

Article

Bayesian Causal Inference for Credit Default Risk

Sello Dalton Pitso *  and Taryn Michael 

Natural & Applied Sciences, Sol Plaatje University, Kimberley 8300, South Africa; taryn.michael@spu.ac.za

* Correspondence: pitsojaden81@gmail.com

Abstract

Banks often assume that higher credit limits increase customer default risk because greater exposure appears to imply greater vulnerability. This reasoning, however, conflates correlation with causation. Whether increasing a customer's credit limit truly raises the likelihood of default remains an open empirical question that this work seeks to answer. We applied Bayesian causal inference to estimate the causal effect of credit limits on default probability. The analysis incorporated Directed Acyclic Graphs (DAGs) for causal structure, d-separation for identification, and Bayesian logistic regression using a dataset of 30,000 credit card holders in Taiwan (April–September 2005). Twenty-two confounding variables were adjusted for, covering demographics, repayment history, and billing and payment behavior. Continuous covariates were standardized, and posterior inference was performed using NUTS sampling with posterior predictive simulations to compute Average Treatment Effects (ATEs). We found that a one-standard-deviation increase in credit limit reduces default probability by 1.44 percentage points (94% HDI: $[-2.0\%, -1.0\%]$), corresponding to a 6.3% relative decline from the baseline default rate of 22.1%. The effect was consistent across demographic subgroups, with homogeneous treatment effects observed for age, education, and gender categories, and remained robust under sensitivity analysis addressing potential unmeasured confounding. The findings suggest that increasing credit limits can causally reduce default risk, likely by enhancing financial flexibility and lowering utilization ratios. These results have practical implications for credit policy design and motivate further investigation into mechanisms and applicability across broader lending environments. These estimates are explicitly interpreted as context-specific causal effects for a pre-crisis consumer credit environment, with external validity assessed conceptually rather than assumed.

Keywords: Bayesian inference; causal inference; credit risk; default prediction; directed acyclic graphs; observational studies; sensitivity analysis



Received: 24 November 2025

Revised: 15 January 2026

Accepted: 26 January 2026

Published: 12 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Credit risk remains a central concern for financial institutions, with borrower default affecting portfolio performance, liquidity management, and regulatory capital requirements. In industry practice, credit limits are commonly treated as a primary driver of exposure, with higher limits assumed to mechanically increase default risk. This assumption underlies conservative limit-setting policies in which increases in available credit are viewed predominantly as sources of additional risk.

However, from a credit-risk-theory perspective, this view is incomplete. Several established mechanisms suggest that higher credit limits may, under certain conditions, reduce the risk of default rather than increase it. First, theories of liquidity smoothing

and buffer stock behavior posit that access to unused credit allows borrowers to absorb transitory income shocks without immediately defaulting (Carroll 1997; Gross and Souleles 2002). Second, utilization-based risk frameworks emphasize that default risk is driven more by the ratio of outstanding balances to available credit than by credit limits alone. Higher limits can therefore lower utilization rates and, in turn, reduce financial distress (Khandani et al. 2010). Finally, behavioral responses to credit constraints suggest that binding limits may amplify repayment stress, increase delinquency clustering, and accelerate default during adverse shocks (Agarwal et al. 2015).

Despite their relevance for credit risk measurement and policy design, these mechanisms are rarely examined through a causal lens. Much of the empirical literature relies on predictive or associative models that cannot disentangle the effect of credit limits from borrower selection or endogenous limit assignment. This study addresses this gap by estimating the causal effect of credit limit increases on default probability using a Bayesian causal inference framework that explicitly models counterfactual outcomes and uncertainty. The contribution of this study is methodological and inferential rather than temporal. It evaluates how credit limit interventions operate under explicitly stated causal assumptions, not how default risk should be predicted in modern datasets.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature and situates this study within the broader credit risk and causal inference landscape. Section 3 describes the data and methodology. Section 4 presents the empirical results, including average treatment effects, subgroup analyses, model validation, and sensitivity checks. Section 5 discusses implications for credit risk management and policy design. Section 6 concludes the paper.

2. Literature Review

This review is intentionally selective, focusing on studies relevant to causal identification, policy interventions, and credit limit mechanisms rather than exhaustively surveying predictive credit scoring models.

Most existing research in credit risk prioritizes prediction over causal explanation. Wang et al. (2021) demonstrated that machine learning models excel at identifying borrowers likely to default, while Kvamme et al. (2018) extended these capabilities to convolutional neural networks for mortgage default prediction. These predictive frameworks, however, cannot answer counterfactual questions such as whether increasing a customer's credit limit would alter their default probability. The distinction matters because prediction identifies patterns in existing data, whereas causal inference evaluates what would happen under alternative policies.

The development of formal causal inference frameworks has provided tools to address these limitations. Pearl's work on Directed Acyclic Graphs (DAGs) (Pearl 2009) established a structural approach to distinguishing causal pathways from spurious associations, offering a graphical language for encoding assumptions about confounding and mediation. While Pearl's framework emerged from computer science and philosophy, Hernán and Robins (2020) extended these principles to observational epidemiology and public health, where randomized experiments are often infeasible. Their target trial framework has since been adopted in economics and social sciences to emulate experiments using observational data.

Despite these methodological advances, applications to credit limit policy remain sparse. Fuster et al. (2022) examined algorithmic fairness in lending decisions, focusing on how machine learning models may perpetuate or mitigate bias across demographic groups. Dobbie et al. (2020) analyzed the long-term consequences of personal bankruptcy using quasi-experimental variation in judge leniency. Both studies engage with causal questions,

yet neither directly investigates whether adjusting credit limits influences default behavior. The gap is striking given that credit limit policies represent one of the most common and consequential interventions available to lenders.

