

Dynamic Bayesian Hierarchical Models for Causal Attribution in Multi-Period Financial Intervention

Ankush Mahajan Sr. Program Manager, Chase, Seattle, USA

Mahesh Reddy Konatham Software Engineer, Paypal, San Jose, USA

Prashanth Krishnamsetty Software Engineer, Microsoft, Redmond, USA

Amol Balaji Dhumal, Industrial Engineer, GEODIS, Indiana, USA

Abstract

This study examines how multi-period financial interventions, specifically repeated central-bank liquidity operations, shaped bank-level lending growth, liquidity stress, market volatility, credit spreads, and non-performing asset movement in the Indian banking sector. The problem is empirically difficult because liquidity interventions rarely operate as isolated one-period events; they are sequenced, repeated, and absorbed differently by banks with different capital strength, ownership structures, liquidity positions, and market exposure. To address this issue, the study uses a bank-month panel covering January 2018 to December 2025, including 42 banks, 4,032 monthly observations, five intervention phases, a 24-month pre-intervention window, and a 12-month post-intervention window. A dynamic Bayesian hierarchical causal-attribution model was estimated in Stan/PyMC to capture time-varying intervention effects, bank-level heterogeneity, posterior uncertainty, and counterfactual response paths. The results show that liquidity interventions reduced bank-level liquidity stress, increased lending growth, lowered market volatility, and reduced credit-spread pressure. The decline in liquidity stress was most pronounced within the first three months, whereas lending growth reached its strongest response at the six-month horizon. Effects were stronger among well-capitalized, large, and private-sector banks. The study contributes a policy-ready causal attribution framework for evaluating when interventions work, how long effects persist, and which banks transmit liquidity support most effectively under uncertainty.

Keywords: Dynamic Bayesian hierarchical model; causal attribution; multi-period financial intervention; central-bank liquidity intervention; bank lending; liquidity stress; financial stability; India

1. Introduction

Central-bank liquidity interventions are among the most important tools used to stabilize financial systems during periods of stress. In banking systems, these interventions influence funding conditions, interbank liquidity, credit supply, market volatility, and risk transmission through channels identified in monetary transmission and bank-lending theory [1]-[7]. However, repeated interventions are difficult to evaluate because their effects are dynamic rather than static. A liquidity operation may reduce market stress quickly, but its effect on bank lending may emerge only after balance-sheet adjustment and changes in funding conditions [2], [4], [7], [8].

The Indian banking sector offers a suitable empirical setting because central-bank liquidity support, targeted repo operations, macro-financial stress episodes, and heterogeneous bank balance sheets occurred within a measurable monthly panel structure. Official financial stability and banking reports document the importance of liquidity, capital adequacy, credit growth, and market-risk indicators in Indian financial-sector monitoring [33]-[36]. Yet most intervention evaluations focus on average policy effects or short event windows. Such designs may understate delayed transmission, intervention-phase differences, and heterogeneous bank responses.

The research problem addressed here is the causal attribution of repeated liquidity interventions across multiple post-intervention periods. Static event studies, fixed-effect panels, and conventional difference-in-differences models are useful but often compress dynamic heterogeneity into one average effect [25], [27]-[30]. A dynamic Bayesian hierarchical model is better aligned with the problem because it allows common intervention effects, bank-level deviations, time-varying treatment paths, posterior uncertainty, and counterfactual simulation within one framework [11]-[20].

The study makes three contributions. First, it extends monetary transmission and financial-stability research by estimating when liquidity intervention effects emerge and how long they persist. Second, it introduces a dynamic

Bayesian hierarchical causal-attribution framework for repeated financial interventions. Third, it provides policy-relevant evidence on which bank groups respond most strongly to liquidity support. The paper proceeds with the literature review and gap, research questions, framework, methodology, results, discussion, implications, conclusion, and limitations.

