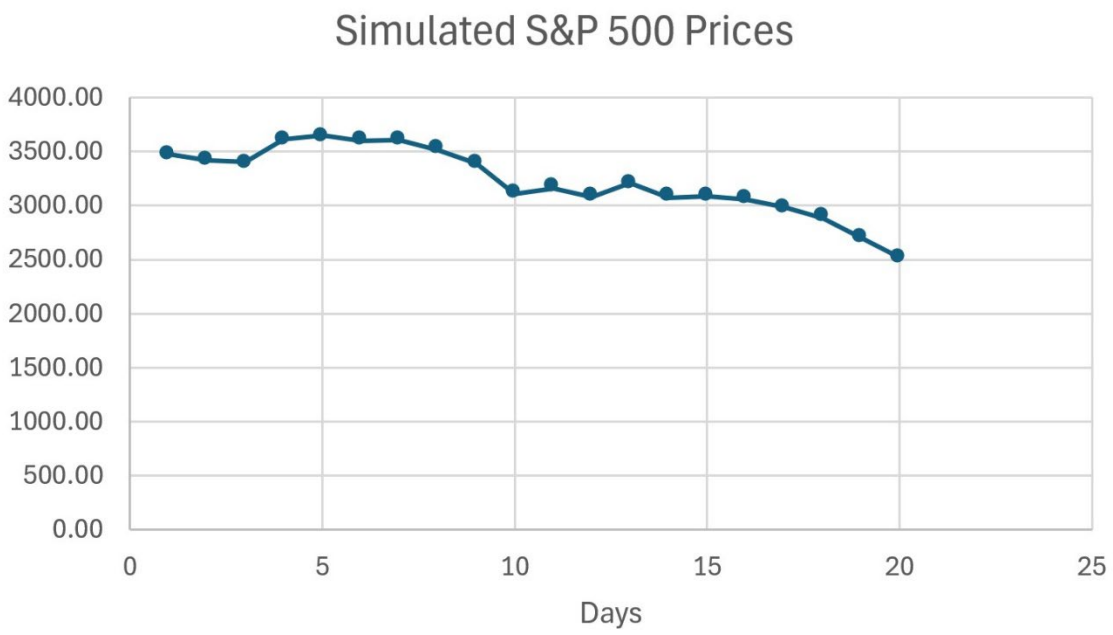

Chart 3

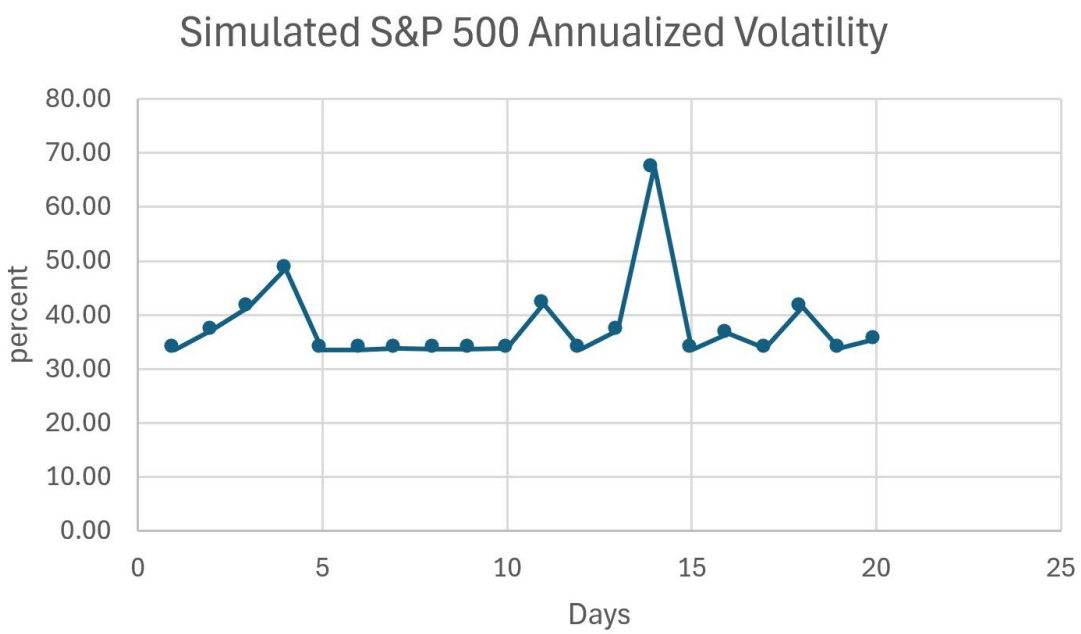

Chart 4

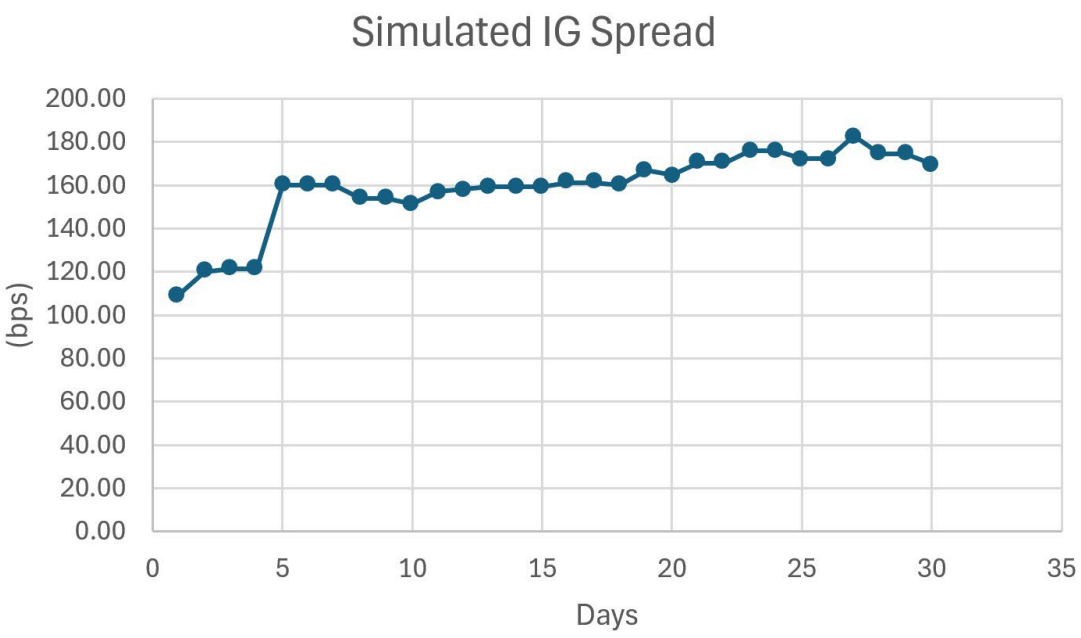

Chart 5

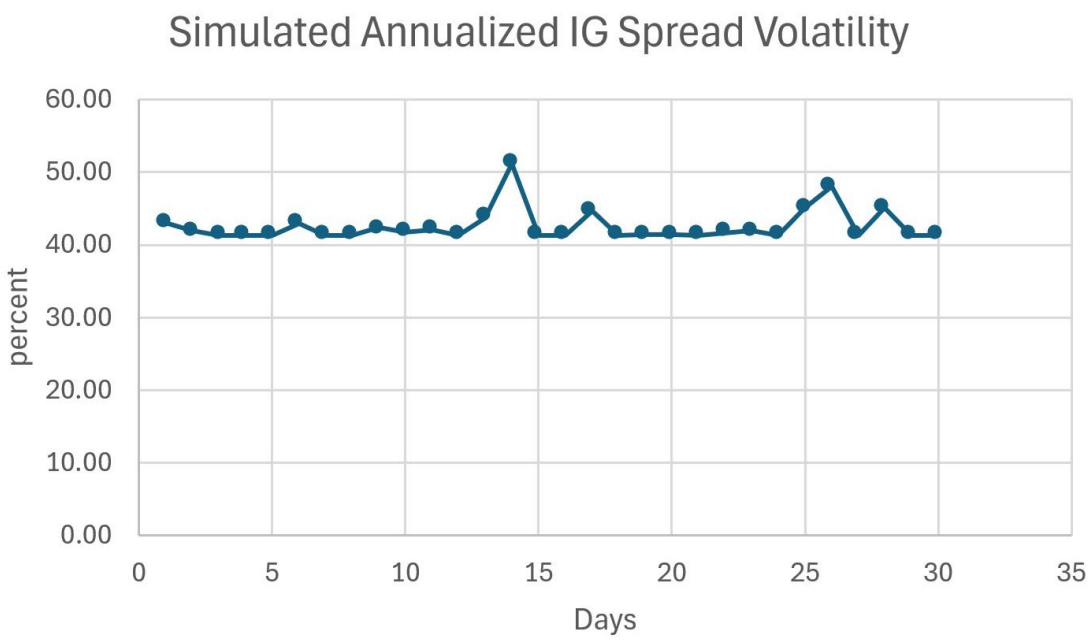

Illiquidity Model

The risk identification analysis suggested that we could break up the period of illiquidity into two periods. Because there is little data on liquidity of assets, especially during a financial crisis, we must keep the model simple so that it can be calibrated through some combination of data, models and judgment.

Accordingly, consistent with the climb of the Libor-OIS spread in chart 1, we specify a period of 10 business days when the market has no liquidity at all. During the first 10 days of the simulation of equity prices or CDS spreads, equity and CDS positions can neither be bought nor sold. In the second period of illiquidity, we assume that only a portion of an asset can be bought or sold each day. We make that proportion stochastic by drawing the amount of the asset that can be bought or sold from a log normal distribution with a mean trade size along with a standard deviation of the trade size, where the trade size is expressed as a fraction of 1. The number of days in the second period is random; it will be as long as necessary to completely liquidate the asset when the amount that can be traded per day is random.