The gap addressed by this study is therefore not the absence of predictive accuracy, nor the lack of sophisticated modeling in credit risk research. Rather, the gap lies in the absence of formally identified causal estimates for credit limit interventions themselves. Existing studies either (i) predict default without counterfactual interpretation or (ii) exploit quasi-experiments that affect credit access indirectly rather than credit limits directly. As a result, the causal effect of adjusting credit limits, a routine, policy-relevant decision faced by lenders, remains largely unquantified. This study directly targets that missing estimand. Within this bounded scope, the contributions are fourfold: (1) implementing DAG-based identification for credit limit interventions, (2) applying Bayesian logistic regression to obtain full posterior uncertainty quantification, (3) evaluating heterogeneity of effects across demographic subgroups, and (4) testing robustness to unmeasured confounding using formal sensitivity analysis. Together, these contributions provide a more rigorous basis for understanding how credit limit adjustments influence borrower behavior and offer practical insights for policy formation within financial institutions.

3. Materials and Methods

3.1. Data

The dataset comes from a financial institution in Taiwan and covers credit card customers from April to September 2005. The dataset contains 30,000 customers with complete data across 24 variables. The data is publicly available and has been used before for credit scoring, though not for causal inference.

The outcome is binary default status (`default.payment.next.month`) which is either 1 for default or 0 for no default. Overall, 22.1% of customers (6630 people) defaulted. The treatment variable is credit limit (`LIMIT_BAL`), measured in New Taiwan dollars, ranging from NT\$10,000 to NT\$1,000,000.

The confounding variables include:

- **Demographics:** AGE (21–79 years), SEX (1 = male, 2 = female), EDUCATION (1 = graduate school, 2 = university, 3 = high school, 4 = others), MARRIAGE (1 = married, 2 = single, 3 = others)
- **Payment History:** PAY_0 through PAY_6 showing repayment status for six months (September 2005 back to April 2005)
- **Billing Information:** BILL_AMT1 through BILL_AMT6 for bill statement amounts across six months
- **Payment Amounts:** PAY_AMT1 through PAY_AMT6 showing actual payment amounts for six months
- **Income:** Unmeasured confounder affecting both credit limit assignment and default risk. While not directly observed in the dataset, the observed financial variables (payment history, billing amounts, payment amounts) and demographics serve as partial proxies for income

The repayment status variables PAY_0 through PAY_6 indicate the number of months of payment delay for the most recent six billing cycles (September 2005 to April 2005). These variables are coded as follows: −1 denotes payment made duly (no delay), 0 denotes revolving credit with no payment delay, and positive integers from 1 to 9 indicate payment delays of one to nine months, respectively, with higher values representing more severe delinquency. These variables therefore capture recent repayment behavior and credit discipline.

Because repayment history directly influences both credit limit assignment and default risk, the variables PAY_0–PAY_6 were included as confounders in the causal adjustment set.

We standardized all continuous variables (z-scoring) before modeling. This makes effect sizes comparable and helps with computational efficiency. All reported coefficients, including the treatment effect, are therefore expressed in terms of standardized variables (per one-standard-deviation change).

3.2. Causal Assumptions and Directed Acyclic Graph

Figure 1 summarizes the assumed causal structure underlying the analysis. Credit limit assignment is influenced by demographic characteristics, repayment history, and observed financial behavior, all of which also affect default risk. Income is treated as an unobserved latent common cause of both credit limits and default. While income is not directly observed and may introduce residual confounding, the assumed adjustment set is designed to account for the primary observed confounding pathways between credit limit and default, allowing the estimated effects to be interpreted causally under the stated assumptions and subject to sensitivity analyses assessing robustness to unmeasured confounding. Observed financial variables and demographic characteristics serve as partial proxies for income and related latent traits, mitigating but not eliminating this source of bias. The Stable Unit Treatment Value Assumption (SUTVA) is assumed to hold, though spillover effects in credit markets may violate this assumption in practice.




Figure 1. Directed acyclic graph representing the assumed causal structure for identifying the effect of credit limit on default. LIMIT_BAL is the treatment variable, default is the outcome, observed variables (demographics, payment history, billing information, payment amounts) are measured confounders, and income (shown in gray/dashed) is an unmeasured confounder.

We assume positivity: all covariate combinations have non-zero probability of observing each credit limit level. Empirical diagnostics in Figure A1 (Appendix A) confirm this assumption holds. All demographic strata (age, education, gender) exhibit credit limits

spanning nearly the full observed range (NT\$10,000 to NT\$1,000,000), with substantial within-group variation. A propensity model predicting credit limit from all 22 confounders achieves $R^2 = 0.36$, indicating sufficient unexplained variation in treatment assignment to support causal identification.

Payment amounts (PAY_AMT1 through PAY_AMT6) are treated as pre-treatment confounders rather than mediators. These variables reflect payment behavior in the six months leading up to the observation period and therefore precede both the credit limit assignment captured in our data and the subsequent default outcome. Credit limits in this dataset represent existing limits at baseline, not newly assigned limits, so historical payment behavior confounds the relationship rather than mediates it.

3.3. Causal Identification Strategy

Causal identification is based on an explicitly stated directed acyclic graph (DAG) that formalizes assumptions about the data-generating process. Credit limit assignment is assumed to depend on observed demographic characteristics, repayment history, and billing behavior, all of which also influence default probability. Household income, while unobserved, is assumed to affect both credit limits and default, introducing a potential source of residual confounding.

