MILLIMAN WHITE PAPER

Understanding mortgage reinsurance loss reserving

Fundamental concepts of mortgage reinsurance risk management

Andrew Netter
Jonathan B. Glowacki, FSA, CERA, MAAA
Michael C. Schmitz, FCAS, MAAA

Milliman

Mortgage reinsurance, with unique claim timing and loss emergence patterns, demands a thoughtful approach when setting loss reserves—with key differences from other P&C lines

Following the 2007 global financial crisis, mortgage reinsurance emerged as a reliable source of capital as well as a prudent risk management tool to support the safety and soundness of the U.S. residential housing market. Over the past 10 years, dozens of global reinsurers began writing mortgage reinsurance. Given this unique and relatively new line of business, reinsurers should deeply understand the varying risk profiles of this exposure prior to developing a loss reserving approach.

This white paper outlines loss reserving considerations that have been developed from both the reserving approach taken by primary mortgage insurers as well as general best practices from broader property and casualty (P&C) reserving. After exploring these two topics, we compile a list of key considerations when setting up a mortgage reinsurance reserving approach.

Note that Milliman analyzes reinsurance structures using standard actuarial methods. The authors of this report are not accountants or auditors and Milliman expresses no opinion as to whether these reinsurance transaction treatments comply with any applicable accounting or auditing standards.

What is the standard primary mortgage insurer approach for loss reserving?

The six U.S. private mortgage insurers (PMIs) establish case reserves for estimated losses when they are first notified of an insured loan becoming delinquent. Generally, the PMIs consider a loan delinquent when it is 60 days or more past due (2+ missed mortgage payments) and has yet to result in a final resolution (e.g., cure, prepayment, or claim event). The industry-standard accounting treatment for PMIs under U.S. GAAP is to only establish loss reserves for delinquent loans and to not establish reserves for future claims on active, nondelinquent, insured loans. Thus, if a loan cures its delinquency, it no longer is part of the population on which reserves are established and any subsequent redefault is considered a new loss occurrence with an accident date corresponding to the new default.

Three ways to improve upon delinquency-based mortgage reinsurance loss reserves

While this is the standard approach for the PMIs, we have identified three key aspects that reinsurers can use to improve upon this delinquency-based approach:

1. Model the risks from the entire pool of covered mortgages—not just the active delinquencies
2. Establish reserves as premium is written (contracts have annual rates on line)—not only after a negative outcome begins
3. Leverage a stochastic simulation or suite of deterministic scenarios to remove reactivity on idiosyncratic delinquency events

MODEL THE RISKS FROM THE ENTIRE POOL OF COVERED MORTGAGES

Typically, mortgage reinsurance transactions have terms of 10+ years with early termination options at five years. This transaction length is generally thought to cover most of the default risk and losses from the underlying pool of 30-year mortgages—mortgages typically reach peak default emergence around three to five years. As a result, a more robust reserving methodology considers the full risk of the underlying mortgage pool (across the full exposure term), not just a point-in-time estimate of the risk to the transactions (as reserving based on active delinquencies accomplishes).

Because reinsurers may not be bound by the same reserving treatment as the PMIs, many have designed their reserving methodologies to consider more holistic views of revenue and expenses, in order to address this known drawback. Many of the transactions—in particular those issued under the Freddie Mac Agency Credit Insurance Structure (ACIS) program—feature multiple excess of loss (XOL) tranches with varying levels of credit enhancement (i.e., reinsurance attachment levels), which can be seen in Figure 1.

