A Ratings-Based Model for Credit Events in MakerDAO

Alex Evans
Placeholder

July 2019

Abstract

This article introduces a ratings-based model for credit events in the MakerDAO lending system. To account for stochastic loan terms, we extend conventional ratings-based models to include two absorbing states, full repayment and delinquency. This extension enables the exploration of the long-term behaviors of individual loans and global system states. The transition rates of the continuous-time Markov Chain are estimated using on-chain data for MakerDAO from Ethereum. The model can be used in credit risk management applications and for pricing derivatives with state-dependent payoffs such as MKR.

1 Introduction

We introduce a ratings-based Markov model for credit events in MakerDAO. Unlike fixed-term loans, Maker allows borrowers to repay ("Wipe") their loans including all accrued interest ("stability fees") at any time. Loans are also subject to liquidation ("Bite") by third-party "Keepers" should the borrower fail to maintain an appropriate loan-to-value ratio (this is akin to a margin call). An additional 10 percent "liquidation fee" on outstanding debt is levied on liquidated loans. The model accounts for stochastic loan terms by including repayment as an absorbing state alongside delinquency. The long-run behavior of each loan can be determined from the transition rates between states. With additional assumptions about birth rates and loan sizes, nearly any aggregate system parameter in Maker can be forecasted.

The model lends itself to convenient statistical estimation using on-chain data and, optionally, values for external covariates. The intended application of the framework is in risk management. The model can be used to compute the expected magnitude of credit exposure for a pool of loans for given covariate values. Both expected and unexpected losses can be computed with additional data on recovery values. The exploration of fee policy choices is also available to users of the model, including forecasting the anticipated effects of stability fee changes. A key feature of the model is the flexibility of its assumptions. Throughout this article, effort is taken to highlight cases where assumptions can be adapted for analysts to refine the model for their intended application.

The remainder of this article will proceed as follows. Section 2 outlines the mechanics and assumptions of the model and examines the long-term behavior of individual loans and global system state. Section 3 discusses strategies for estimating model parameters using on-chain data. Section 4 applies the model to risk management problems using the example of stress testing liquidations. Section 5 provides a brief concluding note on a potential pathway to pricing MKR.

2 The Model

This section introduces a ratings-based Markov model for credit events in the Maker system. Ratings-based credit models allow loans to be in a variety of states, representing different credit ratings, the lowest of which is default. The most popular formulation of this model is found in Jarrow, Lando, and Turnbull (1997), JLT for short. The model presented here draws from this approach, with some important adjustments due to the idiosyncrasies of the Maker system. Individual "Draw" transactions are treated as discrete loans, each of which is modeled according to a Markov Chain on a discrete state space.[1] Each loan is permitted to be in one of four states: Safe, Unsafe, Bitten, or Wiped.[2] Unsafe loans are below the minimum collateralization ratio and may be liquidated by Keepers. Upon creation, each loan begins in the Safe state with transition rates to other states modeled as exponential parameters, forming a continuous-time[3] Markov Chain with generator matrix Q.[4]

$$\mathbf{Q} = \begin{array}{c} \\ \text{Safe} \\ \text{Unsafe} \\ \text{Wiped} \\ \text{Bitten} \end{array} \begin{array}{c} \begin{array}{cccc} \text{Safe} & \text{Unsafe} & \text{Wiped} & \text{Bitten} \end{array} \\ \begin{pmatrix} -q_s & q_{su} & q_{sw} & q_{sb} \\ q_{us} & -q_u & q_{uw} & q_{wb} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \end{array} \tag{1}$$

where q_{ij} represents the instantaneous rate of transition from state i to state j. This assumes $q_{ij} \geq 0$ for $i \neq j$ and $q_i = -\sum_{j=1,j\neq i}^{4} q_{ij}$ for $i \in \{s, u, w, b\}$. q_i is the exponential parameter of the "holding time" at state i, or the time interval that the loan started at i spends in i before transitioning to any of the other three states. This is determined by the transition rates to adjacent states. The holding times for Wiped,[5] w, and Bitten, b, states are zero as these states are "absorbing." Ignoring absorbing states, we can think of the process as starting in state s where it stays for an exponentially distributed amount of time, $1/q_s$ on average. It then moves to a new state, i, other than s, with probability p_{si}. The off-diagonal entries in the generator matrix are $q_{ij} = q_i p_{ij}$. If the loan enters one of the absorbing states, it stays there for the remainder of time.

