International Journal of Advanced Statistics and Probability, 12 (2) (2025) 40-47

SPC

International Journal of Advanced Statistics and Probability

Website: www.sciencepubco.com/index.php/IJASP
https://doi.org/10.14419/zbqm4d90
Research paper

Hidden Markov Models for Analyzing U.S. Bank Failures

Sunil K Sapra *

California State University, Los Angeles
**Corresponding author E-mail: ssapra@calstatela.edu*

Received: September 28, 2025, Accepted: November 3, 2025, Published: November 7, 2025

Abstract

This paper examines the use of Hidden Markov Models (HMMs) to investigate patterns in US bank failures. HMMs are statistical tools designed for handling sequential data and uncovering hidden states, making them particularly useful for studying systemic financial events. The paper provides an application of HMMs to U.S. bank failures data and presents results of an empirical study that analyzes historical data on banking crises in the United States. Four HMMs are employed for the analysis of U.S. bank failures data, and their relative performance is studied with respect to model fitting and forecasting. Empirical results are presented for each of the HMMs employed. Global decoding is employed to predict the most likely state sequence of large bank failure events via the Viterbi algorithm. The findings demonstrate the ability of HMMs to uncover unobservable economic conditions and enhance predictive capabilities for financial stability evaluations.

Keywords: *Hidden Markov Models; State Space Models; Regime Switching Models; Poisson HMM; EM Algorithm; Viterbi Algorithm.*

1. Introduction

Bank failures are critical events that disrupt financial systems and economies. Understanding their patterns and causes is vital for developing effective regulatory frameworks and risk management strategies. Hidden Markov Models (HMMs) offer a powerful method for analyzing time-series data where the driving processes are not directly observable. These models can capture regime-switching behavior, such as switching between expansion and recession cycles in economic data. Rabiner [6] provides an early, insightful account of HMMs focused mainly on applications to speech recognition. Zucchini et al. [9] illustrate the enormous flexibility of HMMs and their applications to a wide range of time series, including continuous-valued, circular, multivariate, binary, bounded, and unbounded counts, and categorical observations.

In this study, we utilize HMMs to analyze the dynamics of bank failures and assess their potential to identify underlying economic conditions. The financial system experiences cycles of stability and distress influenced by macroeconomic conditions, regulatory policies, and sector-specific shocks. These dynamics often lead to patterns of bank failures, which are inherently sequential and shaped by unobserved factors such as market confidence and systemic risks.

This paper aims to introduce the theoretical basis of Hidden Markov Models (HMM), demonstrate their application to the analysis of bank failure data in the United States, and highlight the potential of HMMs in predicting periods of financial instability. The paper is organized as follows. Section 2 describes the theoretical framework of HMMs, section 3 describes U.S. data on bank failures, unemployment rates, recessions, and federal funds rates in the United States utilized in this study, section 4 presents four 2-state HMMs employed for the analysis of U.S. bank failures data, section 5 presents empirical results and their interpretations for all of the HMMs employed, section 6 compares the HMMs in terms of their fit to the data for selecting the best model using the model selection criteria, section 7 presents the predictive out-of-sample performance comparison of the models, and section 8 concludes.

2. Hidden Markov Models: Theoretical Framework

An HMM is a statistical model that captures systems driven by hidden states, with observable outputs influenced by these states, and is defined by the following components:

- m: Number of hidden states in a Markov chain, such as economic stability or distress.
- C(t): Hidden state occupied by the Markov chain at time t, such as economic stability or distress.
- X(t): Observable data generated by the hidden states.
- T: Number of possible observations,
- Γ: State transition probability matrix (m x m), where $\gamma_{ij} = P(C(t+1) = j \mid C(t) = i)$ are transition probabilities of moving from one hidden state i to another hidden state j.

- P(x): Emission probability matrix with ith diagonal element $p_i(x) = P(X(t)= x \mid C(t) = i)$, the probability of observing specific data given a hidden state.
- δ: Initial state distribution (1 x m), where $\delta_i = P(C(1) = i)$, i=1,2,…,m.

An HMM assumes that the distribution of some observed variable(s) is driven by an unobserved state, i.e., each state gives rise to a different distribution for the observed variable. There are, therefore, two model components:

1) a model for the hidden state process {C(t), t = 1,2,…}, which is typically assumed to be a first-order Markov process; and
2) a model for the observation process {X(t): t ϵ N}, which may be multivariate. This model describes the distribution of the observation, conditional on the state process.