Following Pearl (2009), identification relies on the backdoor criterion. An adjustment set satisfies this criterion if it excludes variables affected by the treatment and accounts for confounding pathways linking the treatment to the outcome through observed variables. While income is not directly observed, the included financial and demographic covariates are assumed to capture the primary observable pathways through which income influences both credit limit assignment and default risk.

D-separation analysis was used to assess whether conditioning on the 22 observed covariates is sufficient to control for confounding induced by measured variables under the assumed DAG. This analysis confirms that the proposed adjustment set addresses all backdoor paths arising from observed covariates. Identification therefore rests on the assumption that residual confounding due to unmeasured income is limited, given its strong empirical relationship with observed financial behavior. The robustness of the estimated causal effects to violations of this assumption is evaluated through sensitivity analyses presented in Section 3.

3.4. Bayesian Model Specification

We used Bayesian logistic regression since the outcome is binary and the coefficients are interpretable on the log-odds scale.

Unless otherwise stated, all uncertainty intervals reported in this study are 94% highest density intervals (HDIs), corresponding to the 3rd and 97th posterior percentiles as produced by the Bayesian model outputs.

The model is specified as:

$$\text{DEFAULT}_i \sim \text{Bernoulli}(p_i), \quad (1)$$

$$\text{logit}(p_i) = \alpha + \beta_T \cdot \text{LIMIT_BAL}_i + \sum_{j=1}^{22} \beta_{C_j} \cdot \text{Confounder}_{ij}. \quad (2)$$

Here, $\text{logit}(p_i) = \log\left(\frac{p_i}{1-p_i}\right)$ denotes the log-odds of default for customer i , β_T represents the causal effect of credit limit on default, and β_j , for $j = 1, \dots, 22$, are coefficients corresponding to the confounding variables.

We assigned weakly informative normal priors:

$$\alpha \sim \mathcal{N}(0, 2), \quad (3)$$

$$\beta_T \sim \mathcal{N}(0, 1), \quad (4)$$

$$\beta_{C_j} \sim \mathcal{N}(0, 1), \quad j = 1, \dots, 22. \quad (5)$$

These priors regularize the model but are diffused enough to let the data dominate. The intercept is given a larger variance to accommodate the logit scale, while coefficient priors centered at zero represent weak prior beliefs about effect directions.

3.5. Linearity Assessment

To assess the adequacy of a linear specification for credit limit effects, we examine posterior predictive residuals across the distribution of credit limits. Residuals were computed and binned into deciles of the standardized credit limit variable to detect systematic deviations from linearity.

Under a well-specified linear model, residuals should exhibit random scatter around zero with no pronounced curvature or heteroskedastic patterns. In addition to residual diagnostics, we explored more flexible functional forms, including quadratic terms and restricted cubic splines, to assess the presence of nonlinear effects.

As shown in Appendix A Figure A2, the residual diagnostics support the linear specification. Binned residuals fluctuate around zero across credit limit deciles, with no clear monotonic trend or systematic curvature. More flexible model specifications exhibited substantial computational instability, including a large number of divergent transitions, indicating poorly identified posterior geometry and limiting the reliability of inference under these specifications.

Given the absence of systematic residual patterns and the lack of stable convergence for nonlinear models, we retain the linear specification as a parsimonious and empirically supported approximation. The estimated treatment effect therefore represents an average marginal effect on the log-odds scale, summarizing the association between credit limits and default risk across the observed range of credit limits.

3.6. Posterior Sampling and Convergence Diagnostics

We used PyMC version 5.10 with the No-U-Turn Sampler (NUTS) (Hoffman and Gelman 2014) for posterior sampling. NUTS is a variant of Hamiltonian Monte Carlo that auto-tunes step sizes and trajectory lengths. We ran 4 independent chains with 2000 samples each after discarding 1000 warmup samples. The target acceptance rate was 0.95, ensuring that the sampler explored the posterior thoroughly.

Convergence diagnostics included examining the $\hat{R}$ statistic (the Gelman–Rubin diagnostic (Gelman et al. 2013)) to compare the variance within and between chains, as well as calculating the effective sample size (ESS) to assess the amount of independent information contained in the correlated posterior samples. Trace plots were inspected for all parameters to ensure proper mixing across chains and verified that the number of divergent transitions was zero, as any divergences would indicate potential sampling pathologies.

All parameters hit $\hat{R} \leq 1.01$ and effective sample sizes over 2800. Zero divergent transitions across all chains, indicating that the sampling was successful.

3.7. Causal Effect Estimation

The Average Treatment Effect (ATE) was estimated using posterior predictive simulation under counterfactual interventions. For each of the 8000 posterior draws, baseline default probabilities $p_i^{(0)}$ were first predicted for all individuals under their observed

credit limits. A counterfactual intervention was then simulated by uniformly increasing credit limits by one-standard-deviation (on the standardized scale), and the corresponding counterfactual default probabilities $p_i^{(1)}$ were generated. Individual treatment effects were computed as $p_i^{(1)} - p_i^{(0)}$, and averaged across all individuals to obtain the ATE for each posterior draw.

This procedure yields a full posterior distribution over the ATE, enabling coherent Bayesian uncertainty quantification for the estimated causal effect.

Conditional Average Treatment Effects (CATEs) were estimated using the same procedure within demographic subgroups defined by age (≤ 30 , 31–50, > 50), education (graduate school, university, high school), and gender (male, female). This subgroup analysis allows assessment of treatment effect heterogeneity across populations that may differ in baseline risk exposure and credit access.