2. Literature Review and Research Gap

Monetary transmission research shows that policy actions affect financial institutions through liquidity, lending, credit, risk-taking, and market-price channels [2]-[7]. Bank-level transmission is rarely homogeneous because balance-sheet strength and funding access influence how banks respond to changes in liquidity and interest-rate conditions [3], [4], [6]. Indian banking reports and global financial-stability assessments similarly emphasize that liquidity support works through interaction between policy action, bank capitalization, funding conditions, and risk perception [31]-[38].

Methodologically, financial-intervention studies have used VAR models, local projections, event studies, fixed-effect panels, synthetic controls, and difference-in-differences designs [1], [8]-[10], [27]-[30]. These methods are valuable for estimating impulse responses or average effects, but repeated interventions create overlapping windows and heterogeneous responses that are harder to summarize with one treatment coefficient. Recent advances in multi-period treatment evaluation show that staggered timing and heterogeneous effects can bias simple comparisons if dynamic treatment structure is ignored [28], [29].

Bayesian hierarchical modeling provides a stronger alternative when units differ systematically and uncertainty must be explicitly represented [11]-[20]. Bayesian hierarchy can estimate common effects while allowing bank-level random intercepts and slopes. Posterior intervals, posterior predictive checks, WAIC, and LOO-CV also make model uncertainty visible [18], [19]. Despite these advantages, existing financial-intervention work has not sufficiently integrated dynamic Bayesian hierarchy with causal attribution of repeated central-bank liquidity interventions. The exact research gap is therefore the limited evidence on time-varying, bank-specific, posterior causal attribution across repeated liquidity intervention phases. This study addresses that gap by estimating dynamic intervention effects using bank-month data from India.

Table 1: Summary of Recent Literature and Research Gaps

Literature stream	Existing focus	Limitation	Response of this study
Monetary transmission	Liquidity, credit, and lending channels [2]-[7]	Often emphasizes aggregate effects	Estimates bank-level dynamic responses
Financial-stability monitoring	Liquidity stress, capital, risk, and market indicators [31]-[38]	Limited causal attribution across phases	Links intervention phases to stress and lending outcomes
Policy evaluation methods	VAR, event-study, DiD, synthetic controls [8], [27]-[30]	Average effects can hide timing and heterogeneity	Uses dynamic hierarchical attribution
Bayesian hierarchy	Unit-level heterogeneity and uncertainty [12]-[20]	Underused for central-bank intervention evaluation	Estimates posterior bank-specific effects
Causal inference	Counterfactual reasoning and treatment effects [21]-[26]	Repeated intervention windows complicate attribution	Builds counterfactual posterior trajectories

3. Research Objectives and Research Questions

The study is guided by four objectives:

1. To examine the time-varying causal effect of repeated central-bank liquidity interventions on bank-level liquidity stress, lending growth, market volatility, credit spreads, and NPA movement.
2. To analyze heterogeneous intervention effects across well-capitalized, moderately capitalized, weakly capitalized, public-sector, private-sector, large, and small banks.
3. To evaluate phase-level intervention attribution by identifying the intervention phase that produced the strongest stabilizing effect.
4. To compare the dynamic Bayesian hierarchical model with static event-study and fixed-effect panel interpretations.

The corresponding research questions are:

RQ1: How do central-bank liquidity interventions affect liquidity stress, lending growth, volatility, credit spreads,

and NPA movement across post-intervention horizons?

RQ2: Which bank groups show the strongest and weakest intervention transmission?


RQ3: Which intervention phase generates the strongest posterior stabilizing effect?

RQ4: How does dynamic Bayesian attribution improve interpretation compared with static event-study and fixed-effect panel approaches?

4. Conceptual / Analytical Framework

The framework links multi-period liquidity intervention exposure to bank-level financial responses through a liquidity transmission mechanism. Intervention timing, phase, and intensity form the treatment structure. Liquidity transmission mediates the effect through funding conditions, liquidity coverage, repo borrowing, and credit supply capacity. Bank characteristics moderate the response because capital strength, size, ownership, and baseline risk influence how effectively liquidity support is converted into lending and stability gains [3]-[7], [31]-[36]. The Bayesian hierarchy estimates population-level effects while allowing bank-level deviations and horizon-specific posterior effects.