The model works by taking the equity prices or CDS spreads and simulating them over a random number of days, using the GMS starting date to derive the starting value of the asset. Over the first 10 days, no trades are possible. For each day after the 10th, we draw a trade size from the log normal distribution[6], calculate the gain or loss when we sell it and then move to the next day. We then repeat the process. We finish the simulation path when we have fully sold the asset.

We derive the GMS shocks by finding the total liquidation losses on each liquidation path over 50,000 simulated liquidation paths. We then find the 99th percentile worst loss and then find the equivalent shock that would produce that loss. For example, to find the investment grade (IG) shock, we perform 50,000 simulations of the liquidation of a five-year CDS trade assumed to be at-the-money during the

[6] We use the log normal distribution for the sake of illustration. The model designer may decide to use an alternative distribution. For example, a poisson distribution could also be used. The constraint is that the distribution must always generate positive trade opportunities

simulation. We take the 99th percentile worst liquidation loss and then find the CDS shock that would produce that loss.

Table 1 illustrates how the calculation would work on a single simulated path for the S&P 500 with a 10-day period of no liquidity and a 10-day period of limited liquidity. We assume that we have 100 shares with a starting price of 3583.07, so the total value of the position is 3583.07 X 100 = 358,307. We first simulate the equity price for at least 20 days, since we do not know how long it will take to fully sell off the position after waiting 10 days. Then we simulate the daily trading opportunities. We assumed the mean trading opportunity was 35,830.07 (1/10th of the total) with a standard deviation of 17,915.35 (one half of the average or 1/20th of the total). We simulate the daily trading opportunities with that mean and standard deviation until we have fully liquidated the position. The simulation may take fewer or more than 10 days. In this case, we show an example that takes exactly 10 days.

