

Expected-Loss Provisioning and the Distribution of Mortgage Credit

David Olsen

david.olsen@utahtech.edu

Utah Tech University

Cindy Greenman

Utah Tech University

Research Article

Keywords: CECL, expected-loss provisioning, mortgage credit, bank risk management, credit allocation, HMDA

Posted Date: August 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10696408/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Expected-loss provisioning can affect more than when banks recognize credit losses; it may also be associated with how banks allocate credit. We examine whether bank-year provision pressure after implementation of the Current Expected Credit Loss (CECL) standard is associated with mortgage denial outcomes across institutions, loan purposes, and borrower groups. We link 2018–2024 Home Mortgage Disclosure Act (HMDA) applications to the Philadelphia Fed HMDA Lender File and bank Call Reports. Provision pressure is measured as loan-loss provision expense scaled by gross loans, and the main tests use application-weighted cell regressions. The evidence is concentrated among banks with at least \$10 billion in assets and in nonpurchase mortgage channels, especially home improvement and refinancing. In these settings, higher post-CECL provision pressure is associated with larger Black-White denial differences. Common-sample debt-to-income tests compare identical DTI-defined cells with and without DTI controls, and the focal Black interaction remains positive and statistically significant in all three principal channels. The findings are associational because provision pressure is not randomly assigned and public HMDA does not provide borrower credit scores or banks' internal underwriting information. The study identifies a financial-services setting in which expected-loss provisioning is systematically associated with the distribution of mortgage credit, consistent with a role for formal bank decision systems.

1. Introduction

Expected-loss provisioning can change more than the timing of credit-loss recognition. When banks incorporate expected-loss estimates into underwriting, risk limits, capital planning, and portfolio management, the provision burden associated with credit risk can influence which applications are approved and which are denied. This possibility matters because a bank can maintain aggregate lending while changing the composition of the credit it supplies across products and borrowers. We examine whether bank expected-loss provision pressure is associated with the distribution of mortgage credit following implementation of the Current Expected Credit Loss (CECL) standard.

CECL provides a useful setting for this question because it requires banks to recognize expected lifetime credit losses at or near loan origination rather than waiting for losses to become probable. The change increased the role of forward-looking loss estimates in bank financial reporting and risk governance. Our central mechanism is straightforward. When expected losses become more salient or costly through current-period provisioning, banks may respond by tightening risk tolerances, changing approval thresholds, or reallocating credit toward exposures that impose less expected-loss pressure. These responses need not occur through an explicit rule tied to any particular borrower characteristic. They can operate through the formal systems banks use to evaluate risk, manage portfolios, and control credit exposure.

Prior research establishes that loan-loss accounting can affect bank lending and risk taking. Beatty and Liao (2011) show that delayed loss recognition is associated with sharper reductions in bank lending during downturns, and Bushman and Williams (2012) link provisioning practices to bank risk taking. More recent CECL research likewise shows that expected-loss accounting can affect credit-market behavior. Yang (2025) reports changes in loan growth and mortgage origination strategy around CECL, while Granja and Nagel (2025) show that CECL-induced reserve requirements affected consumer-loan pricing, with stronger effects among more constrained banks and underprivileged borrowers. These studies establish that expected-loss accounting can affect how much banks lend and the price at which they lend. They leave open a related question that is central to financial intermediation: does expected-loss pressure also affect how banks allocate credit across applications, mortgage products, and borrower groups?

We test that question in three steps. First, we examine whether post-CECL provision pressure is associated with mortgage denial outcomes. Second, we test whether the association is stronger in settings where formal bank risk-management systems should transmit expected-loss pressure more directly into credit decisions. This prediction points especially to large banks, which rely more heavily on centralized underwriting, hard information, formal risk models, and portfolio-level controls (Stein 2002; Berger and Udell 2002; Berger et al. 2005; Liberti and Petersen 2019). We also examine mortgage-purpose channels because purchase, refinancing, and home-improvement credit serve different household and portfolio functions. Third, after identifying the settings in which the provision-pressure relation is strongest, we test whether the association is distributed differently across borrower groups. This sequence is important: the paper first establishes where expected-loss pressure is associated with credit allocation and then asks who bears more of that association.