FIGURE 1: FREDDIE MAC ACIS 2021-SAP10 TRANSACTION STRUCTURE

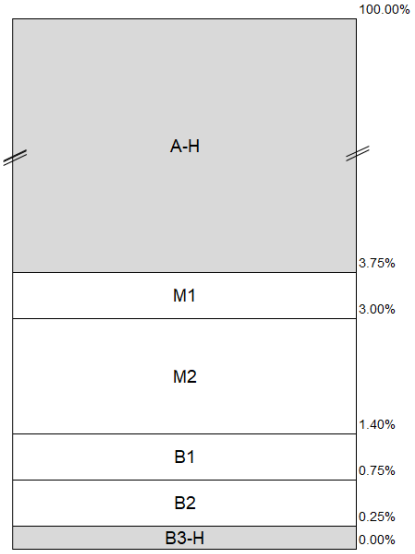

The percentages on the right of Figure 1 represent the level of loss that must be realized – as a % of the total underlying pool of mortgages – in-order to attach the reinsurance coverage. Given all mortgages reinsured in these transactions begin as performing loans, with a borrower making their full monthly payments, loans must become delinquent in the future to create a loss to reinsurers. A reserving methodology based solely on active delinquencies will typically result in little forewarning of risks to the risk-remote tranches—and will not reflect any reserves until a loss-producing event can be imminently caused by the active delinquencies, which is likely to be in the later years of the runoff of the contract, when premium earnings are lowest due to runoff of the pool of mortgages.

ESTABLISH RESERVES AS PREMIUM IS WRITTEN (CONTRACTS HAVE ANNUAL RATES ON LINE)

Unlike other P&C lines, where high-loss events can occur at any time, mortgage loans typically reach peak defaulting periods around three to five years. As a result, the periods before this (e.g., years 1 and 2) typically result in loans paying high amounts of premium with few losses.

To further illustrate this mismatch, we extracted future government-sponsored enterprise (GSE) Credit Risk Transfer (CRT) transaction performance estimates from M-PIRe.[1] These forecasts are based on models fit on over 20 years of historical mortgage performance. Figure 2 shows M-PIRe's underlying mortgage reference pool forecasts for a Freddie Mac ACIS transaction (2021-SAP10). The blue series reflects the percentage of the unpaid principal balance (UPB) remaining over time, and the orange line reflects the percentage of the UPB that is 60+ days delinquent. The blue series is the driver of premium (as the limit on GSE CRT transactions declines as a function of the underlying UPB), and the orange series is the driver of reserves or losses. One can see that, when the orange series begins to plateau (reflecting minimal new loans entering the population 60+ days delinquent), approximately 40% of the underlying UPB has already paid off. Thus, using delinquencies as the sole driver of loss reserves will result in a substantial mismatch between the revenue and expenses of the contract.

FIGURE 2: FREDDIE MAC ACIS 2021-SAP10 PROJECTED CASH FLOWS

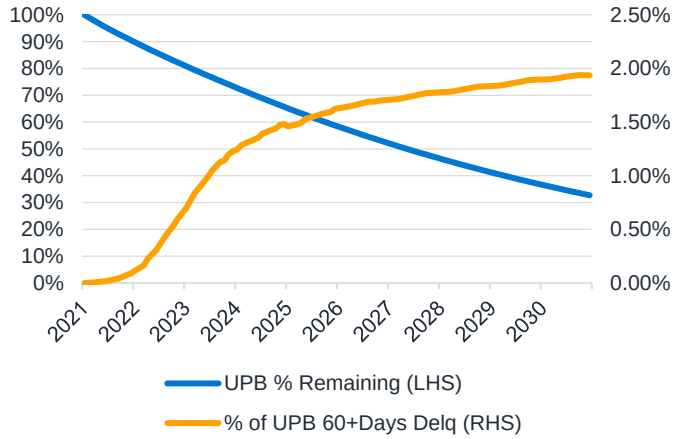

Establishing reserves as premium is earned can create a more robust reserving approach over time and the mismatch in booked premium and losses can be reduced.

LEVERAGE A STOCHASTIC SIMULATION OR SUITE OF DETERMINISTIC SCENARIOS TO REMOVE REACTIVITY ON IDIOSYNCRATIC DELINQUENCY EVENTS

The discussions above centered around delinquencies as the events generating ultimate insurance losses (and thus reserves). Although this is true, delinquencies do not necessarily lead directly to losses; occasionally events cause large spikes in delinquencies that lead to very little losses. This can leave the framework overexposed to this indicator of stress.

Two such events over the recent past that led to large spikes in delinquencies without any discernible increase in ground-up net losses to date have been (1) the 2017 Atlantic hurricane season (Harvey, Irma, and Maria), and (2) the economic disruption following the onset of the 2020 COVID-19 pandemic.