[1] Analysts may opt for a different unit of analysis than Draw transactions such as individual dai units or Collateralized Debt Positions (CDPs). There are trade-offs to different approaches. While providing lower resolution, CDPs may be more suitable units of analysis when studying default correlation. Smaller units may show multiple loans defaulting simultaneously when they are in fact part of the same CDP. One drawback is that the sizes of individual CDPs vary significantly across time, while modeling loans based on Draw transactions allows for at least a fixed maximum size. Focusing the analysis on individual dai units would enable higher resolution but would sacrifice computational efficiency.

[2] This state space is also at the discretion of the analyst. For example, the analyst may choose to expand the number of permitted states to include different collateralization levels for each loan. In section 4, the state space is contracted to three states to simplify the model for estimation.

[3] A discrete-time formulation is also valid and may better match the reality of blockchains where transactions occur in discrete chunks. The continuous-time formulation makes the tools of stochastic calculus available to the analyst.

[4] Informally, the Markov property states that the future evolution of the process only depends on the current state. For much of this analysis the homogeneity assumption is also maintained. Informally, this states that the future evolution does not depend on the starting time. More sophisticated analysts are encouraged to explore relaxing both assumptions.

[5] In the model used here, Wipes are recorded only when the entire loan is paid back. This choice is motivated by simplicity and is not necessarily realistic. The analyst may wish to relax this assumption to allow for the loan to be paid back over time.

The generator matrix depicts the instantaneous transition rates associated with the transition function, $\mathbf{P}(t)$, i.e. $\mathbf{P}'(0) = \mathbf{Q}$. This transition matrix is given by the matrix exponential

$$\mathbf{P}(t) = e^{t\mathbf{Q}} = \sum_{n=0}^{\infty} \frac{1}{n!}(t\mathbf{Q})^n \tag{2}$$

[see Dobrow (2016) p. 278)]. The entry $\mathbf{P}_{ij}(t)$ of the transition function gives the probability of transitioning from initial state i to j at t.

2.1 Long-Term Convergence

Loans in Maker do not have fixed terms but are instead either Wiped by their owners or Bitten by Keepers. Unlike most conventional loans, their term is stochastic. This requires some adjustments to traditional credit models such as JLT in order to examine long-term convergence. The long-term dynamics of each loan are given by the limiting distribution of its associated Markov Chain, defined as $\lim_{t\to\infty} \mathbf{P}_{ij}(t) = \pi_i$ for all states i and j. This distribution allows analysts to predict the long-term behavior of each loan in the system and forecast loan-specific risks. To compute the average time until a loan is absorbed in either a Wiped or Bitten state, we start by rewriting $\mathbf{Q}$ in block matrix form:

$$\mathbf{Q} = \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ * & \mathbf{V} \end{pmatrix}$$

where $\mathbf{V}$ is a 2×2 matrix on the set of transient states, s and u. Next, we define the fundamental matrix, $\mathbf{F}$, where the ij^{th} entry is the expected time, for a process started in i, that the process is in j until absorption. Therefore, $\mathbf{F} = -\mathbf{V}^{-1}$ [for a proof, see Dobrow (2016, p. 289)]. This is given by:

$$\mathbf{F} = -\begin{pmatrix} -q_s & q_{su} \\ q_{us} & -q_u \end{pmatrix}^{-1} = \frac{1}{q_s q_u - q_{su} q_{us}} \begin{pmatrix} q_u & q_{su} \\ q_{us} & q_s \end{pmatrix} \tag{3}$$

For the chain started in s, the mean time until absorption is given by $\mathbf{F}_{su} + \mathbf{F}_{ss}$, the sum of the first row of $\mathbf{F}$. This time is

$$\frac{1}{\kappa} = \frac{q_u + q_{su}}{q_s q_u - q_{su} q_{us}} \tag{4}$$

The long-run probability that a loan ends up Wiped or Bitten is easiest to compute using the embedded chain $\tilde{\mathbf{P}}$ of the loan Markov Chain with generator $\mathbf{Q}$[6]

$$\tilde{\mathbf{P}} = \begin{array}{l} \\ \text{Safe} \\ \text{Unsafe} \\ \text{Wiped} \\ \text{Bitten} \end{array} \begin{array}{c} \begin{array}{cccc} \text{Safe} & \text{Unsafe} & \text{Wiped} & \text{Bitten} \end{array} \\ \begin{pmatrix} 0 & p_{su} & p_{sw} & p_{sb} \\ p_{us} & 0 & p_{uw} & q_{wb} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array} \tag{5}$$