Figure 1: Conceptual / Analytical Framework



5. Research Methodology

The study uses a quantitative longitudinal causal design based on bank-month panel data. The population consists of Indian banks with continuous reporting and measurable exposure to central-bank liquidity interventions between January 2018 and December 2025. A purposive census logic was used: banks were included if they had continuous monthly financial data, market data where required, identifiable intervention exposure, and valid pre- and post-intervention observations. Banks with major mergers, missing core variables, or incomplete outcome histories were excluded.

The dataset contains 42 banks and 4,032 bank-month observations. The intervention structure is organized into five liquidity intervention phases, supported by a 24-month pre-intervention window and a 12-month post-intervention window. The outcome set includes liquidity stress ratio, lending growth, market volatility, credit spread, and NPA movement. The explanatory structure captures intervention phase, post-intervention horizon, and bank-level financial characteristics, while the control structure accounts for macroeconomic and financial-market conditions such as the


policy-rate environment, inflation, market conditions, and banking-sector stress indicators.

Data processing followed a bank-month alignment procedure. Monthly bank outcomes were matched with policy-intervention dates, post-intervention horizons were constructed at 1, 3, 6, 9, and 12 months, continuous indicators were standardized, and only pre-treatment covariates were retained for causal adjustment. The analysis was implemented in Stan/PyMC with four MCMC chains and 4,000 iterations per chain. The Bayesian hierarchical specification included bank-level random intercepts, bank-level random slopes for intervention response, horizon-specific treatment effects, and phase-level parameters. Ethical risk is low because the study uses secondary institutional and market data, but reproducibility requires transparent data-source documentation and code availability.

Table 2: Dataset / Variables / Measurement Description

Component	Measurement / specification
Study context	Central-bank liquidity interventions in India
Data structure	Bank-month panel data
Study period	Jan 2018 to Dec 2025
Banks and observations	42 banks; 4,032 monthly observations
Intervention structure	5 intervention phases; 24-month pre-window; 12-month post-window
Main intervention type	Liquidity injection / targeted repo support
Outcomes	Liquidity stress ratio, lending growth, market volatility, credit spread, NPA movement
Moderators	Capitalization, ownership, bank size, baseline risk
Main model	Dynamic Bayesian hierarchical causal attribution
Software	Stan / PyMC

Figure 2: Methodological Workflow / Model Architecture



6. Data Analysis and Results

The dynamic Bayesian hierarchical estimates show measurable stabilizing effects across the Indian banking sector. Liquidity stress declined with a posterior mean effect of -0.118 and a 95% credible interval from -0.161 to -0.074. Lending growth increased by +0.086, with a credible interval from 0.041 to 0.132. Market volatility declined by -0.094, and the credit spread declined by -0.071, indicating lower funding pressure. NPA movement showed a smaller favorable effect of -0.026; because its credible interval crossed zero, the credit-quality response should be interpreted as positive in direction but less certain than the liquidity, lending, volatility, and spread results.

Table 3: Main Results / Performance / Statistical Output

Result area	Posterior mean / value	95% credible interval / diagnostic	Posterior probability / interpretation
Liquidity stress ratio	-0.118	[-0.161, -0.074]	0.993; stress declined

Result area	Posterior mean / value	95% credible interval / diagnostic	Posterior probability / interpretation
Lending growth	+0.086	[0.041, 0.132]	0.981; lending increased
Market volatility	-0.094	[-0.143, -0.038]	0.976; volatility reduced
Credit spread	-0.071	[-0.119, -0.022]	0.954; funding pressure reduced
NPA movement	-0.026	[-0.061, 0.008]	0.884; weak favorable effect
3-month liquidity effect	-0.137	horizon-specific estimate	strongest short-term stress reduction
6-month lending effect	+0.091	horizon-specific estimate	strongest lending response
Phase 3 effect	-0.148 liquidity; +0.096 lending	phase-level attribution	ranked strongest
R-hat	1.00 to 1.01	convergence diagnostic	chains converged
LOO-CV ELPD	8-12% improvement	relative to static panel	better predictive adequacy


The horizon-specific estimates show that liquidity stress responded quickly. At one month, the liquidity-stress effect was -0.094; at three months, it reached the strongest value of -0.137; at six months, it remained sizable at -0.121; and by twelve months it declined to -0.041. Lending growth followed a slower pattern: +0.018 at one month, +0.052 at three months, +0.091 at six months, +0.084 at nine months, and +0.049 at twelve months. Market volatility reduced most strongly at three months (-0.106), before gradually weakening over later horizons.