During the first ten days, no trading opportunities are available. On day 11, we simulate a trading opportunity of 31,560.24, which translates into an opportunity to sell 8.8 of the 100 shares in the position. When we sell them, we incur a loss of the simulated price on that day minus the initial price times 8.8 shares, or -3,758.17. We continue the liquidation process for 10 days, obtaining a cumulative loss of -58,426.75, or a percent loss of -16.31 percent. On this path, the GMS S&P 500 shock would then be -16.31 percent.

Table 1

			Starting value of 100 shares = 3583.07 X 100 = 358,307			
Day	S&P 500	Simulated Trading Opportunities	Calculation	Shares Liquidated	Calculation	Loss
0	3583.07					
1	3479.23	0				
2	3416.78	0				
3	3400.58	0				
4	3610.40	0				
5	3645.84	0				
6	3597.57	0				
7	3603.62	0				
8	3513.85	0				
9	3391.22	0				
10	3108.71	0				
11	3155.86	31,520.64	(31,520.64/358,307) X 100 =	8.80	(3155.86-3583.07) X 8.80 =	-3,758.17
12	3079.98	24,664.43		6.88		-3,463.07
13	3205.44	47,842.32		13.35		-5,042.28
14	3073.54	31,822.68		8.88		-4,525.37
15	3089.16	41,161.87		11.49		-5,673.99
16	3055.36	33,041.68		9.22		-4,866.30
17	2983.25	41,869.95		11.69		-7,009.21
18	2884.22	61,934.46		17.29		-12,079.75
19	2700.44	23,505.67		6.56		-5,790.25
20	2519.21	20,943.30		5.85		-6,218.36
	Total trades	**358,307.00**	**Total shares liquidated**	**100.00**	**Total losses**	**-58,426.75**
					Percent loss of 358,307	-16.31%

Table 2 reports the results of the simulations under various assumptions on the average number of days to liquidate, including the 10-day no-liquidation period. We perform 50,000 simulations as calculated in Table 1 above and choose the 99th percentile worst shock.

Table 2

		Average days to liquidate			
	Standard deviation as multiple of average trading opportunity	20	30	40	50
	0.5	-24.75%	-28.77%		
S&P 500	1	-25.05%			
	2	-26.50%			
	0.5		47.44%	54.75%	
IG CDS	1		48.33%		
	2		48.09%		
	0.5			44.23%	52.33%
HY CDS	1			44.57%	
	2			47.65%	

In Table 1, we have assumed that the S&P 500 has the most liquidity, taking 20 days on average to liquidate a position during a financial panic. We assume IG CDS takes 30 days on average, while HY CDS is the least liquid, taking 40 days on average to sell out of a trade completely. The results show that the shocks are essentially invariant to different assumptions about the standard deviation of the trading opportunities per day. For example, for equities a standard deviation that is half the mean trading opportunity per day yields a shock that is very close to the assumption that the standard deviation is twice the mean trading opportunity per day. The table also suggests that the calibration of the mean trading opportunity size per day is the most important parameter. The effect is largest for HY CDS. An assumption that it takes on average 40 days to liquidate a five-year HY CDS produces a shock between 44 percent and 47 percent. Changing the average days to liquidate a representative trade to 50 days produces a shock of 52 percent.

Risk 2: High Correlation in the Equity and Credit Shocks

The simplest way to impose very high correlation in the credit and equity shocks is to calculate each shock independently as we did above without correlating the residual draws for the simulations. On the other hand, if we wanted to impose some kind of explicit correlation between shocks, we could impose historical correlations in the simulations by drawing the estimated residuals for each daily simulation simultaneously. For example, if we were simulating an equity index as well as an IG and HY index, we could draw each residual for equities, IG, and HY on the same historical day, thus preserving the historical correlations. If we wanted to break the historical correlations but not impose something as extreme as completely independent simulations of the risk factors, then we could order all the historical

innovations from largest positive to largest negative. Then, we would draw the equity, IG and HY innovations according to some assumed correlation[7] that would ensure that if a large innovation is drawn in the IG simulation, large innovations are also drawn for the HY and equity simulation.