$\tilde{\mathbf{P}}$ is the transition matrix if one is to ignore time spent in each state. The p_{ij} entry denotes the probability that the next step in the process for a chain started in i is j. In block matrix form, the

[6] Technically, $\mathbf{Q}$ and $\tilde{\mathbf{P}}$ should not allow direction transitions from Safe to Bitten, as loans have to visit the Unsafe state first. These parameters are included in the model nonetheless in case the analyst considers longer discrete periods but can easily be replaced with zero otherwise.

embedded chain can be rewritten as

$$\tilde{\mathbf{P}} = \begin{pmatrix} \mathbf{S} & \mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix}$$

where $\mathbf{S}$ is the matrix of transitions between transient states, $\mathbf{R}$ is the matrix of transitions from transient to absorbing states, $\mathbf{0}$ is 2×2 matrix of zeros, and $\mathbf{I}$ is the 2×2 identity matrix. In n steps, the transition matrix is

$$\widetilde{\mathbf{P}^n} = \begin{pmatrix} \mathbf{S}^n & (\mathbf{I} + \mathbf{S} + \cdots + \mathbf{S}^{n-1})\mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \tag{6}$$

The limiting distribution is given by

$$\begin{aligned} \lim_{n\to\infty} \widetilde{\mathbf{P}^n} &= \lim_{n\to\infty} \begin{pmatrix} \mathbf{S}^n & (\mathbf{I} + \mathbf{S} + \cdots + \mathbf{S}^{n-1})\mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \\ &= \begin{pmatrix} \lim\limits_{n\to\infty} \mathbf{S}^n & \lim\limits_{n\to\infty} (\mathbf{I} + \mathbf{S} + \cdots + \mathbf{S}^{n-1})\mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{0} & (\mathbf{I} - \mathbf{S})^{-1}\mathbf{R} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \end{aligned} \tag{7}$$

Where the last equation follows the result from linear algebra that $\sum_{n=0}^{\infty} \mathbf{A}^n = (\mathbf{I} - \mathbf{A})^{-1}$ for a square matrix $\mathbf{A}$ that satisfies $\mathbf{A}^n \to 0$ as $n \to \infty$. The first row of the limiting submatrix $(\mathbf{I} - \mathbf{S})^{-1}\mathbf{R}$ gives the long-term probabilities of transitioning from a Safe state to either a Wiped or Bitten state.

$$\begin{aligned} (\mathbf{I} - \mathbf{S})^{-1} \mathbf{R} &= \begin{pmatrix} 1 & -p_{su} \\ -p_{us} & 1 \end{pmatrix}^{-1} \begin{pmatrix} p_{sw} & p_{sb} \\ p_{uw} & p_{ub} \end{pmatrix} \\ &= \frac{1}{1 - p_{su}p_{us}} \begin{pmatrix} p_{sw} + p_{su}p_{uw} & p_{sb} + p_{su}p_{ub} \\ p_{uw} + p_{us}p_{sw} & p_{ub} + p_{us}p_{sb} \end{pmatrix} \end{aligned} \tag{8}$$

This matrix derivation will be especially convenient if the analyst is expanding the state space, for example by including multiple "Safe" states representing different collateralization levels or segmenting "Bitten" into varying degrees of severity. Using this result, one can make predictions about the probability that a loan started in a Safe state will be Bitten (entry 1,2) or be Wiped (entry 1,1).

2.2 Investigating Global State

By estimating the entries of the generator matrix, analysts can forecast any loan level parameter in Maker. These parameters may include the mean time that a loan stays open, the probability that it is bitten, or the probability that it is wiped by the borrower. With additional assumptions about the birth process, the framework can be extended to estimate global system parameters, such as total dai outstanding, total stability fees accrued, and total liquidation fees.