3.8. Sensitivity Analysis

Causal inference from observational data relies on the untestable assumption of no unmeasured confounding. To assess the robustness of our findings to violations of this assumption, we conducted a sensitivity analysis following the framework of VanderWeele and Ding (2017), which evaluates how strong an unmeasured confounder would need to be to fully explain away the estimated causal effect.

We simulated hypothetical unmeasured confounders U with varying strengths by allowing the correlation between the confounder and the treatment, $\rho(U, \text{LIMIT_BAL})$, to range from 0 to 0.5. The effect of the confounder on the outcome was parameterized by β_U , which was allowed to vary between 0 and 2.0 on the log-odds scale. This grid of values captures a wide range of plausible confounding scenarios and enables systematic assessment of potential bias in the estimated effect.

For each combination of confounder strength and outcome influence, a bias-adjusted Average Treatment Effect was computed, and the threshold at which the estimated effect was attenuated to zero or reversed in sign was identified. In addition, we calculated E-values (VanderWeele and Ding 2017), which quantify the minimum strength of association that an unmeasured confounder would need with both the treatment and the outcome, on the risk ratio scale, to fully explain the observed effect.

3.9. Model Validation

Posterior predictive checks were conducted to assess whether the fitted model adequately reproduces key features of the observed data. Specifically, 4000 replicated datasets were generated from the posterior predictive distribution and compared with the observed outcomes.

Model validation focused on several complementary dimensions. First, overall fit was evaluated by comparing the posterior predictive default rate with the observed default rate of 22.1%. Second, calibration was assessed by comparing predicted default probabilities with observed default frequencies across probability bins. Third, discrimination was examined by evaluating the separation of predicted probabilities between defaulters and non-defaulters. Finally, residual plots were inspected to identify systematic patterns or evidence of model misspecification.

Together, these diagnostics indicate that the model provides an adequate representation of the observed data-generating process for the purposes of prediction and causal effect estimation.

4. Results

4.1. Model Convergence and Parameter Estimates

Table 1 shows posterior summary statistics for key parameters. All parameters achieved $\hat{R} = 1.0$ with effective sample sizes over 2800. Zero divergent transitions across 8000 samples.

Table 1. Posterior summary statistics for key model parameters.

Parameter	Mean	SD	HDI 3%	HDI 97%	ESS	$\hat{R}$
Intercept (α)	−1.468	0.016	−1.499	−1.439	6356	1.0
Treatment ($\beta_{\text{LIMIT_BAL}}$)	−0.099	0.020	−0.137	−0.061	6510	1.0

The intercept posterior mean of −1.468 represents baseline log-odds of default. The treatment effect coefficient of −0.099 (94% HDI: [−0.137, −0.061]) shows that higher credit limits is associated with lower default probability after adjusting for all 22 confounders. The entire credible interval excludes zero, indicating strong evidence against no effect.

4.2. Average Treatment Effect

The ATE is −0.0144 (94% HDI: [−0.020, −0.010]), meaning a one-standard-deviation increase in credit limit (roughly NT\$100,000) is estimated to cause a 1.44 percentage point drop in default probability. This represents a 6.3% relative reduction in defaults, dropping the baseline rate from 22.1% to a counterfactual rate of 20.7% under intervention. Figure 2 corroborates this result: the left panel shows the posterior distribution of the ATE, which lies entirely below zero, providing strong evidence of a consistently negative (protective) effect of higher credit limits on default risk. The right panel further illustrates the dose–response relationship, demonstrating that progressively larger increases in credit limits are associated with proportionally greater reductions in default probability, supporting a monotonic and economically meaningful treatment effect rather than a threshold-driven response.




Figure 2. Average Treatment Effect analysis. **Left:** Posterior distribution of the ATE (mean −0.0144; 94% HDI [−0.020, −0.010]), indicating a consistently negative and protective effect. **Right:** Dose–response curve illustrating that larger credit limit increases lead to proportionally greater reductions in default risk.

4.3. Conditional Average Treatment Effects

Table 2 reports conditional average treatment effects (CATEs) across demographic subgroups defined by age, education level, and gender. The estimated effects are remarkably consistent in both magnitude and uncertainty across all subgroups, with posterior means clustered around −0.014 and overlapping 94% highest density intervals. This indicates that a one–standard-deviation increase in credit limit produces a similar reduction in

default probability regardless of demographic characteristics, providing little evidence of meaningful treatment effect heterogeneity.

This homogeneity is further supported by Figure 3. The posterior distribution shown in the left panel remains tightly concentrated below zero, reinforcing the stability of the estimated treatment effect across conditioning variables. The right panel illustrates a smooth and monotonic dose–response relationship between standardized credit limit increases and default probability, suggesting that the causal effect operates uniformly across the population rather than being driven by specific subgroups.

Table 2. Conditional Average Treatment Effects by demographic subgroup.

Subgroup	N	CATE Mean	94% HDI
Age ≤ 30	11,013	−0.0150	[−0.020, −0.010]
Age 31–50	16,718	−0.0142	[−0.019, −0.009]
Age > 50	2269	−0.0148	[−0.021, −0.009]
Graduate School	10,585	−0.0131	[−0.018, −0.008]
University	14,030	−0.0144	[−0.019, −0.009]
High School	4917	−0.0143	[−0.020, −0.009]
Male	11,888	−0.0141	[−0.019, −0.009]
Female	18,112	−0.0145	[−0.020, −0.009]




Figure 3. Average Treatment Effect (ATE) of credit limit increases on default probability. **Left:** Posterior distribution of the ATE (probability units) with mean -0.0144 (1.44 percentage point reduction) and 94% HDI $[-0.020, -0.010]$, corresponding to a one-standard-deviation increase in credit limit (approximately NT dollars 100,000). **Right:** Dose-response curve showing the causal change in default probability as a function of standardized credit limit increases.