Risk 3: A Drop in Rates with the Term Structure Steepening

To define an empirically reasonable stress test for the term structure of U.S. Treasury rates, we can use a standard model, principal components analysis (PCA). PCA will allow us to isolate the most common moves in the yield curve historically. Once we have estimated those common moves, we can use them to design a set of stress shocks.

To illustrate how to use PCA to design an interest rate shock, we obtained historical daily data on the level of U.S. Treasury rates for the one-month, three-month, six-month, one-year, two-year, three-year, five-year, seven-year, 10-year, 20-year and 30-year maturities over the period Jan. 1, 2005 to April 22, 2025. After estimating PCA, chart 6 shows the resulting first three principal components representing a one-standard deviation daily move in rates, measured in basis points.[8] The first principal component is a parallel shift of the yield curve of approximately 65-70 bps. The second principal component represents a change in curvature, in which the short end of the curve rises more than the long end. The third principal component represents a twist in the yield curve, in which the short and long ends of the curve rise more than the middle maturities. We can also define these one-standard-deviation moves in the opposite direction by changing all the signs of each maturity.

Chart 6

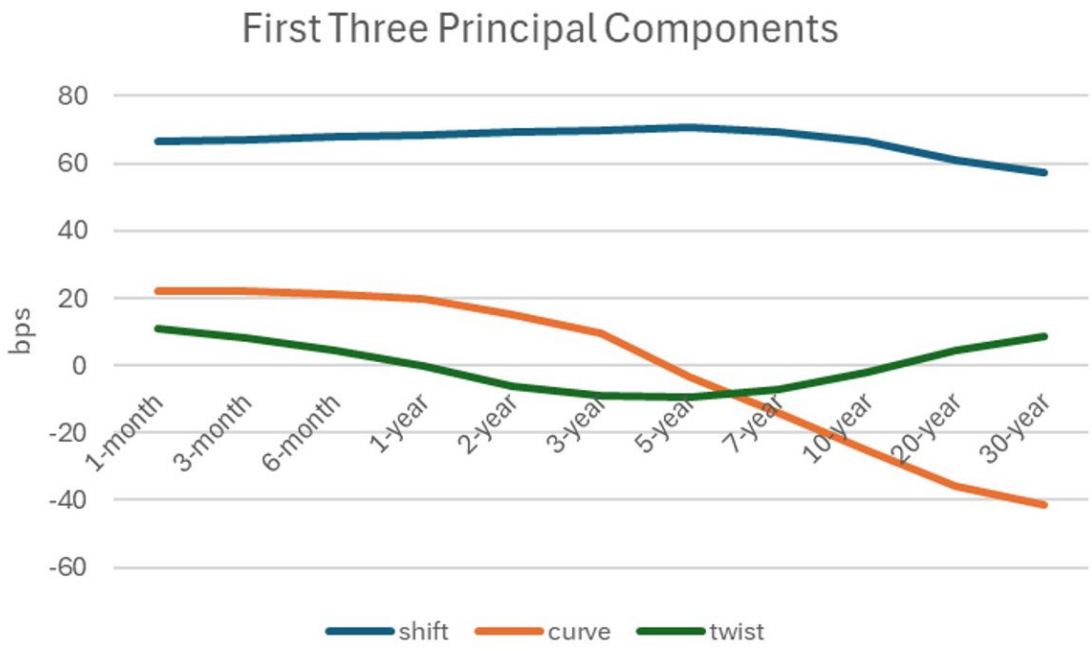

To capture the risk of a downward shift in interest rates as well as a change in curvature, we take a minus-two-standard-deviation move of the first principal component and add to it a positive one standard deviation move of the curvature, resulting in a net three standard deviation daily move in

[7] A Gaussian or t-copula could be employed for this purpose.
[8] The eigenvectors are equal to the principal components and the eigenvalues equal the variance of each component. We calculate a one-standard deviation move by multiplying the square root of a component's eigenvalue by the elements of the principal component.

interest rates. The positive curvature will attenuate the downward shocks in the short end of the yield curve, reducing the shock sizes relative to the long rates. Table 3 shows the resulting stress test. In this methodology, the number of standard deviations can be adjusted to create a more or less severe stress test.