The empirical design links three public data sources that do not answer the research question separately. Home Mortgage Disclosure Act (HMDA) records provide mortgage application outcomes, borrower race and ethnicity, loan purpose, loan type, income, loan amount, and property geography. The Philadelphia Fed HMDA Lender File provides the identifier bridge needed to connect HMDA reporters to insured depository institutions. FDIC/FFIEC Call Reports provide bank-year accounting and financial-condition measures, including loan-loss provisions, allowances, charge-offs, past-due loans, capital, profitability, and bank size. The resulting data infrastructure contains more than

122 million mortgage application records in the national build, more than 43 million applications matched to FDIC-insured banks, and approximately 31.6 million applications represented in the principal weighted-cell models. Our primary bank-year measure is provision pressure, defined as loan-loss provision expense scaled by gross loans.

The results are concentrated rather than uniform. The strongest evidence appears among banks with at least \$10 billion in assets, consistent with the prediction that formalized credit-risk systems provide an important transmission setting. Across mortgage purposes, the most consistent results occur in nonpurchase credit, especially home-improvement and refinancing loans, rather than home-purchase lending. These channels matter because mortgage credit affects households after the initial purchase decision. Refinancing can reduce borrowing costs and help households manage payment risk, while home-improvement credit can support property maintenance, collateral preservation, and the continuing financial value of homeownership. The concentration of the results in these channels therefore points to an allocation effect that extends beyond entry into homeownership.

The distributional tests show that the provision-pressure association is especially pronounced for Black applicants relative to White applicants in the large-bank and nonpurchase settings where the overall relation is strongest. Results for Hispanic and Asian applicants are less consistent, so the evidence does not support a broad claim that expected-loss pressure affects all non-White groups in the same way. We also do not interpret the Black-White result as evidence that CECL caused intentional discrimination. A common tightening in credit standards can have unequal effects if borrower groups differ in their distribution around approval margins because of differences in wealth, collateral, income stability, prior credit access, geography, or other risk-related conditions. The empirical result is therefore a distributional finding about credit allocation, not a claim about lender intent.

We address several alternative explanations directly. The analysis uses observable borrower, loan, geographic, tract, and bank controls and evaluates the results under alternative clustering structures, top-bank exclusions, leave-one-bank-out diagnostics, alternative treatments of extreme provision observations, placebo timing tests, and alternative credit-loss measures. Additional specifications examine observable underwriting risk and compare applicants within bank-year settings. These tests reduce concerns that the results are driven solely by lender concentration, outliers, persistent pre-CECL bank differences, or broad bank distress. They do not eliminate the limits of public mortgage data. In particular, public data do not observe all information used in underwriting, including borrower credit scores. For that reason, CECL timing and provision pressure are not treated as randomly assigned shocks, and the estimates are interpreted as conditional associations rather than causal estimates of discrimination or borrower treatment.

The paper contributes to research on financial intermediation by showing that expected-loss provisioning is associated not only with aggregate credit supply or pricing, but also with the allocation of household credit. The evidence identifies a setting in which bank accounting and risk-management systems intersect: current-period provision pressure is most strongly associated with mortgage allocation at large institutions and in specific lending channels, and the resulting association is not evenly distributed across applicants. This perspective shifts the question from whether expected-loss accounting changes the total amount of credit banks provide to whether it changes which applications and products are more likely to receive credit. More broadly, the findings suggest that provisioning systems can have distributional consequences when accounting measures become inputs to formal credit-risk management.