A simple model for loan creation is to assume Draw transactions follow a birth process with Poisson parameter, λ. The number of loans open at time $t > 0$ is therefore determined by a birth and death process with deaths occurring at a rate κn, where κ is derived from (4) and n

is the number of loans outstanding in the system. Assuming the birth and death processes are independent, this forms a $M/M/\infty$ queue [see Kulkarni (2011) p. 210]. Starting from zero loans at $t = 0$, the expected number of loans in the system for any time t is given by

$$\frac{\lambda}{\kappa}\left(1 - e^{-\kappa t}\right) \tag{9}$$

As $t \to \infty$, the limiting distribution for the individuals in the system is simply a Poisson random variable with rate λ/κ. Thus, in steady state, the expected number of loans in the system is λ/κ. This feature of the model prevents the size of outstanding loans from growing without bound. With these parameters and additional assumptions for fees and loan sizes, the analyst can forecast nearly any global system parameter in Maker.

Example 3.2.1 Expected Liquidation and Stability Fees

Up to time t, the expected number of closed loans is

$$C_t = \lambda t - \frac{\lambda}{\kappa}\left(1 - e^{-\kappa t}\right) \tag{10}$$

From (8) we have the expected number of bitten loans is given by

$$C_t \frac{p_{sw} + p_{su}p_{uw}}{1 - p_{su}p_{us}} \tag{11}$$

Thus, the expected cumulative liquidation fees accrued by t is given by

$$fDC_t \frac{p_{sw} + p_{su}p_{uw}}{1 - p_{su}p_{us}} \tag{12}$$

where D is the expected loan size and f is the liquidation fee (assuming liquidation fees, loan sizes, and birth and death rates are independent). For large t, the expected lifespan of loans closed before t approaches the mean lifespan $1/\kappa$. For a fixed stability fee this would yield total expected stability fees

$$DC_t \int_0^\infty \left((1+r)^x - 1\right)\kappa e^{-\kappa x}\, dx \tag{13}$$

Note that Maker uses discrete, per-second compounding when calculating stability fees owed.

3 Estimating the Empirical Generator Matrix

The Markov model specified in Section 2 lends itself to convenient statistical estimation. Using on-chain data for state transitions of Maker loans, each off-diagonal element in the generator matrix can be estimated by

$$q_{ij} = \frac{N_{ij}(T)}{\int_0^T Y_i(s)\, ds} \tag{14}$$

[see Jarrow et al (1997) p. 504] where $N_{ij}(T)$ is the total number of transitions from i to j over the interval $[0, T]$ and $Y_i(s)$ is the number of loans in state i at time s. In other words, the per second transition rate from i to j is given by the number of transitions from state i to j divided by the total number of "loan seconds" spent in state i from 0 to T. Descipher has built a python tool for estimating the generator using (14).[7] D5 also has an implementation[8] that can separate estimates by month as a proxy for different collateral price regimes.

3.1 Empirical Generator with Covariates

Analysts may also allow the transition intensities to depend on a set of external covariates. In this section, the intensities are assumed to be given by $q_{ij}\left(z\left(t\right)\right)$, where $z\left(t\right)$ can be any time-dependent or individual-specific explanatory variable (or vector of covariates). This is particularly useful if the analyst is attempting to explain the effect of different risk parameters and collateral behaviors on overall system performance. Here we focus on the stability fee and collateralization ratio as covariates due to their simplicity, but nearly any well-motivated covariate choice is available to the analyst.

The approach to specifying the transition intensities is similar to that of survival analysis techniques used in disease modeling. To make the analysis tractable we focus on the transition to just Bitten and Wiped states. This reduces the generator to a 3×3 matrix with two absorbing states, termed "competing risks" in survival analysis. Safe and Unsafe states are merged into an "Open" state, denoted o. To estimate the effect of covariates on the transition rates, we use the msm package for the statistical software R. The documentation and capabilities of this package are described in Jackson (2019). The estimation used in msm is based on a proportional hazards model. In the formulation below, the transition rates to the Wiped and Bitten states are determined by separate covariates, but this choice is at the discretion of the analyst. Specifically, the transition from Open to Bitten is assumed to be given by

$$q_{ob}\left(c\left(t\right)\right) = q_{ob}^{0} e^{\beta_{ob}^{T}\left(c(t)\right)} \tag{15}$$