Zucchini et al. [9] describe a hidden Markov model (HMM) {X(t): t ϵ N} as a dependent mixture model consisting of two stochastic processes: (i) a hidden (unobserved) state process (parameter process) {C(t), t = 1,2,…} satisfying the Markov property and (ii) the observed state dependent process {X(t), t=1,2,..} for which the distribution of X(t) depends only on the current state C(t) and not on the previous states or observations. X(t) is defined as an m-state HMM if the Markov chain C(t) has m states. Let $X^{(t)}$ = (X(1), X(2),…, X(t)) and $C^{(t)}$ = (C(1), C(2),…, C(t)) represent respectively the histories of the observed and state processes from time 1 to time t. The simplest hidden Markov model (HMM) can be represented by the following equations (Zucchini et al. [9], pp. 36-38):

Hidden State Process: First-order Markov Chain

The hidden (unobserved) state process{C(t), t = 1, 2,…} satisfies the Markov property

$$P\left(C(t) \middle| \mathbf{C}^{(t-1)}\right) = P\left(C(t) \middle| C(t-1)\right), t = 2,3, \ldots, \quad (1)$$

Parametrized in terms of an (mxm) transition matrix with row sums equal to 1:

$$\mathbf{\Gamma} = \begin{pmatrix} \gamma_{11} & \cdots & \gamma_{1m} \\ \vdots & \vdots & \vdots \\ \gamma_{m1} & \cdots & \gamma_{mm} \end{pmatrix},$$

where $\gamma_{ij} = P(C(s+t) = j|C(s)=i)$. The specification of the Markov chain also requires an (mx1) vector of m initial probabilities $\delta^{(1)}$ = (P(C(1)) = 1, P(C(1)) = 2, …, P(C(1)) = m), with sum equal to 1.

Observation Process: At each time step, X(t) depends on the current state C(t) through the state-dependent distribution

$$P\left(X(t) \middle| X^{(t-1)}, C^{(t)}\right) = P\left(X(t) \middle| C(t)\right), t \epsilon N. \quad (2)$$

Given the HMM characterized by the parameter vector, **λ** = (**Γ**, **P(x)**, **δ**), HMMs are solved using algorithms such as:

- **Forward-Backward Algorithm:** For state probability estimation, since direct evaluation is $O(m^T)$. The recursive nature of the likelihood function evaluation via forward and backward probabilities $\alpha_t(i)$ and $\beta_t(i)$ is computationally more efficient than direct maximization over all state sequences:

Forward Algorithm: $\alpha_t(i) = P(x(1), x(2), ..., x(t), c(t) = i \mid \lambda)$

Backward Algorithm: $\beta_t(i) = P(x(t+1), x(t+2), ..., x(T) \mid c(t) = i, \lambda)$

- **Viterbi Algorithm:** To identify the most probable sequence of hidden states and find the most likely hidden state sequence, C* = argmax P(C(T) = c(T)| X(T) = x(T)).

The likelihood function of an HMM is given by

$$L = P(X^{(T)} = x(T)) = \delta P(x(1))\Gamma P(x(2)) \ldots \Gamma P\big(x(T)\big)1', \quad (3)$$

Where δ is the initial distribution (that of C(1)) and P(x) is the (mxm) diagonal matrix with ith diagonal element equal to the state-dependent probability or density $p_i(x)$ (Zucchini et al. [9]). In principle, we can therefore compute $L = \alpha_T 1'$ recursively as follows:

$$\alpha_1 = \delta P(x_1)$$

and

$$\alpha_t = \alpha_{t-1}\Gamma P(x_t) \text{ for } t = 2,3,\ldots,T.$$

An alternative to direct maximization of the likelihood function for obtaining the maximum likelihood estimates of the HMM parameters is the EM algorithm (Dempster et al. [1]), which is known as the Baum-Welch algorithm for HMMs. The EM algorithm makes use of forward and backward probabilities α_t and β_t.

3. Data Description and Variable Definitions

The dataset analyzed in this study consists of data on US bank failure counts and macroeconomic indicators, including unemployment rate, interest rates, and recessions spanning the time period 2001-2024. Data on US bank failure counts are taken from the Federal Deposit Insurance Corporation (FDIC [10]) and publicly available economic databases. Data on Federal Funds Effective Rate (RIFSPFFNA) for this time period are retrieved from the Federal Reserve Bank of St. Louis (FRED [11]) website and are seasonally unadjusted. The data on UNPCT (US unemployment rate) are from Investopedia (Rief [7]) and the US Bureau of Labor Statistics (BLS [12]) and are seasonally unadjusted as well. Historic unemployment rates for the month of December in years dating back to 2001 are used as the annual unemployment rate (UNPCT) for each year. Data on RECESSION (US recessions) for this time period are retrieved from the Federal Reserve Bank of St. Louis (FRED [11]) website.