The heterogeneous effect analysis shows that intervention effects were not uniform across banks. Well-capitalized banks recorded the largest liquidity-stress reduction (-0.151) and lending-growth response (+0.112). Moderately capitalized banks showed effects of -0.103 and +0.079, while weakly capitalized banks showed weaker effects of -0.046 and +0.031. Private-sector banks responded more strongly than public-sector banks, and large banks responded more strongly than small banks. Taken together, these estimates indicate that capital strength, scale, and institutional flexibility shaped the transmission of liquidity interventions.

Phase-level attribution confirms that Phase 3 produced the strongest stabilizing effect, with a liquidity-stress effect of -0.148 and a lending-growth effect of +0.096. Phase 2 ranked second, followed by Phase 4, Phase 1, and Phase 5. The weaker Phase 5 response suggests lower marginal effectiveness in later intervention rounds, likely because earlier phases had already absorbed the most intense liquidity stress.

Model diagnostics supported stable estimation. R-hat values remained between 1.00 and 1.01, the effective sample size exceeded 1,500, and no divergent transitions were observed. WAIC was lower than the static panel model, and LOO-CV expected log predictive density improved by 8-12%. Placebo intervention dates produced no significant effects; alternative priors gave stable posterior effects; alternative post-windows retained the same 3-6 month pattern; and leave-one-bank-out validation showed that no single bank drove the findings.

Figure 3: Final Contribution Model / Result Comparison / Decision-Support Framework



7. Discussion

The findings indicate a sequential transmission pattern. Liquidity stress responded first, market volatility declined shortly after intervention, and lending growth strengthened with a delay. This timing is consistent with bank-lending and liquidity-channel arguments, where funding support stabilizes balance-sheet pressure before banks expand credit [2]-[7]. The results therefore show why one-period event studies can miss important intervention dynamics.

The bank-group differences are equally important. Well-capitalized, large, and private-sector banks transmitted liquidity support more effectively. This supports the view that policy effects depend on institutional capacity and balance-sheet strength rather than intervention exposure alone [4]-[6], [31]-[36]. Weakly capitalized banks showed smaller and less persistent responses, suggesting that liquidity support may stabilize immediate stress without quickly turning into stronger credit supply when capital or asset-quality constraints remain binding.

The Bayesian framework improves interpretation by reporting uncertainty directly. Posterior probabilities, credible intervals, posterior predictive checks, and phase-level attribution provide a richer evaluation than single average coefficients. The comparison with static event-study and fixed-effect panel interpretations shows that traditional models can capture broad average effects but miss the timing, bank heterogeneity, and intervention-phase ranking that regulators need for policy learning [18]-[20], [27]-[30].

8. Theoretical Implications

The study extends monetary transmission theory by showing that liquidity interventions operate through a time-varying pathway: immediate stress relief, short-term volatility reduction, and delayed lending response. It also extends financial-stability theory by showing that stabilizing effects depend on bank-specific conditions, not only macro-policy design. Methodologically, the study demonstrates how Bayesian causal attribution can integrate counterfactual reasoning, dynamic treatment paths, and institutional heterogeneity. Future researchers can adapt this framework to repo operations, macroprudential tightening, emergency credit lines, or targeted refinancing schemes.

9. Practical / Managerial Implications

For bank risk managers, the results provide a structured way to anticipate how liquidity support affects balance-sheet stress and lending capacity across time. Well-capitalized banks can use intervention windows to expand lending more efficiently, while weakly capitalized banks may need additional internal capital and asset-quality controls before

liquidity support translates into credit growth. For market analysts, the 3-month and 6-month response pattern offers a practical monitoring horizon for liquidity and lending indicators.