[1] M-PIRe is a software product designed specifically to evaluate and manage mortgage reinsurance. For more information, see https://www.milliman.com/en/products/milliman-m-pire.

Although the underlying drivers are multifaceted, based on recent historical experience, natural disasters typically cause borrowers to become delinquent, but ultimately result in relatively benign loss outcomes. Borrowers impacted by the natural disaster often have the impacts mitigated by homeowners and flood insurance as well as local and federal assistance programs.

During the COVID-19 pandemic, the U.S. federal government took unprecedented action to assist borrowers facing temporary financial hardship. Although not fully resolved, this largely resulted in many borrowers becoming delinquent as they took advantage of forbearance, but ultimately resulted in no loss payment for most delinquencies, much like the above natural disaster situation. Many of the loans that were extended forbearance during the pandemic were resolved as a cure or prepayment rather than a claim.

To further illustrate these events using actual GSE CRT performance, we extracted historical delinquency and loss data from M-PIRe. Figure 3 shows historical mortgage reference pool forecasts for eight ACIS transactions issued in 2019. The 2019 transactions were squarely impacted by the COVID-19 pandemic. The blue series reflects the conditional percentage of the unpaid principal balance (UPB) becoming 60-days delinquent in each month, and the orange series reflects the cumulative net loss rates from those same pools of mortgages. In theory, spikes in the blue series should result in future increase in the orange series—as increases in delinquencies should lead to increases in losses, all else equal. However, cumulative losses to date on these pools of mortgages have been less than 0.01% of the pool's issuance UPB. If a reserving framework led to an increase in loss reserves proportionate to increases in delinquencies, the result would be an overreaction given the exceptional nature of these delinquencies.

FIGURE 3: FREDDIE MAC ACIS 2019 TRANSACTION TO-DATE PERFORMANCE

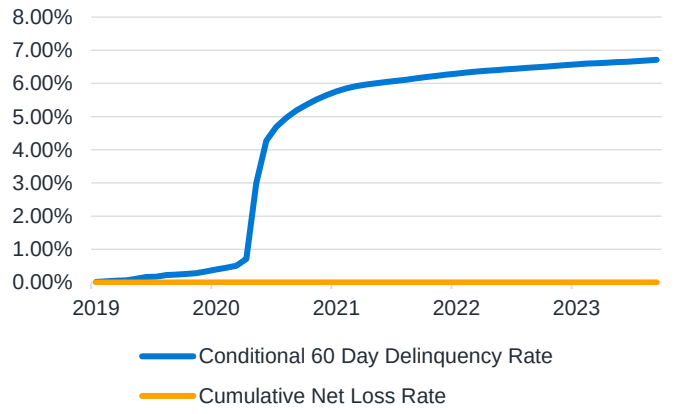

Review of potential alternative reserving methodologies

In this section, we outline other P&C reserving methods—and general principles and considerations—that could be applied in a reserving framework for mortgage reinsurance. The alternative approaches discussed in this section can be used to resolve some of the drawbacks identified with the PMI approach above.

In this section we discuss:

1. Loss development method
2. Expected loss ratio method
3. Dynamic and hybrid loss ratio methods

LOSS DEVELOPMENT METHOD

In general, loss development methods of reserving assume that historical loss development patterns can accurately predict future loss development patterns. As a result, historical data is summarized in development triangles, and claim development factors are selected. Essentially, experience to date coupled with the selected development factor to ultimate inform the total estimated losses.

We have previously published research on why the loss development method is not well-suited for mortgage credit exposures (Schmitz and Mrotek, 2010).[2] The below commentary summarizes some of the main drawbacks of this approach.

- Given that mortgage insurance is a long-tail line of business and the reinsurance programs—as they exist today—were established in 2013 or later (less than 10 years), there is not a large amount of data from this time period that illustrates the risks inherent in these exposures. Since 2013, mortgage performance has been extremely robust, leading to very little realized losses for almost all exposures to date (virtually all GSE CRT tranches have a 0% loss ratio to date).
- Historically—looking further back in history to analyze underlying mortgage performance—mortgage insurance claims experience has been varied, with many years of good experience and relatively fewer years with high levels of mortgage defaults and subsequent delays until ultimate resolution of claims. Specifically, during the 2007 global financial crisis (GFC), mortgage insurers experienced high numbers of claims (loss ratios of greater than 200%).