4.4. Model Validation

Posterior predictive checks indicate that the model adequately reproduces key features of the observed data, including the marginal default rate, calibration across risk strata, and discrimination between defaulters and non-defaulters. While mild residual nonlinearities are visible at the extremes of the credit limit distribution, these deviations are limited in magnitude and concentrated in regions with sparse data. As such, they do not materially affect the estimation of the average treatment effect, supporting the use of the fitted model for causal effect estimation under the stated assumptions.

4.5. Sensitivity Analysis

While sample sizes vary across demographic subgroups (Age > 50: $n = 2269$ vs. Age 31–50: $n = 16,718$), all groups retain sufficient statistical power. The 94% credible intervals for all subgroups exclude zero and exhibit comparable widths (Table 2), indicating similar estimation precision across demographics. The smallest subgroup (Age > 50, interval width = 0.012) shows only a marginally wider interval than the largest subgroup (Age

31–50, interval width = 0.010), suggesting that differences in sample size do not materially compromise the robustness of the conditional effect estimates.

Sensitivity analysis indicates that completely nullifying the observed ATE of -1.44 percentage points would require an unmeasured confounder that is strongly associated with both credit limit assignment and default risk. Specifically, such a confounder would need to exhibit a correlation with the treatment of at least $\rho \geq 0.30$ and an effect size on default probability of $\beta_U \geq 1.0$ on the log-odds scale—substantially larger than the effects of any observed covariates. The corresponding E-value of approximately 1.15 further suggests that only a confounder of considerable strength could plausibly eliminate the estimated causal effect.

Model adequacy is further assessed using posterior predictive checks shown in Figure 4. The top-left panel compares the posterior predictive distribution of the default rate with the observed default rate. The observed rate (red dashed line) and the posterior predictive mean (blue dashed line) coincide at 22.1%, resulting in visual overlap; although both reference lines are present, the red line is not separately visible because it lies exactly beneath the blue line. This overlap indicates excellent calibration of the model at the aggregate level.




Figure 4. Posterior predictive checks for model validation. The posterior predictive default rate closely matches the observed default rate of 22.1%. Calibration plots compare predicted default probabilities (unit interval) with observed frequencies across probability bins. Predicted probabilities exhibit good discrimination between defaulters and non-defaulters. Residual diagnostics show a mild V-shaped pattern at the extremes of the credit limit distribution, suggesting weak nonlinearities that do not materially affect the estimated average treatment effect.

The calibration plot (top-right panel) demonstrates close agreement between predicted default probabilities and observed frequencies across probability bins, with deviations

confined to the extremes. The bottom-left panel shows clear separation between predicted probabilities for defaulters and non-defaulters, indicating good discriminatory power. Residual diagnostics (bottom-right panel) reveal a mild V-shaped pattern at the boundaries of the predicted probability range, suggesting weak nonlinearities; however, these deviations are limited to the extremes and do not materially affect the estimated average treatment effect or the substantive conclusions of the analysis.

The implications of this sensitivity analysis are visualized in Figure 5. The left panel shows the bias-adjusted ATE across increasing strengths of an unmeasured confounder, with the null effect boundary (ATE = 0) indicated for reference. Even under moderate confounding scenarios, the estimated treatment effect remains negative, indicating substantial robustness of the causal conclusion. The right panel delineates the robustness and reversal regions, highlighting the combinations of confounder–treatment correlation and confounder–outcome effect size (on the log-odds scale) required to overturn the observed effect. Consistent with the E-value analysis, only confounders with implausibly strong associations with both credit limit assignment and default risk would be sufficient to eliminate the estimated causal impact.




Figure 5. Sensitivity analysis for unmeasured confounding. **Left:** Bias-adjusted Average Treatment Effect (ATE) in probability units across varying confounder strengths, with the null effect boundary (ATE = 0) shown. **Right:** Robustness and reversal regions indicating combinations of confounder treatment correlation and confounder outcome effect size (log-odds scale) required to eliminate the observed causal effect.

5. Discussion

5.1. Principal Findings

This study provides evidence consistent with a causal effect that increasing a customer’s credit limit reduces their probability of default. The results show that a one-standard-deviation increase in credit limit leads to an estimated causal reduction of 1.44 percentage points in default probability (94% HDI: [−2.0%, −1.0%]). The posterior intervals exclude zero; all parameters achieved perfect convergence ($\hat{R} = 1.0$), and the effect was consistent across demographic subgroups defined by age, education, and gender. Moreover, sensitivity analyses indicate that the findings are robust even to strong forms of unmeasured confounding. Collectively, these results challenge the conventional assumption that higher credit limits inherently increase default risk.

Once confounding factors that simultaneously influence credit limit assignment and default behavior are properly accounted for, the evidence instead indicates that higher limits exert a protective effect against default. This magnitude represents a policy-relevant effect within the studied context. For a lender with 100,000 customers at the baseline default rate of 22.1%, a one-standard-deviation credit limit increase would prevent approximately 1440 defaults. Given typical loss-given-default rates in consumer credit markets, this translates to meaningful reductions in credit losses and improved portfolio performance,

justifying strategic credit limit policies informed by causal evidence rather than exposure-based reasoning alone.