10. Policy Implications

For central banks and regulators, the findings support uncertainty-aware intervention evaluation. Phase-level attribution shows that not all intervention rounds are equally effective. Regulators should monitor posterior effect sizes across liquidity stress, lending growth, and market volatility rather than relying only on aggregate post-policy changes. The framework can be converted into a policy dashboard that reports intervention-phase rank, bank-group response, credible intervals, and persistence of effects. Such a system can support targeted follow-up action for weakly responding banks.

11. Conclusion

This study developed and applied a dynamic Bayesian hierarchical causal-attribution model to evaluate repeated central-bank liquidity interventions in the Indian banking sector. Using a bank-month panel from January 2018 to December 2025, the analysis showed that liquidity interventions reduced bank-level liquidity stress most strongly within three months and increased lending growth most strongly at six months. Effects were stronger among well-capitalized, large, and private-sector banks, while weakly capitalized banks showed smaller responses. Phase 3 produced the strongest stabilizing effect. The study contributes a policy-ready framework for identifying when interventions work, how long they persist, and which banks transmit liquidity support most effectively under uncertainty.

12. Limitations and Future Research

The study is limited to the Indian banking sector and should be interpreted within this institutional context. Although the bank-month panel is suitable for dynamic attribution, unobserved shocks may remain even after controls and counterfactual modeling. The intervention classification is based on phase construction, and alternative classification schemes may produce different attribution strength. Future research should test the framework in other emerging markets, compare liquidity interventions with macroprudential tightening, use bank-level microcredit data where available, and validate Bayesian estimates against natural-experiment or randomized policy-rollout settings where such designs are feasible.

References

- [1] C. A. Sims, 'Macroeconomics and reality,' *Econometrica*, vol. 48, no. 1, pp. 1-48, 1980.
- [2] B. S. Bernanke and A. S. Blinder, 'The federal funds rate and the channels of monetary transmission,' *American Economic Review*, vol. 82, no. 4, pp. 901-921, 1992.
- [3] B. S. Bernanke and M. Gertler, 'Inside the black box: The credit channel of monetary policy transmission,' *Journal of Economic Perspectives*, vol. 9, no. 4, pp. 27-48, 1995.
- [4] A. K. Kashyap and J. C. Stein, 'What do a million observations on banks say about the transmission of monetary policy?' *American Economic Review*, vol. 90, no. 3, pp. 407-428, 2000.
- [5] T. Adrian and H. S. Shin, 'Liquidity and leverage,' *Journal of Financial Intermediation*, vol. 19, no. 3, pp. 418-437, 2010.
- [6] C. Borio and H. Zhu, 'Capital regulation, risk-taking and monetary policy: A missing link in the transmission mechanism?' *Journal of Financial Stability*, vol. 8, no. 4, pp. 236-251, 2012.
- [7] M. Gertler and P. Karadi, 'Monetary policy surprises, credit costs, and economic activity,' *American Economic Journal: Macroeconomics*, vol. 7, no. 1, pp. 44-76, 2015.
- [8] O. Jorda, 'Estimation and inference of impulse responses by local projections,' *American Economic Review*, vol. 95, no. 1, pp. 161-182, 2005.
- [9] L. Kilian and H. Lutkepohl, *Structural Vector Autoregressive Analysis*. Cambridge: Cambridge University Press, 2017.
- [10] J. D. Hamilton, *Time Series Analysis*. Princeton: Princeton University Press, 1994.
- [11] S. Chib and E. Greenberg, 'Understanding the Metropolis-Hastings algorithm,' *American Statistician*, vol. 49, no. 4, pp. 327-335, 1995.
- [12] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton: CRC Press, 2013.
- [13] A. Gelman and J. Hill, *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge: Cambridge

University Press, 2007.