[2] Schmitz, M. & Mrotek, K. (2010). An Analysis of the Limitations of Utilizing the Development Method for Projecting Mortgage Credit Losses and Recommended Enhancements. Casualty Actuarial Society E-Forum, Vol. 2. Retrieved April 23, 2024, from https://www.casact.org/sites/default/files/database/forum_10fforumpt2_schmitz_mrotek.pdf?msclkid=12b563fdaba211ec99100bb73dc518c1.

These aspects of the business line make historical loss development methods relatively poor choices. Loss ratio methods—which often use sophisticated statistical models to control for various differences in the historical data over time—are likely better suited and discussed below.

EXPECTED LOSS RATIO METHOD

In general, loss ratio methods estimate expected claims as a percentage of the expected earned premium. A variety of techniques can be used to obtain the expected claims and premiums. The approach is intuitive, as when writing reinsurance—one expects a portion of the premiums received to be used to pay future losses. This approach uses that ratio to establish a loss reserve.

Over time, as paid losses develop, they are subtracted from the ultimate reserve estimate to obtain a future estimate. This method works well when the business line is not well suited for a development method—often due to lack of robust historical data or generalizable development trends (see discussion above). Potential drawbacks of this methodology can be its lack of sensitivity to the exposure performance over time—to the extent that performance is developing better or worse than the original forecast.

In theory, this method could work well when reserving for mortgage exposures. Given that mortgage is a long-tail line and actual performance is not known for many years, booking an expected loss ratio can yield a comprehensive view of the risk from day 1 of the exposure. Depending on how one calculates and defines the expected loss ratio, the value can also encapsulate the asymmetric distribution of mortgage performance. In particular, if a practitioner averages over a variety of scenarios, the expected value can be comprised of many outcomes of benign performance and relatively fewer years of more extreme, poor performance—which is the historical distribution of mortgage performance, as outlined above.

Given that there are many robust mortgage performance models commercially available (this includes Milliman's M-PIRe software), an expected loss ratio type of approach can be easily implemented using the mortgage forecasting tools readily available.

The main drawback of this type of approach is the lack of dynamic updates when actual performance begins to stray from expected performance. Given the distribution of mortgage performance described above, actual outcomes of an exposure will likely either be better, or significantly worse, than the expected value (given that the expected value is a blend of many years of the former and fewer years of the latter). Thus, it is important to have an adjustment mechanism to align to development of actual performance. The final category of methods, outlined below, offers potential solutions to this potential drawback.

DYNAMIC AND HYBRID LOSS RATIO METHODS

Following on from the discussion of the expected loss ratio method, one could improve on the static and potentially stale nature of such an approach by reperforming the estimation on a regular basis. For instance, if a practitioner was to "reprice" the exposure, given known premiums to date and known losses to date, then the approach could be adjusted to incorporate the most current data on a regular basis. For the sake of this discussion, we refer to this method as a "dynamic loss ratio method."

Although this method does not have any glaring deficiencies (as it fully accounts for the risks covered under the transactions as well as performance to date), some reinsurers may be cautious to implement this degree of dynamism. For example, if a particular transaction has an initial period of strong or benign performance, this approach would adjust reserves downward as the loss ratio was repriced. If the economy experienced a shock shortly thereafter, the loss ratio would likely be revised upward (potentially to a level higher than the initially expected loss ratio). Some reinsurers may prefer—out of a degree of prudency—that the initially expected loss ratio was not revised downward, in light of the early positive performance that may not be a credible indication of future performance. Some might use a static period until enough credible experience has occurred to reliably recalibrate the pricing.

Addressing the above valid potential concern, another potential method is to use a hybrid approach between the expected loss ratio method and the dynamic loss ratio method—which we will call the "hybrid loss ratio method." This method would generally consist of beginning on the expected loss ratio method and transitioning—at some period of time in the future—to the dynamic loss ratio method. The transition can be done over multiple periods by calculating a weighted average to ensure a smooth transition.