where q_{ob}^{0} is the baseline hazard and $c\left(t\right)$ is a factor for different levels of aggregate system collateralization. The factor is coded as "0" when Maker's aggregate collateralization ratio is above 280%, "1" when between 250% and 280%, "2" when below 250%. The particular ranges of these collateralization ratio "buckets" can be modified by the analyst as desired. The collateralization ratio is defined as the USD value of total locked collateral (in this case "Pooled ETH") divided by total outstanding debt. All data was obtained from on-chain transactions and oracle-submitted prices through June 24^{th} 2019 and is available to analysts in the accompanying repository.[9] For the transition rate from Open to Wiped we have

$$q_{ow}\left(i\left(t\right)\right) = q_{ow}^{0} e^{\beta_{ow}^{T}\left(r(t)\right)} \tag{16}$$

Where $r\left(t\right)$ represents the stability fee[10] at time t. The estimation of (15) and (16) is summarized for daily values in Table 1.[11]

[7] https://github.com/visavishesh/dai_markov_chains

[8] https://github.com/askeluv/dai-transitions/blob/master/notebooks/dai-state-transitions.ipynb

[9] github.com/ahevans/Maker-Risk-Model

[10] The data records the time when the stability fee changes were *enacted*, but the effect of stability fee change *announcements* may also be of interest.

[11] Note that this is imperfect analysis intended for demonstration and almost certainly unreliable for real-world forecasting. More experienced statisticians are encouraged to create their own models with the accompanying data.

Table 1: Estimation Results for Equations (15) and (16)

States	**Baseline**	**Stability Fee**	**Coll 250–280%**	**Coll <250%**
Open-Open	-0.010679 $(-0.010808, -0.010552)$			
Open-Wiped	0.008598 $(0.008485, 0.008712)$	72.39 $(59.76, 87.69)$	1.000	1.000
Open-Bitten	0.002081 $(0.002023, 0.002141)$	1.000	4.23 $(4.022, 4.448)$	14.29 $(13.38, 15.26)$
$-2 *$ log-likelihood:	360130.8			

In the baseline model, the transition rate is 0.0086 from Open to Wiped and 0.0021 to Bitten. Dividing by the holding time in the Open state gives the result that approximately 80.5% of loans are expected to be Wiped and 19.5% Bitten. The covariate estimates for Open to Wiped indicate that a 100% rise in the annual stability fee would increase the rate of CDP closures by a factor of 72.39 or approximately $e^{4.2821}$ (in such an event, the model predicts that nearly all existing loans would be closed within a few days). A 10% increase in the annual fee would increase the transition rate to the wiped state by $e^{0.42821}$ or 53.45%. The aggregate collateralization ratio also has a similarly significant effect on the rate at which loans are bitten, particularly when collateralization ratios are below 250%. The transition rate is 4.23 times higher when aggregate system collateralization is between 250% and 280% and 14.29 times higher when below 250%. The figures in parenthesis in Table 1 indicate 95% confidence intervals for the estimated values.

Plotting the survival curves for individual loans allows visual exploration of these effects. In Figure 1, different collateralization level assumptions result in noticeably altered survival curves when holding stability fees constant. Stability fees exhibit a similar pattern, where higher rates reduce the expected lifespan of each loan.

3.2 Estimating Loan Creation Parameters

As discussed in Section 2.2, certain assumptions about the loan creation process are necessary when forecasting global system parameters. If the analyst models the loan birth process using the Poisson approach discussed in Section 2.2, there are several options available for empirical estimation. First, the parameter can be set to its average in the dataset. Second, it can be assumed to depend on a set of fixed or time-varying covariates, $\lambda\left(z(t)\right)$ as with the entries of the generator matrix. In the case of time-varying covariates, the non-homogenous birth process would have mean $E\left[N_t\right] = \int_0^t \lambda\left(x\right)\, dx$ where N_t is the number of loans created before t. It is important to note that all the formulas presented in Section 2.2 rely on the birth and death processes being independent. Depending on the specification, the independence assumption can still hold conditionally for given values of the covariates. Finally, the birth parameter can be chosen such that the long-term debt in the system converges to a desired value based on the limit of (9) as $t \to \infty$. For example, with an expected loan size $D = 7K$ (the mean in the dataset) a birth rate of $\lambda = 382$ would predict that the single-collateral system stabilizes at 250M dai at long-run equilibrium.

An important variable related to the birth process is the relative distribution for loan sizes.