The variables are defined as follows:

BANK FAILURES = Annual number of bank failures in the US
UNPCT = Annual unemployment percentage rate
RIFSPFFNA = Annual Federal Funds Effective Rate
RECESSION = 1 if the US economy is determined to be in recession in a year based on Hamilton's GDP-based recession indicator index and 0 if this is not the case.
The following table summarizes the data.

Data Summary

Variable	BANK FAILURES	UNPCT	RIFSPFFNA	RECESSION
Mean	23.6666667	5.64583333	1.747917	0.125
Std Error	8.90584132	0.37377719	0.374877	0.06896
Median	4.5	5.2	1.24	0
Mode	0	5.7	#N/A	0
Standard Deviation	43.6295339	1.8311268	1.836516	0.337832
Variance	1903.53623	3.35302536	3.372791	0.11413
Kurtosis	4.74160736	0.05505442	-0.49769	4.210266
Skewness	2.36257411	0.92558379	0.938817	2.42186
Range	157	6.4	5.25	1
Minimum	0	3.5	0.08	0
Maximum	157	9.9	5.33	1
Sum	568	135.5	41.95	3
Sample Size	24	24	24	24

4. Methodology

4.1 Model specification

Hidden State Process
We modeled bank failure counts as a two-state HMM with either no covariates or a single covariate, assuming that the observed number of bank failures arises from one of two distributions, depending on the underlying state. The hidden states are the following.

- State 1 (Distress): Marked by significant spikes in failures, often coinciding with recessions or financial crises.
- State 2 (Stable): Periods of low bank failures correlated with strong economic indicators.

Observation model
The observation process was assumed to follow a Poisson distribution, suitable for count data. Transition probabilities between states were estimated from the time series data.
Here, the response variable, BANK FAILURES, is a count variable, so we use a Poisson distribution in each state. That is, we assume that $\{X(t)|C(t) = j\} \sim \text{Poisson}(\lambda_j)$, where $\lambda_j > 0$, $j = 1, 2$.
Hidden Markov models for U.S. Bank Failures
An HMM is a combination of the hidden state and the observation models, and so an HMM object is created using these two components. In cases where the model includes covariates, i.e., the parameters are time-varying, these initial values correspond to the first row of the data set (Michelot [4] and [5]). We study the following hidden Markov models for bank failures in the United States for the period 2001-2024.

HMM 1: $\ln(\lambda_j) = \beta_0^{(j)}$,

HMM 2: $\ln(\lambda_j) = \beta_0^{(j)} + \beta_1^{(j)}\text{RIFSPFFNA}$,

HMM 3: $\ln(\lambda_j) = \beta_0^{(j)} + \beta_1^{(j)}\text{UNPCT}$,

HMM 4: $\ln(\lambda_j) = \beta_0^{(j)} + \beta_1^{(j)}\text{RECESSION}$,

Where λ_j is the average number of bank failures, $\beta_0^{(j)}$ Is the intercept, and $\beta_1^{(j)}$ Is the slope for the covariate in state $j \in \{1, 2\}$.
Training and Validation

- **Training:** The Baum-Welch algorithm was employed to estimate the model's parameters.
- **Validation:** The model's predictive accuracy was tested using out-of-sample likelihood measures and forecasting exercises.

5. Fitting The HMMs to The U.S. Bank Failure Data: Empirical Results

We employed four Poisson HMMs, each with two hidden states. These models included models with no covariates as well as models with a single covariate. The results for these models and comparisons among them are discussed next.

5.1. Identified states and economic insights

Two-state Poisson HMMs
The HMMs effectively differentiated between two latent states: State 1 (Distress) and State 2 (Stable), defined above.
Four hidden Markov models (HMM 1-HMM 4) were fitted to the data on US bank failures from 2001 through 2024. R packages hmmTMB (Michelot [4] and [5]) and MSwM (Sanchez-Espigares and Lopez-Moreno [8]) were employed for all computations in this paper. These models are compared using their marginal AICs for model fit and selection in section 6 and RMSEs on test data for predictive performance

in section 7. Below, we present the results of fitting these HMMs and the prediction of most likely states by each HMM via the Viterbi algorithm.