Hybrid reserving approaches have been developed in the past to incorporate the benefits of multiple approaches. For example, the Bornhuetter-Ferguson (BF) loss reserving technique was developed to blend the loss development method with the expected loss ratio method. The method uses an a priori expected loss ratio (ELR) and, as time progresses, assigns a lower and lower weight to the ELR. The weight is a function of the loss development factor from the current age to ultimate—i.e., as the factor is higher, the weight to the ELR is higher, and vice versa. We have previously outlined frameworks for modifying the BF method for use in mortgage (Schmitz and Mrotek, 2010). This underscores the point that hybrids between two distinct approaches have been used successfully in the past.

Implementing a comprehensive approach

After surveying alternative approaches, we generally find that a dynamic or hybrid loss ratio method that either directly uses repriced estimates representing the firm's best estimate of loss reserves, or uses some blended method with repriced results as well as a priori results. Both will work well to address the drawbacks identified with the PMI methodology and lead to a robust reserving framework.

In this section, we further discuss:

A. Methods to select loss ratios
B. Premium forecasting
C. Transition times and weights to repriced results
D. Frequency of repricing

METHODS TO SELECT LOSS RATIOS

Typically, industry practitioners use a loan-level mortgage performance model to forecast individual loan performance for the loans underlying the reinsurance transactions. After completing this step, the results from the model are aggregated and the reinsurance structure is overlaid on top of the forecasted loan performance.

Typically, mortgage performance can be forecast under either deterministic scenarios or a stochastic simulation. In deterministic scenarios, users specify the economic conditions they would like the model to reflect. House prices, interest rates, and unemployment rates are the key drivers of mortgage performance. Thus, for a constructed deterministic scenario, a user will define the future paths for each of these variables. The mortgage performance model—which is often a regression-based model using these three economic variables as a portion of the drivers—will produce performance estimates taking these factors into account.

Typically, selecting a single deterministic scenario to set loss ratios is a poor approach—it does not account for the asymmetry in the distribution of expected mortgage performance. However, weighting several deterministic scenarios and selecting the weighted-average loss ratio result is a method that others in the industry have employed. For instance, a practitioner may select the following weighting and scenarios based on its assessment of the odds of particular economic outlooks for home prices: 75% "Baseline," 10% "House Price Index (HPI) Down 10%," 5% "HPI Down 20%," 5% "HPI Down 30%," and 5% "HPI Down 40%" (five events with weights summing to 100%). An approach using deterministic scenarios has the benefit of being easily interpretable by a variety of stakeholders—as it is relatively easy to explain the underlying scenarios being weighted together.

Alternatively, most modeling approaches also feature a stochastic simulation of future mortgage performance. The major benefit to the stochastic approach is that many more data points (simulated trials) can be used to approximate the tail of the underlying mortgage performance distribution more granularly. Instead of weighting a small number of punitive scenarios to reflect the tail risks, the stochastic approach can generate hundreds (or thousands) of tail scenarios and ascribe probability weights to them. Reinsurers can select the expected value of the entire distribution, or a conditionally weighted portion of the distribution reflecting their outlook for future mortgage credit performance. We are aware of reinsurers using stochastic approaches when selecting expected loss ratios as well.

PREMIUM FORECASTING

Inherent in an ultimate loss ratio approach is an aspect of premium forecasting. GSE CRT transactions pricing is generally defined by an annual rate on line (ROL)—or rate on the loan balance of the pool—paid monthly.

Because premium is paid over time and as a function of the reinsurance limit outstanding, estimations of premium can differ based on modeling assumptions. In general, as underlying mortgage performance impacts both reinsurance premium and losses, best practice is to estimate premiums using the same methods as used when estimating losses—i.e., with the same scenarios and percentiles selected.

Because premium is written over time, and not as a single premium at contract inception, establishing a loss reserve recognizing the estimated ultimate losses would result in a mismatch to premium. Thus, applying expected loss ratios to earned premium to date to calculate current incurred losses addresses the mismatch between premium earnings and losses incurred. This will ensure incurred losses are established or built up proportionally to the earned premium (and loss reserves are calculated by subtracting cumulative paid losses to date from incurred losses).