5.2. Causal Mechanisms

The protective effect of higher credit limits likely operates through empirically documented mechanisms grounded in consumer financial behavior. Credit utilization, a key component of credit scoring models that reflects the ratio of outstanding balances to available credit, is strongly associated with default risk (Gross and Souleles 2002). By increasing credit limits while holding balances constant, borrowers automatically reduce their utilization ratios, which signals responsible credit management to lenders and may reduce the perceived risk of future default.

Beyond utilization effects, financial flexibility plays a critical role. Field experiments in credit markets demonstrate that repayment flexibility leads to substantial improvements in business outcomes and lower default rates by providing insurance against income fluctuations and enabling better resource allocation during adverse shocks (Battaglia et al. 2024). Corporate finance research establishes that financial flexibility in the form of unused debt capacity allows firms to meet funding needs during financial deficits without forcing immediate rebalancing (Denis and McKeon 2012), a mechanism that extends naturally to consumer credit behavior.

For individual borrowers, higher available credit provides a liquidity buffer during temporary income disruptions or unexpected expenses. Consumer credit research shows that households respond to credit constraints by adjusting consumption and borrowing behavior (Gross and Souleles 2002), suggesting that adequate credit access enables smoother financial management and reduces the likelihood of payment delinquency. Additionally, lower credit utilization ratios improve credit scores, which facilitate access to refinancing options and additional credit sources at more favorable terms, further reducing default risk. Research on credit report information confirms that improved credit profiles lead to economically significant increases in credit access and reductions in subsequent default probability Dobbie et al. (2020).

While our data do not contain direct measurements of utilization ratios before credit limit increases, the consistency of our findings with this established theoretical and empirical literature supports financial flexibility and utilization effects as plausible causal pathways. Future research employing mediation analysis with time-varying utilization data and direct measures of liquidity constraints would strengthen causal claims regarding these specific mechanisms.

5.3. External Validity and Temporal Considerations

This study relies on credit card data from Taiwan covering April to September 2005, a period preceding major structural changes in global credit markets and risk management practices. Several considerations are therefore relevant when interpreting the external validity of our findings in contemporary settings.

First, the 2008 Global Financial Crisis fundamentally altered lending practices, regulatory oversight, and consumer credit behavior. Post-crisis regulatory frameworks, including Basel III, introduced stricter capital requirements and more conservative risk management standards, which may have affected both credit limit assignment and the relationship between limits and default risk. As a result, our estimates reflect credit dynamics in a pre-crisis environment characterized by different underwriting practices and macroeconomic conditions.

Second, advances in credit risk assessment since 2005 have transformed lending technologies. The widespread adoption of machine learning models, alternative data

sources, and real-time monitoring systems allows lenders to more precisely differentiate borrower risk. Modern credit limit assignment may therefore incorporate information not available in the period studied, potentially modifying the strength of the causal relationship we estimate. Nevertheless, the underlying mechanism of financial flexibility provided by unused credit capacity remains economically relevant, even as risk measurement becomes more granular.

Third, macroeconomic conditions play a central role in shaping credit dynamics. The study period was characterized by relatively stable inflation and moderate interest rates. In environments marked by high inflation, elevated borrowing costs, or heightened economic uncertainty, such as during the COVID-19 pandemic, the protective effect of higher credit limits may either weaken or strengthen. Higher interest rates reduce the value of unused credit as a liquidity buffer, while greater uncertainty may amplify its insurance role. The net effect is therefore context-dependent and requires empirical validation using contemporary data.

Fourth, cross-country generalizability warrants caution. Taiwan's institutional context, including its regulatory environment, consumer protection framework, and cultural attitudes toward debt, may differ from those of other economies. While the magnitude of the estimated effects may not transfer directly to other markets, the directional relationship between credit limits and default risk is likely to be relevant across developed consumer credit systems.

These considerations motivate several avenues for future research. Replication using more recent datasets across diverse institutional settings would help assess whether the estimated effects persist under modern lending technologies and regulatory regimes. Examining heterogeneity across interest rate environments and periods of economic stress would further clarify the conditions under which financial flexibility mechanisms are most effective. In addition, incorporating modern machine-learning-based risk scores as covariates could test whether the causal effect of credit limits remains after accounting for contemporary risk assessment tools.

Despite these limitations, the core finding remains: under the stated causal assumptions, increasing credit limits is estimated to reduce default risk, providing meaningful insight for credit policy design. Financial flexibility operates through fundamental features of consumer credit markets such as liquidity constraints, utilization dynamics, and access to emergency funds that remain relevant across time and institutional contexts. Our results suggest that overly conservative credit limit policies may inadvertently increase default risk, highlighting the importance of policy design informed by causal evidence rather than exposure-based reasoning alone.

Recent advances in causal inference for credit risk further support this perspective. Ren (Ren 2025) demonstrates that DAG-based causal frameworks can identify domain-invariant relationships that generalize across contexts, achieving high performance retention when transferring models between financial credit and health insurance domains. This evidence suggests that causal structures identified through explicit modeling assumptions, such as those employed in this study, may retain relevance beyond specific datasets and time periods, even when predictive relationships evolve.

5.4. Comparison with Existing Literature

Our findings contribute to credit risk literature by applying causal inference methods to estimate treatment effects of credit limit changes, complementing existing predictive and quasi-experimental approaches.

When contrasting predictive modeling against causal inference, we find that most credit risk research focuses on predictive accuracy through machine learning architectures.