Table 1: Maximum Likelihood Estimates for HMM 1: Two-State Poisson HMM with no Covariates

	Mean	Std Error	z-statistic
Obs Rate: State 1	4.700479	0.04767315	98.598*
Obs Rate: State 2	1.856301	0.08838819	21.0017*
S1>S2.(Intercept)	-1.098458	1.15465591	-0.9612
S2>S1.(Intercept)	-2.890198.	1.02732246	-2.8133

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
log L = -157.902 (df=5)
Stationary Probabilities: δ_0: State 1: 0.17392, State 2: 0.82608
Maximum gradient component: 0.0004626072
Marginal AIC: State 1: 319.2049, State 2: 318.7751

As expected, the average number of bank failures in state 1 is $\lambda_1 = e^{4.700479} = 109.9978$, which is significantly higher than that in state 2, for which the average is $\lambda_2 = e^{1.856301} = 6.40$. The negative intercepts in the transition probability are not problematic, as these are inverted to derive the transition probability matrix from the estimated link function.

Estimated Transition probability matrix (tpm) with confidence intervals

The transition probabilities are estimated by modeling them with multinomial logit functions of parameters. This ensures that these estimated probabilities fall within the (0,1) interval. The estimated transition probability matrix (TPM) for HMM 1 is

$$\Gamma = \begin{pmatrix} 0.7069 & 0.2931 \\ 0.0782 & 0.9218 \end{pmatrix}.$$

The transition probabilities of staying in states 1 and 2 are high, which means that the process is more likely to stay in its current state than to transition to an alternative state. Furthermore, the probability of staying in state 2 (stable) is significantly higher relative to state 1 (distress). The probability of transitioning from state 1 to state 2 is significantly higher than the probability of transitioning from state 2 to state 1. As such, the process is more likely to transition to state 2 from state 1 and to stay in state 2, which is consistent with our intuition. To assess the precision of the transition probabilities, we present the lower and upper confidence limits for a 95% confidence interval.
The lower 95% confidence limit matrix is

$$\Gamma_l = \begin{pmatrix} 0.2520 & 0.0338 \\ 0.0073 & 0.6954 \end{pmatrix}$$

And the upper 95% confidence limit matrix is

$$\Gamma_u = \begin{pmatrix} 0.9662 & 0.7480 \\ 0.3046 & 0.9927 \end{pmatrix}$$

Stationary distribution

The stationary distribution is (0.17392, 0.82608).

Most likely state sequence

The most likely state sequence based on the Viterbi algorithm is 2 2 2 2 2 2 2 2 2 1 1 1 1 2 2 2 2 2 2 2 2 2 2 2 2 2. Occupancy of these states and transitions between them were consistent with the 2008 Global Financial Crisis, but the short-lived COVID-19 recession of 2020 was missed.

Table 2: Maximum Likelihood Estimates for HMM 2: Two-State Poisson HMM with A Covariate: Federal Funds Rates (RIFSPFFNA)

	Hidden State 1		
	Estimate	Std Error	z-statistic
Intercept	3.3728282	0.26975611	12.5033*
RIFSPFFNA	8.9263898	1.74361331	5.1194*
S1>S2.(Intercept)	-1.0973093	1.15432466	-0.9506
	Hidden State 2		
	Estimate	Std Error	z-statistic
Intercept	2.3354306	0.12008801	19.4477*
RIFSPFFNA	-0.2894201	0.06233476	-4.6430*
S2>S1.(Intercept)	-2.8898489	1.02716198	2.8134

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
log Likelihood = -130.9785 (df=7)
Stationary Probabilities: δ_0: State 1: 0.1738438, State 2: 0.8261562.
Maximum gradient component: 0.003454295
Marginal AIC: State 1: 277.3536, State 2: 275.9569

Conditional means

The conditional means for states 1 and 2 are given, respectively, by

$$\ln_t \lambda_1 = 3.3728282 + 8.9263898 \text{ RIFSPFFNA}_t,$$

$$\ln_t \lambda_2 = 2.3354306 - 0.2894201 \text{ RIFSPFFNA}_t,$$

Where t is the year. The positive sign of the coefficient estimate of RIFSPFFNA in state 1 (distress) and its negative sign in state 2 (stable) suggest that an increase in interest rates is likely to result in more bank failures on average due to the higher risk for loan defaults in state 1 and a moderate risk in state 2 when the economy is stable. As expected, the predicted (average) number of bank failures for the average

RIFSPFFNA = 1.74% in state 1 is $\lambda_1 = e^{18.90424}$, which is significantly higher than those in state 2, for which the predicted (average) is $\lambda_2 = e^{1.832} = 6.246$ failures.