TRANSITION TIMES AND WEIGHTS TO REPRICED RESULTS

When establishing a "hybrid" loss ratio approach, as described above, a key choice is the is the transition mechanism from the "expected" loss ratio to the "dynamic" loss ratio. This method can be defined via many different adjustment mechanisms, ranging from the slowest transition (i.e., no transition, which would be equivalent to the expected loss ratio method) to the fastest transition (i.e., immediately, using the repriced loss ratio, which would be equivalent to the dynamic loss ratio method).

Best practice will be informed by the exposures it is being applied to and a balance of the stability of the initial expected loss ratio against the responsiveness of the repriced loss ratio in light of the uncertainty surrounding the repriced loss ratio forecast. The term

of the transaction and the early termination options embedded in the transaction will have an impact on the expected time ”on risk.” These factors should be taken into account to ensure that the transition to repriced results occurs at appropriate times (i.e., in recognition of the credibility of the repriced loss ratio in order to balance stability and responsiveness).

Furthermore, the Freddie Mac ACIS program ceded multiple different tranches, each with different expected weighted average lives (WALs). The tranche limits pay down sequentially, with the M1 tranche being ”off-risk” before the M2 tranche’s limit is reduced at all. This behavior will lead to the different layers of the transactions being impacted by the transition methodology relatively differently.

In order for an M1 tranche to experience a loss, a considerable stress event must start almost immediately after transaction issuance. If six to 12 months pass without a stress to mortgage performance, and the M1 risk begins paying down, then most modeled scenarios reflecting M1 will not experience a loss. This dynamic may provide support for a faster transition to repricing for tranches similar to M1 in favorable scenarios with high mortgage prepayments.

On the opposite extreme, looking at the B2 tranche—where the limit amortizes slowest or last—a slow transition may be reasonable. B2 tranches typically do not reduce at all over the entire term of the transaction. Furthermore, the tranche has the lowest attachment point out of the tranches ceded. As a result, it likely takes a longer period of time to develop certainty about the B2 tranche outcome in favorable scenarios.

In summary, many different transition/repricing structures may be defensible. However, practitioners must be mindful of how these choices interplay with the exposures they are writing. If a reinsurer mainly wrote M1 business, it is likely that a different (quicker) transition approach would be preferred in favorable scenarios compared to a reinsurer that mainly wrote B2 business.

FREQUENCY OF REPRICING

Lastly, when using (or factoring in) a dynamic repriced loss ratio, one must select a frequency by which to reprice the risk. Typically, this is done quarterly, as that periodicity matches financial reporting periods.

However, if a reinsurer was concerned about potential “noise” caused by quarterly updates in its approach, it could reprice the loss ratio on an annual basis and simply apply the annually estimated loss ratio to quarterly updated earned premium to date to develop a loss reserve.

Conclusion

We consider the dynamic and hybrid loss ratio methods to be the most robust when reserving for reinsurers’ mortgage exposures. Given the advancement of econometric mortgage performance models and the detailed loan-level disclosures, loss ratio methods can provide tailored assessments of the risks, taking many factors into consideration (performance to date, loan characteristics, macroeconomics). Furthermore, we have witnessed and/or directly assisted clients with these types of approaches for mortgage reinsurance reserving in the past. However, these methodologies are not prescriptive in how they should be implemented. There are many details that must be considered and selected when implementing the approach on a particular set of exposures. As always, prior to implementation reinsurers should consult their accountants and auditors to ensure these treatments comply with any applicable accounting or auditing standards.

Milliman

Milliman is among the world’s largest providers of actuarial, risk management, and technology solutions. Our consulting and advanced analytics capabilities encompass healthcare, property & casualty insurance, life insurance and financial services, and employee benefits. Founded in 1947, Milliman is an independent firm with offices in major cities around the globe.

milliman.com

CONTACT

Andrew Netter
andrew.netter@milliman.com

Jonathan Glowacki
jonathan.glowacki@milliman.com

Michael Schmitz
michael.schmitz@milliman.com