Table 3

Term	1-month	3-month	6-month	1-year	2-year	3-year	5-year	7-year	10-year	20-year	30-year
Shock (bps)	-111	-112	-115	-117	-124	-131	-144	-153	-158	-158	-156

Using These Techniques Would Improve GMS Coherence

The model results reported above for equities and credit used as starting values for the simulation the risk factor values observed on Oct. 14, 2022, the 2023 CCAR date. It is interesting to compare the 2023 GMS shocks to the shocks produced by the illiquidity model. In the 2023 CCAR, the U.S. equity shock was -26.3 percent whereas the model produced shocks in the neighborhood of -25 percent, essentially the same. Thus, the equity shocks could be explained by the assumption of a 10-day period of no liquidity, followed by a 10-day period of partial liquidity, under the assumption of high correlation between credit and equity risk factors. However, the difference between the CCAR GMS shocks and the illiquidity model GMS shocks for IG and HY index spreads is huge. The 2023 CCAR shock for IG was 177 percent versus 48 percent from our model. For HY, the 2023 CCAR shock was 80 percent whereas the model produced a shock of about 45 percent.

The current Fed-derived GMS CDS shocks seem very large when compared to those generated by an explicit liquidation model. We can reproduce the CCAR IG GMS shock of 177 percent in the model by assuming a 10-day period of no liquidity and a 300-day period of partial liquidity. The CCAR GMS HY shock of 80 percent can be reproduced in the model by assuming a 10-day period of no liquidity and a 110-day period of partial liquidity. These implicit periods of partial liquidity of CDS that are imposed in the GMS are quite implausible.

The illiquidity model results strongly suggest that the 2023 GMS shock sizes for CDS are not coherent. HY is generally less liquid than IG and yet the 2023 GMS is implicitly assuming that HY is much more liquid than IG. The CCAR HY GMS shock is calibrated to about a half-year period of illiquidity whereas the CCAR IG GMS shock is calibrated to a period of illiquidity that is greater than one year. As another point of comparison, the FRTB assumes that the period of illiquidity for HY is 60 days while the period of illiquidity for IG is 40 days, a much shorter period of illiquidity than in CCAR with HY less liquid than IG. We see similar unusually large shocks in GMS exercises in other years, including the most current 2025 test, suggesting that lack of coherence is a persistent problem in the GMS.

Using an explicit illiquidity model to calibrate the shock sizes would improve the coherence of the GMS, because it will be clear which risk factors are assumed to be less liquid and which are assumed to be more liquid, and by how much. In the current GMS, it is very difficult to understand exactly what illiquidity assumptions have been imposed. Similarly, using a PCA model to define the interest rate tests clarifies exactly what assumptions are being made.

Conclusion

The introduction of the CCAR stress test by the Fed was a major regulatory risk management innovation. However, the GMS has consistently lacked scenario coherence, dramatically limiting the ability of the Fed, financial markets, policymakers and the public to understand the effectiveness of the GMS in

reducing systemic risk to the financial system while not harming market liquidity or reducing economic growth. In this note, we have suggested a methodology for making the GMS coherent. Besides its other benefits, coherency requires model transparency, which when coupled with a notice and comment period, will dramatically improve the GMS, allowing the Fed to review and, if appropriate, incorporate a wide number of views and perspectives from market participants, the academic community, and banks with expertise in particular risk areas. This open review process is bound to improve the GMS specification, reducing the chance that no risk is brewing on the horizon that is missing from the stress test. In making the GMS coherent, the Federal Reserve should also consider the risks already captured in other regulatory capital processes, so that risks are not duplicated by the GMS. The GMS will be a more effective regulatory tool if it captures different risks than other capital methodologies.