Estimated transition probabilities (t = 1): The estimated transition probability matrix (tpm) is

$$\Gamma = \begin{pmatrix} 0.7497 & 0.2503 \\ 0.0527 & 0.9473 \end{pmatrix}$$

The transition probabilities of staying in states 1 and 2 under HMM 2 are higher relative to HMM 1, which means that the process is even more likely to stay in its current state than to transition to an alternative state under HMM 2. Furthermore, the probability of staying in state 2 (stable) is significantly higher relative to state 1 (distress) under HMM 2. Additionally, the probabilities of transitioning from state 1 (distress) to state 2 (stable) as well as the probability of transitioning from state 2 to state 1 are lower under HMM 2 relative to HMM 1. Nevertheless, the probability of transitioning from state 1 to state 2 is higher than the probability of transitioning from state 2 to state 1, implying that the process is more likely to transition to state 2 and to stay in state 2 when interest rates are relatively low.

To assess the precision of the transition probabilities, we present the lower and upper confidence limits for a 95% confidence interval.

The lower 95% confidence limit matrix is

$$\Gamma_l = \begin{pmatrix} 0.2520 & 0.0338 \\ 0.0073 & 0.6954 \end{pmatrix}$$

And the upper 95% confidence limit matrix is

$$\Gamma_u = \begin{pmatrix} 0.9662 & 0.7480 \\ 0.3046 & 0.9927 \end{pmatrix}.$$

Stationary distribution

The stationary distribution is $\boldsymbol{\delta_0}$ = (0.1738438, 0.8261562).

Most likely state sequence

The most likely state sequence based on the Viterbi algorithm is 2 2 2 2 2 2 2 2 2 1 1 1 1 2 2 2 2 2 2 2 2 2 2 2 2 2, which is identical to the sequence under HMM 1. As under HMM1, occupancy of these states and transitions between them under HMM2 were consistent with historical events, such as the 2008 Global Financial Crisis, but the short-lived COVID-19 recession of 2020 was missed.

Table 3: Maximum Likelihood Estimates for HMM 3: Two-State Poisson HMM with A Covariate: Unemployment Rate (UNPCT)

	Hidden State 1		
	Estimate	Std Error	z-statistic
Intercept	-1.2368901	0.32839171	-3.7665*
UNPCT	0.6736548	0.03918265	17.1927*
S1>S2.(Intercept)	-1.4409678	1.15774025	1.2446
	Hidden State 2		
	Estimate	Std Error	z-statistic
Intercept	-3.3760441	0.49846236	-6.7729*
UNPCT	0.8425761	0.05334072	15.7961*
S2>S1.(Intercept)	-1.3219394	0.95425136	-1.3853

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
log Likelihood = -68.63286 (df=7)
Stationary Probabilities: δ_0: State 1: 0.523763, State 2: 0.476237
Maximum gradient component: 0.1317196
Marginal AIC: State 1: 171.4028, State 2: 151.2657

Conditional means

The conditional means for states 1 and 2 are given, respectively, by

$$\ln_t\lambda_1 = -1.2368901 + 0.6736548\ \mathrm{UNPCT_t},$$

$$\ln_t\lambda_2 = -3.3760441 + 0.8425761\ \mathrm{UNPCT_t},$$

Where t is the year. The positive coefficient estimate of UNPCT in both states suggests that an increase in the unemployment rate is likely to result in more bank failures due to the higher potential for loan defaults and stress, regardless of the hidden state of the bank. As expected, the predicted (average) number of bank failures for the average UNPCT = 5.65% in state 1 is $\lambda_1 = e^{2.569} = 13.0527$, which is significantly higher than those in state 2, for which the predicted (average) is $\lambda_2 = e^{1.3845} = 3.9928$ failures.