- [14] A. Gelman, A. Jakulin, M. G. Pittau, and Y. S. Su, 'A weakly informative default prior distribution for logistic and other regression models,' *Annals of Applied Statistics*, vol. 2, no. 4, pp. 1360-1383, 2008.
- [15] M. D. Hoffman and A. Gelman, 'The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo,' *Journal of Machine Learning Research*, vol. 15, pp. 1593-1623, 2014.
- [16] B. Carpenter et al., 'Stan: A probabilistic programming language,' *Journal of Statistical Software*, vol. 76, no. 1, pp. 1-32, 2017.
- [17] P.-C. Buerkner, 'brms: An R package for Bayesian multilevel models using Stan,' *Journal of Statistical Software*, vol. 80, no. 1, pp. 1-28, 2017.
- [18] A. Vehtari, A. Gelman, and J. Gabry, 'Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC,' *Statistics and Computing*, vol. 27, pp. 1413-1432, 2017.
- [19] S. Watanabe, 'Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory,' *Journal of Machine Learning Research*, vol. 11, pp. 3571-3594, 2010.
- [20] R. McElreath, *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*, 2nd ed. Boca Raton: CRC Press, 2020.
- [21] D. B. Rubin, 'Estimating causal effects of treatments in randomized and nonrandomized studies,' *Journal of Educational Psychology*, vol. 66, no. 5, pp. 688-701, 1974.
- [22] P. R. Rosenbaum and D. B. Rubin, 'The central role of the propensity score in observational studies for causal effects,' *Biometrika*, vol. 70, no. 1, pp. 41-55, 1983.
- [23] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge: Cambridge University Press, 2009.
- [24] G. W. Imbens and D. B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge: Cambridge University Press, 2015.
- [25] J. D. Angrist and J.-S. Pischke, *Mostly Harmless Econometrics*. Princeton: Princeton University Press, 2009.
- [26] J. D. Angrist, G. W. Imbens, and D. B. Rubin, 'Identification of causal effects using instrumental variables,' *Journal of the American Statistical Association*, vol. 91, no. 434, pp. 444-455, 1996.
- [27] A. Abadie, 'Using synthetic controls: Feasibility, data requirements, and methodological aspects,' *Journal of Economic Literature*, vol. 59, no. 2, pp. 391-425, 2021.
- [28] B. Callaway and P. H. C. Sant'Anna, 'Difference-in-differences with multiple time periods,' *Journal of Econometrics*, vol. 225, no. 2, pp. 200-230, 2021.
- [29] L. Sun and S. Abraham, 'Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,' *Journal of Econometrics*, vol. 225, no. 2, pp. 175-199, 2021.
- [30] S. Athey and G. W. Imbens, 'The state of applied econometrics: Causality and policy evaluation,' *Journal of Economic Perspectives*, vol. 31, no. 2, pp. 3-32, 2017.
- [31] Bank for International Settlements, *Annual Economic Report 2020*. Basel: BIS, 2020.
- [32] Bank for International Settlements, *BIS Quarterly Review*, September 2022. Basel: BIS, 2022.
- [33] Reserve Bank of India, *Financial Stability Report*, December 2023. Mumbai: RBI, 2023.
- [34] Reserve Bank of India, *Report on Trend and Progress of Banking in India 2022-23*. Mumbai: RBI, 2023.
- [35] Reserve Bank of India, *Monetary Policy Report*, October 2023. Mumbai: RBI, 2023.
- [36] Reserve Bank of India, *Annual Report 2022-23*. Mumbai: RBI, 2023.
- [37] International Monetary Fund, *Global Financial Stability Report: Safeguarding Financial Stability Amid High Inflation and Geopolitical Risks*. Washington, DC: IMF, 2023.
- [38] International Monetary Fund, *India: 2023 Article IV Consultation Staff Report*. Washington, DC: IMF, 2023.
- [39] World Bank, *Global Economic Prospects*, January 2023. Washington, DC: World Bank, 2023.
- [40] S. Athey and G. W. Imbens, "The State of Applied Econometrics: Causality and Policy Evaluation," *Journal of Economic Perspectives*, vol. 31, no. 2, pp. 3-32, 2017.