Studies such as (Wang et al. 2021) address class imbalance through ensemble methods, and (Kvamme et al. 2018) apply convolutional neural networks to mortgage default prediction. Benchmark studies (Lessmann et al. 2015) systematically compare classification algorithms, prioritizing discriminative performance metrics such as AUC and accuracy. While these approaches excel at identifying high-risk borrowers, they estimate associations rather than causal effects. Our analysis employs formal causal inference methods (Hernán and Robins 2020; Imbens and Donald 2015; Pearl 2009) to estimate the effect of credit limit interventions, addressing a complementary question relevant to policy evaluation.

Recent work in causal inference in consumer credit markets has begun applying causal frameworks to credit market questions. (Fuster et al. 2022) investigate how algorithmic credit screening affects lending decisions and fairness outcomes. Dobbie et al. (2020) use quasi-experimental variation from bankruptcy flag removal to show that improved credit reports lead to increases in credit access and reduced default rates. Their mechanism operates through information revelation rather than direct credit limit changes. (Battaglia et al. 2024) demonstrate in a microfinance context that repayment flexibility reduces default rates, supporting financial flexibility mechanisms. Our analysis estimates the causal effect of credit limit changes specifically, using observational data with explicit adjustment for measured confounders.

Our study integrates DAG-based causal identification (Pearl 2009) with Bayesian inference (Gelman et al. 2013) to quantify posterior uncertainty over causal effects. We combine structural causal models with Bayesian computation (Hoffman and Gelman 2014) and test robustness to unmeasured confounding using sensitivity analysis (VanderWeele and Ding 2017). This approach provides explicit uncertainty quantification and formal assessment of identifying assumptions.

Our findings are consistent with evidence from consumer finance showing that liquidity constraints affect household behavior (Gross and Souleles 2002) and that unused debt capacity provides financial flexibility during adverse shocks (Denis and McKeon 2012). The observed negative treatment effect suggests that higher credit limits may reduce default probability through financial flexibility mechanisms, though our data limitations prevent direct testing of specific pathways such as utilization ratio changes.

Consistent with recent MDPI studies that employ the Taiwanese credit card default dataset as a benchmarking standard (Bhandary and Ghosh 2025), the present work uses this historical dataset to support methodological investigation rather than contemporaneous risk prediction. Whereas Bhandary and Ghosh (Bhandary and Ghosh 2025) focus on predictive performance benchmarking across statistical and machine learning classifiers, our study makes a distinct contribution by applying Bayesian causal inference with an explicitly specified directed acyclic graph (DAG), enabling causal effect estimation, principled posterior uncertainty quantification, and credible interval analysis that are not attainable within purely predictive frameworks.

By estimating the causal effect of an actionable policy variable (credit limits) within a formally specified causal model, our findings highlight that credit risk dynamics cannot be fully understood through exposure-based or predictive reasoning alone.

6. Conclusions

This paper investigated the effect of credit limit increases on default risk using a Bayesian causal inference framework. By combining DAG-based identification, Bayesian logistic regression, and posterior predictive simulation on a dataset of 30,000 credit card customers, we estimated that a one-standard-deviation increase in credit limit is associated with a reduction of approximately 1.44 percentage points in default probability under a

causal intervention. After adjusting for key confounders, the results provide evidence consistent with a protective causal effect of higher credit limits on default outcomes.

Relative to the existing literature, which has largely emphasized predictive performance rather than causal mechanisms, this study contributes by isolating the effect of changes in an actionable policy variable, namely credit limits. Prior research has documented strong associations between credit behavior and default risk, but comparatively few studies have applied formal causal frameworks to evaluate credit limit policies. Our findings therefore extend the literature by providing a more rigorous assessment of how credit limit adjustments influence borrower outcomes.

From a practical perspective, the results suggest that strategically increasing credit limits may reduce default rates by enhancing borrowers' financial flexibility and lowering utilization ratios. These insights can inform credit policy design, supporting risk management strategies that balance exposure considerations with borrower resilience and portfolio performance.

Several limitations should be noted. The analysis relies on data from a single financial institution during a specific historical period and therefore does not capture the full range of economic environments, regulatory regimes, or institutional practices observed in modern credit markets. Although sensitivity analyses indicate robustness to moderate unmeasured confounding, the possibility that omitted factors influence both credit limit assignment and default behavior cannot be fully excluded.

We further assumed a linear relationship between credit limits and default probability on the logit scale. Diagnostic checks suggest that this specification provides a reasonable approximation for estimating average treatment effects, although it may not fully capture localized nonlinearities or heterogeneous effects across the credit limit distribution. More flexible specifications could not be reliably estimated in this setting, and future work with richer data may allow for a more detailed characterization of nonlinear dynamics.

A key remaining concern relates to unmeasured income. While observed financial behavior variables serve as strong proxies and effect estimates are stable across demographic subgroups, income may represent a stronger confounder than can be fully addressed in the present analysis. As a result, the magnitude of the estimated effect should be interpreted with caution, even though reversing the direction of the effect would require implausibly strong confounding. Replication using datasets with direct income measures would strengthen causal conclusions.

Future research could extend this work by examining heterogeneous or nonlinear treatment effects, replicating the analysis across diverse lenders and macroeconomic contexts, and evaluating dynamic credit limit adjustments over time. Greater attention to behavioral mechanisms, including utilization responses and liquidity constraints, would further clarify the channels through which credit limit policies affect default risk.

Overall, the findings suggest that credit limits play a more nuanced role in credit risk management than exposure-based reasoning alone implies. When viewed through a causal lens, increases in credit limits may reduce default risk by improving borrower resilience, highlighting the value of policy decisions informed by causal evidence rather than predictive associations alone.