Estimated Transition probability matrix with confidence intervals

The estimated TPM and its lower and upper confidence limits are

$$\Gamma = \begin{pmatrix} 0.7579 & 0.2421 \\ 0.2445 & 0.7555 \end{pmatrix}$$

with a lower 95% confidence limit matrix

$$\Gamma_l = \begin{pmatrix} 0.3063 & 0.0238 \\ 0.0331 & 0.3797 \end{pmatrix}$$

and the upper 95% confidence limit matrix

$$\Gamma_u = \begin{pmatrix} 0.9762 & 0.6937 \\ 0.6203 & 0.9669 \end{pmatrix}.$$

As with HMM 2, the transition probabilities of staying in states 1 and 2 are high, which means that the process is more likely to stay in its current state than to transition to an alternative state. Nevertheless, in contrast with HMM 2, under HMM 3, the probability of staying in either state is about the same. The probability of transitioning from state 1 to state 2 is also very close to the probability of transitioning from state 2 to state 1. As such, the process is no more likely to transition to state 2 from state 1 than to transition from state 1 to state 2 under HMM 3.

Stationary distribution

The stationary distribution is $\boldsymbol{\delta_0}$ = (0.523763, 0.476237).

Most likely state sequence

The most likely state sequence based on Viterbi algorithm is 2 2 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 2 1 2 2 2 1 1 and differs somewhat from the sequences under HMM 1 and HMM 2. Nevertheless, the occupancy of these states and transitions between them under HMM 3 would have predicted the Global Financial Crisis in 2009 and a more prolonged crisis than HMM 1 and HMM 2, which is slightly at odds with real-world events of this period.

Table 4: Maximum Likelihood Estimates for HMM 4: Two-State Poisson HMM with One Covariate: RECESSION

	Hidden State 1		
	Estimate	Std Error	z-statistic
Intercept	4.6051457	0.05773573	79.7625*
RECESSION	0.3365253	0.10235265	3.2879*
S1>S2.(Intercept)	-1.0967543	1.15416470	- 0.9503
	Hidden State 2		
	Estimate	Std Error	z-statistic
Intercept	1.7047468	0.10050384	16.9620*
RECESSION	0.9694538	0.21114445	4.5914*
S2>S1.(Intercept)	-2.8897910	1.02713539	-2.8134

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
log Likelihood = -144.0189 (df=7)
Stationary probabilities: δ_0: State 1: 0.173792, State 2: 0.826208
Maximum gradient component: 0.004010065
Marginal AIC: State 1: 277.3536, State 2: 275.9569

Conditional means

The conditional means for states 1 and 2 are given, respectively, by

$$\ln_t\lambda_1 = 4.6051457 + 0.3365253\ \mathrm{RECESSION}_t,$$

$$\ln_t\lambda_2 = -3.3760441 + 0.8425761\ \mathrm{RECESSION}_t,$$

Where t is the year. The positive coefficient estimate of RECESSION in both states suggests that a recession, whether it is severe under state 1 (distress) or mild under state 2 (stable), is likely to result in more bank failures due to the higher potential for loan defaults and stress, regardless of the hidden state of the bank. Furthermore, the predicted (average) number of bank failures for the average RECESSION = 1 in state 1 is $\lambda_1 = e^{4.942} = 140.05$, which is significantly higher than those in state 2, for which the predicted (average) number is $\lambda_2 = e^{2.674} = 14.50$ failures.

Transition probability matrix with confidence intervals

The estimated TPM and its lower and upper confidence limits are

$$\Gamma = \begin{pmatrix} 0.6988 & 0.3012 \\ 0.0784 & 0.9216 \end{pmatrix}$$

with a lower 95% confidence limit matrix

$$\Gamma_l = \begin{pmatrix} 0.2702 & 0.0285 \\ 0.0068 & 0.6870 \end{pmatrix}$$

and the upper 95% confidence limit matrix

$$\Gamma_u = \begin{pmatrix} 0.9715 & 0.7298 \\ 0.3130 & 0.9932 \end{pmatrix}$$

In contrast with HMM 3, under HMM 4, the probability of staying in state 2 is much higher than the probability of staying in state 1. The probability of transitioning from state 1 to state 2 is much higher than the probability of transitioning from state 2 to state 1. As such, the process is more likely to transition to state 2 (stable) from state 1 (distress) than to transition from state 1 to state 2. This is consistent with the data and historical events, such as the 2008 Global Financial Crisis.

Stationary distribution

The stationary distribution is $\boldsymbol{\delta_0}$ = (0.173792, 0.826208).