Figure 1: Loan Survival Functions

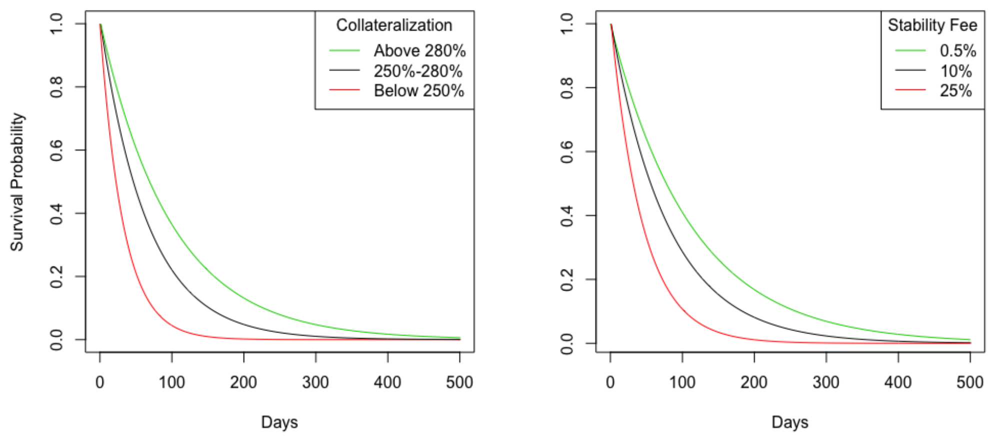

The simplest description of the loan size distribution is to assume an equally weighted homogenous portfolio where all loans are of size D. The analyst may instead specify a distribution for loan sizes. It is important to remember when the independence assumption applies as in (12) and (13). It may also be useful to include loan size as a covariate to investigate potential effects on transition rates.

4 Risk Management Applications

Analysts can use the estimation of the loan Markov Chain to stress test the resilience of the system under different scenarios. This can be applied to managing both portfolio credit risk and fee policy for Maker.

A straightforward application of the model outlined in section 3 is the analysis of credit exposure for a given pool of loans. A key concept in portfolio credit risk management and the pricing of multi-name credit derivatives is the loss distribution function. This function specifies the probability distribution for default levels in a given portfolio.[12] Informally, for a given portfolio, the loss distribution function assigns probabilities to different percentage default levels. Here, we focus on the conditional loss distribution function where default rates are assumed to be independent conditional on the values of the covariates.

In the 3×3 matrix specified in section 4, the mean time to absorption simplifies to $\kappa = q_o$ and the probability that the loan will end up bitten is simply $p_{ob} = q_{ob}/q_o$. The default probability is assumed to depend on collateralization ratio as specified in Section 3. Conditional on the values of the covariates, individual bites are pairwise uncorrelated. As a result, for a given number of loans, I, the number of bites in the portfolio follows a binomial distribution with success probability p_{ob}

[12] For an introduction to loss distribution functions and default correlations in portfolio credit risk modeling see chapter 21 in Bomfim (2016).

Figure 2: Conditional Loss Distribution

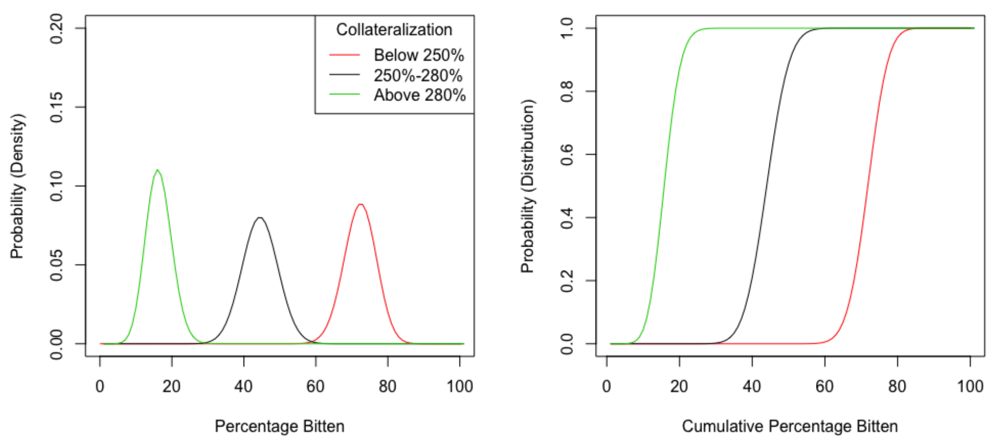

[see Bombfim (2016) p. 239]. If we let L denote the percentage of bites in the portfolio and N the number of bites, the conditional probability of a given percentage of loans getting liquidated is given by