Author Contributions: Conceptualization, S.D.P.; methodology, S.D.P.; software, S.D.P.; validation, T.M.; formal analysis, S.D.P.; investigation, S.D.P.; resources, S.D.P.; data curation, S.D.P.; writing—original draft preparation, S.D.P.; writing—review and editing, T.M.; visualization, S.D.P.; supervision, T.M.; project administration, T.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study uses publicly available secondary data.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are publicly available. The credit default dataset can be accessed from Kaggle: <https://www.kaggle.com/datasets/uciml/default-of-credit-card-clients-dataset> (accessed on 1 November 2025). Analysis code is available at <https://github.com/pitsojaden/Bayesian-Causal-Analysis---Credit-Default> (accessed on 1 November 2025).

Acknowledgments: The authors acknowledge the use of Claude (Anthropic) for assistance with literature review organization and LaTeX formatting. The authors have reviewed and edited all output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ATE	Average Treatment Effect
CATE	Conditional Average Treatment Effect
DAG	Directed Acyclic Graph
ESS	Effective Sample Size
HDI	Highest Density Interval
MCMC	Markov Chain Monte Carlo
NUTS	No-U-Turn Sampler
SUTVA	Stable Unit Treatment Value Assumption

Appendix A




Figure A1. Positivity diagnostics. **Top:** Credit limit distributions (NT\$) by demographics show substantial overlap across all strata. **Middle:** Propensity model ($R^2 = 0.36$) predicting credit limit from 22 confounders. **Bottom:** Covariate balance (SMD) and credit limit ranges by group. All demographic strata span nearly the full credit limit range (NT\$10K–1M), supporting positivity.




Figure A2. Linearity diagnostics. **Left:** Posterior predictive residuals versus standardized credit limit show random scatter around zero with no systematic curvature. **Right:** Binned residuals across credit limit deciles oscillate around zero (range: -0.015 to $+0.03$) with no monotonic trend, supporting the linear specification.

References

- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebe. 2015. Regulating Consumer Financial Products: Evidence from Credit Cards. *The Quarterly Journal of Economics* 130: 111–64. [\[CrossRef\]](#)
- Battaglia, Marianna, Selim Gulesci, and Andreas Madestam. 2024. Repayment Flexibility and Risk Taking: Experimental Evidence from Credit Contracts. *The Review of Economic Studies* 91: 2635–75. [\[CrossRef\]](#)
- Bhandary, Rakshith, and Bidyut Kumar Ghosh. 2025. Credit Card Default Prediction: An Empirical Analysis on Predictive Performance Using Statistical and Machine Learning Methods. *Journal of Risk and Financial Management* 18: 23. [\[CrossRef\]](#)
- Carroll, Christopher. D. 1997. Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics* 112: 1–55. [\[CrossRef\]](#)
- Denis, David J., and Stephen B. McKeon. 2012. Debt Financing and Financial Flexibility: Evidence from Proactive Leverage Increases. *The Review of Financial Studies* 25: 1897–929. [\[CrossRef\]](#)
- Dobbie, Will, Paul Goldsmith-Pinkham, Neale Mahoney, and Jae Song. 2020. Bad Credit, No Problem? Credit and Labor Market Consequences of Bad Credit Reports. *The Journal of Finance* 75: 2377–419. [\[CrossRef\]](#)
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. Predictably Unequal? The Effects of Machine Learning on Credit Markets. *The Journal of Finance* 77: 5–47. [\[CrossRef\]](#)
- Gelman, Andrew, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. 2013. *Bayesian Data Analysis*, 3rd ed. Boca Raton: CRC Press.
- Gross, David B., and Nicholas S. Souleles. 2002. Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data. *The Quarterly Journal of Economics* 117: 149–85. [\[CrossRef\]](#)
- Hernán, Miguel A., and James M. Robins. 2020. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC.
- Hoffman, Matthew D., and Andrew Gelman. 2014. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research* 15: 1593–623.
- Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge: Cambridge University Press.
- Khandani, Amir E., KAdlar J. Kim, and Andrew W. Lo. 2010. Consumer Credit-Risk Models via Machine-Learning Algorithms. *Journal of Banking & Finance* 34: 2767–87. [\[CrossRef\]](#)
- Kvamme, Håvard, Nikolai Sellereite, Kjersti Aas, and Steffen Sjursen. 2018. Predicting Mortgage Default Using Convolutional Neural Networks. *Expert Systems with Applications* 102: 207–17. [\[CrossRef\]](#)
- Lessmann, Stefan, Bart Baensens, Hsin-Vonn Seow, and Lyn C. Thomas. 2015. Benchmarking State-of-the-Art Classification Algorithms for Credit Scoring. *European Journal of Operational Research* 247: 124–36. [\[CrossRef\]](#)
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge: Cambridge University Press.
- Ren, Luqing. 2025. Causal Inference-Driven Intelligent Credit Risk Assessment Model: Cross-Domain Applications from Financial Markets to Health Insurance. *Academic Journal of Computing & Information Science* 8: 8–14. [\[CrossRef\]](#)

- VanderWeele, Tyler J., and Peng Ding. 2017. Sensitivity Analysis in Observational Research: Introducing the E-Value. *Annals of Internal Medicine* 167: 268–74. [[CrossRef](#)] [[PubMed](#)]
- Wang, Hong, Qingsong Xu, and Lifeng Zhou. 2021. Large Unbalanced Credit Scoring Using Lasso–Logistic Regression Ensemble. *PLoS ONE* 16: e0246598. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.