Most likely state sequence

The most likely state sequence based on the Viterbi algorithm is 2 2 2 2 2 2 2 2 2 2 2 2 2 1 1 1 1 1 2 2 2 2 1 1 1 under HMM 4 and is consistent with the sequence under HMM 3. Both HMM 3 and HMM 4 would have predicted the large bank failure events resulting from the Global Financial Crisis in 2009 instead of 2008, and a more prolonged crisis than HMM 1 and HMM 2.

Comparing the most likely state sequences for HMM 1-HMM 4

The most likely state sequences based on the Viterbi algorithm for the four HMMs are displayed in Table 5. The sequences predicted by HMM 1 and HMM 2 are identical, as are those predicted by HMM 3 and HMM 4. Nevertheless, the identical sequences predicted by HMM 3 and HMM 4 are somewhat different from those predicted by HMM 1 and HMM 2. While HMM 1 and HMM 2 successfully predict bank failures due to the 2008 financial crisis, they fail to predict bank failures resulting from the short-lived 2020 Covid-19 recession. In comparison with these models, HMM 3 and HMM 4 predict the 2008 and 2020 large bank failure events nearly successfully, but predict large bank failure events in 2023 and 2024 as well, which did not occur.

Table 5: Comparison of Most Likely State Sequences Predicted by the HMMS

MODEL*	Most likely state sequence
HMM 1 (No covariates)	2 2 2 2 2 2 2 2 2 1 1 1 1 2 2 2 2 2 2 2 2 2 2 2
HMM 2 (RIFSPFFNA)	2 2 2 2 2 2 2 2 2 1 1 1 1 2 2 2 2 2 2 2 2 2 2 2
HMM 3 (UNPCT)	2 2 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 2 1 2 2 2 1 1
HMM 4 (RECESSION)	2 2 2 2 2 2 2 2 2 1 1 1 1 1 1 1 1 2 1 2 2 2 1 1

* List of covariates in parentheses.

6. Comparison of Model Fitting Performance of HMM 1-HMM4 and Model Selection

Table 6 displays the results on the relative model fitting performance of the four HMMs, including an HMM with no covariates and three HMMs with one covariate each, to data on US bank failures from 2001-2024. The table displays the marginal AICs for both states for each model. Among the 4 models studied, HMM 3 with the covariate, UNPCT, has the lowest marginal AIC in state 1 at 171.4028 and 151.2657 in state 2, while HMM 1 with no covariates has the highest marginal AIC in both states, with 319.2049 in state 1 and 318.7751 in state 2. This indicates that of the four HMMs studied, HMM 3 with the covariate, Unemployment Rate (UNPCT), provides the best fit, while HMM 1 with no covariates provides the poorest fit to the data. Model 3 with the covariate, UNPCT, with the lowest AICs in both states, is therefore selected as the best model.

Table 6: Model Fitting Results for HMM 1-HMM 4

MODEL*	Marginal AIC: State 1	Marginal AIC: State 2	-log L
HMM 1	319.2049	318.7751	157.902 (df=5)
HMM 2 (RIFSPFFNA)	277.3536	275.9569	130.9785 (df=7)
HMM 3 (UNPCT)	171.4028	151.2657	68.63286 (df=7)
HMM 4 (RECESSION)	277.3536	275.9569	144.0189 (df=7)

* List of covariates in parentheses.

7. Comparing the Predictive Performance of The Four Poisson HMMs

We created a random training and test set (18,6) split of the 24 observations, trained each HMM on the training set, and evaluated its performance on the test set. The relative prediction performance of the four HMMs was studied by comparing their training and test set root mean squared errors (RMSEs).

Table 7 displays the training and test set RMSEs of forecasts based on the four HMMs. The second column of the table presents the training set RMSEs of forecasts based on these models. HMM 3 with the single covariate UNPCT has the lowest training set RMSE at 0.9511972, while HMM 4 with the single covariate RECESSION has the largest training set RMSE at 1.502328 among all HMMs compared. Since better training set performance does not always guarantee superior out-of-sample performance, we now compare the test set forecasting performance of the four HMMs. Accordingly, the third column of the table presents the test set RMSEs of forecasts based on the four HMMs. HMM 3 with the single covariate UNPCT has the lowest test set RMSE at 1.057775, suggesting that HMM 3 outperforms the other three HMMs on the test set as well. HMM 4 with the single covariate RECESSION has the largest test set RMSE at 2.736184, indicating possible overfitting and poor forecasting performance. This underscores the importance of the unemployment rate in predicting bank failures. It also suggests that while a recession may mean slower economic activity, it may not always be a good predictor of bank failures.