$$P\left(L = \frac{N}{I} \;\middle|\; c\right) = \frac{I!}{n!\,(I-n)!} p_{ob}\left(c\right)^{n} \left(1 - p_{ob}(c)\right)\right)^{I-n} \tag{17}$$

Figure 2 illustrates the use of the model in assessing loss probabilities in a portfolio of 100 loans under different scenarios involving the collateralization ratio covariate (i.e. scenario stress testing). The left side of Figure 2 shows the probability of conditional losses for different collateralization levels. In the scenario where collateralization remains above the 280% threshold, the probability that less than 5% of the portfolio will be bitten is small as is the probability of bites above 25%. In contrast, when assuming collateralization levels below 250% the probability is nearly zero for small losses and very high for large losses. The impact is more clearly shown in the right-hand side which depicts the probability that bites will be less than a given percentage of the loan portfolio. In the case where the collateralization ratio is below 250%, large losses are a near certainty with more than 99.5% probability at least 65% of loans the portfolio will be bitten. It is important to note the risks associated with higher variance of default percentages, which is highest when collateralization is between 250%–275%. The variance is important when analyzing unexpected losses in the portfolio or credit value at risk.

While these numbers are exclusively for demonstration, identifying high-risk scenarios like these is a key function of risk management. The next task for the analyst would be to identify expected losses from undercollateralized bites. This requires additional assumptions about the percentage of liquidated CDPs that will be bitten below 100% collateralization and their expected recovery values. Collateralization ratios at bite are included in the accompanying dataset for analysts who wish to explore these effects. The calculation of value bitten from the loss distribution function necessitates further assumptions about the birth process and loan sizes. Finally, the portfolio can be segmented into separate collateral types as multi-collateral dai is activated. One approach to studying portfolio credit risk in a multi-collateral system would be to include collateral types and associated parameters as covariates.

In Figure 2, the covariates values are assumed constant for the observation period, but this need not be the case. More sophisticated models may allow covariates to vary deterministically or may specify a stochastic process for their evolution. The latter approach would allow the analyst to specify a probability distribution over future covariate values and use it compute a probability-weighted average of conditional Bite probabilities. This would yield the unconditional loss distribution which offers a more complete account of default correlations. The case of allowing covariates to vary across time would produce non-homogenous transition rates. Allowing these to vary stochastically over time produces a so-called doubly stochastic model. While this exercise is beyond the scope of this work, numerous stochastic processes for interest rates and prices from mathematical finance are available to analysts in modeling stability fees, collateral prices, and other covariates of choice.

5 Future Work: Note on Pricing

One of the most appealing features of JLT is that it can efficiently handle the valuation of credit derivatives whose payouts are state-dependent. This is particularly interesting in the case of MKR as loans have different payouts to MKR holders depending on whether they are closed in the Safe or Bitten states. As a thought experiment, one might think of MKR as a Portfolio Default Swap written on the set of outstanding loans at any future time. Applying JLT would require computing the "risk-neutral" generator matrix from observed market prices for comparable loans and the empirical generator matrix in (1) using the recursive method described in Jarrow et al. For now, this remains a theoretical exercise as no directly comparable products to CDP loans are traded from which appropriate risk premia can be computed.

References

[1] Bomfim, A. (2016). *Understanding Credit Derivatives and Related Instruments.* Waltham, MA: AP.

[2] Dobrow, R. (2016). *Introduction to Stochastic Processes with R.* Hoboken, NJ: Wiley.

[3] Jackson, C. (2019). *Multi-state modelling with R: the msm package.*

[4] Jarrow, R. and Lando, D. and Turnbull, S. (1997) *A Markov Model for the Term Structure of Credit Risk Spreads.* Review of Financial Studies, Vol. 10, No. 2.

[5] Kulkarni, V.G. (2011). *Introduction to Modeling and Analysis of Stochastic Systems.* New York, NY: Springer.