Table 7: Training and Test Sample Forecast RMSEs

MODEL*	TRAINING SAMPLE FORECAST RMSE	TEST SAMPLE FORECAST RMSE
HMM 1 (NO COVARIATES)	0.960083	1.549353
HMM 2 (RIFSPFFNA)	1.493238	1.593751
HMM 3 (UNPCT)	0.9511972	1.057775
HMM 4 (RECESSION)	1.502328	2.736184

* List of covariates in parentheses.

8. Conclusions

The paper studied four two-state HMMs and their relative performance via their application to the US data on bank failures, unemployment rates, federal fund rates, and recessions for the period 2001-2024. These HMMs included a model with no covariates and three models with a single covariate each. A comparison of the model fit performance of these HMMs using the marginal AIC measure suggests that HMM 3 with the single covariate, UNPCT, is the best model among the four HMMs studied. A comparison of the forecasting performance of these HMMs suggests that while the difference between HMM 3 and HMM 1 with no covariates is not significant on the training sample, HMM 3 significantly outperforms HMM 1 as well as the other HMMs on the test sample. The predicted most likely state sequence obtained via the Viterbi algorithm was not consistent across the four HMMs. HMM 3 with the covariate, UNPCT, and HMM 4 with the covariate, RECESSION, yielded identical and the most accurate, most likely predicted state sequence, which was nearly consistent with observed large bank failure events in the United States. Across all HMMs, the transition probability of moving from state 1 (distress) to state 2 (stable) is much higher than the transition the other way around, and persistence in state 2 is higher relative to state 1, underscoring that bank failures have become a rare event in recent years in the United States. Possible extensions of the HMMs employed in this paper are to consider non-homogeneous Markov chains with dependence of transition probabilities on covariates, as well as nonlinearities and dynamics in covariates in both hidden and observed process models. Employing larger time series data may provide further improvements in the model-fitting and forecasting performance with nonlinear data patterns.

The paper demonstrated that HMMs are potentially an effective analytical tool for studying patterns in bank failures. By uncovering latent economic states and enabling predictive insights, HMMs can enhance financial risk assessment and inform regulatory decisions. Future work can build on this study by incorporating multivariate data, testing alternative distributional assumptions, or applying the HMMs to cross-country analyses of bank failures.

References

[1] Dempster, A.P., Laird, N.M. and Rubin, D., Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society: Series B*, 39, (1977),1-38. https://doi.org/10.1111/j.2517-6161.1977.tb01600.x.

[2] Hamilton, J. D., Time Series Analysis. Princeton University Press, (1994). https://doi.org/10.1515/9780691218632.

[3] Hamilton, James, Dates of U.S. recessions as inferred by GDP-based recession indicator [JHDUSRGDPBR], Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/JHDUSRGDPBR, September 21, 2025.

[4] Michelot, Théo, HmmTMB: Hidden Markov Models with Flexible Covariate Effects in R." arXiv Preprint arXiv:2211.14139, (2022).

[5] Michelot, Theo, Analysing time series data with hidden Markov models in hmmTMB, (2024).

[6] Rabiner, L. R., A tutorial on Hidden Markov Models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2), (1989), pp. 257–286. https://doi.org/10.1109/5.18626.

[7] Reiff, Nathan, Historical U.S. Unemployment Rate by Year, Retrieved from https://www.investopedia.com/historical-us-unemployment-rate-by-year-7495494, September 18, 2025.

[8] Sanchez-Espigares, Josep A., and Alberto Lopez-Moreno, R-package MSwM, (2021).

[9] Zucchini, Walter, Iain L MacDonald, and Roland Langrock, Hidden Markov Models for Time Series: An Introduction Using R, Second Edition. CRC, Boca Raton (2017). https://doi.org/10.1201/b20790.

[10] Bank Failures in Brief-Summary, Retrieved from FDIC: https://www.fdic.gov/resources/resolutions/bank-failures/in-brief/index, September 18, 2025.

[11] Board of Governors of the Federal Reserve System (US), Federal Funds Effective Rate [FEDFUNDS], Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/FEDFUNDS, September 18, 2025.

[12] U.S. Bureau of Labor Statistics, Databases, Tables & Calculators by Subject. "Labor Force Statistics from the Current Population Survey: Unemployment Rate," select years 2001 to 2024. Retrieved from Bureau of Labor Statistics Data, September 18, 2025.