2. Institutional Setting, Theory, and Predictions

2.1 Expected-Loss Provisioning and Bank Credit Decisions

Loan-loss accounting affects more than the timing of reported credit losses. Because provisions and allowances enter capital planning, risk governance, and lending decisions, the way banks recognize credit deterioration can affect credit supply. Earlier research shows that banks often provision too little in expansions and too much in downturns, creating procyclical loss recognition (Laeven and Majnoni 2003; Bikker and Metzmakers 2005). Beatty and Liao (2011) show that banks with more delayed loss recognition reduce lending more sharply during recessions. Kim (2022) extends this logic to housing markets, finding that areas more exposed to banks with delayed loss recognition experienced larger mortgage-supply contractions, more distressed sales, and larger house-price declines. Together, this work establishes a direct link between credit-loss recognition and real lending outcomes.

The Current Expected Credit Loss (CECL) standard shifts that link forward. CECL requires banks to recognize expected lifetime credit losses rather than waiting for losses to become probable (Financial Accounting Standards Board 2016). Supervisory guidance likewise connects expected-loss measurement to credit-risk management, governance, and regulatory oversight (Bank for International Settlements, Basel Committee on Banking Supervision 2015; U.S. Office of the Comptroller of the Currency et al. 2016). In this setting, current-period provision expense can become an operating signal of expected credit risk rather than merely an after-the-fact accounting outcome.

We capture that signal with provision pressure, defined as loan-loss provision expense scaled by gross loans. Provision pressure is an income-statement measure and differs from allowance balances, which are stock measures, and from charge-offs and past-due loans, which more directly reflect realized or already-emerging credit deterioration. Prior research shows that provisions reflect expected credit risk as well as regulatory incentives, managerial judgment, and reporting discretion (Kanagaretnam et al. 2003; Liu and Ryan 2006; Fonseca and González 2008; Bushman and Williams 2012; Beatty and Liao 2014). We therefore do not treat provision pressure as a pure measure of borrower risk. Instead, we use it as a bank-year measure of the credit-loss pressure being recognized through the bank's expected-loss system.

The mechanism follows the broader real-effects literature: accounting measures can change behavior when they enter contracts, regulation, risk management, investment, or financing decisions (Zhang 2008; Kanodia and Saprà 2016; Roychowdhury, Shroff and Verdi 2019). Expected-loss accounting is especially relevant because it changes the ex ante cost and salience of holding credit-risky assets. Mahieux, Saprà and Zhang (2023) show theoretically that expected-loss accounting can affect loan origination. International evidence under expected-credit-loss regimes shows changes in the information content of provisions and in credit contraction when expected-loss provisions rise (López-Espinosa, Ormazabal and Sakasai 2021; Chen and Huang 2026). In the United States, Yang (2025) finds that capital-constrained banks changed loan growth and mortgage origination strategy around CECL, while Granja and Nagel (2025) show that CECL-induced reserve requirements increased consumer-loan interest rates, with stronger effects for weakly capitalized banks and underprivileged borrowers.

These studies show that expected-loss accounting can affect how much banks lend and the price of credit. Our question is different but closely related: whether expected-loss provision pressure is also associated with which mortgage applications receive credit. A bank can maintain aggregate lending while changing approval thresholds, product mix, or the set of marginal applications it accepts. If higher provision pressure increases the measured or perceived cost of credit risk, banks may tighten risk tolerances or reallocate credit even when total lending changes little.

Prediction 1. In the post-CECL period, higher provision pressure is associated with tighter mortgage-credit allocation, measured through higher denial outcomes or larger denial differences in the settings where expected-loss pressure is most likely to enter lending decisions.

2.2 Large-Bank Credit Systems and Provisioning Transmission

The transmission of provision pressure should depend on how banks make lending decisions. Banking research distinguishes between soft, relationship-based information and hard, codified information that can be transmitted through larger organizations. Stein (2002) develops a theory in which hierarchical organizations favor hard information. Berger and Udell (2002) argue that large banks are less suited to relationship lending, while Berger, Miller, Petersen, Rajan and Stein (2005) document systematic differences in lending practices across large and small banks. Petersen and Rajan (2002) show that lending distance increased with information technology, and Liberti and Petersen (2019) synthesize the broader evidence on hard versus soft information in lending.

This distinction gives the large-bank tests a direct theoretical role. Large institutions are more likely to use centralized underwriting, formal credit models, portfolio-level controls, and standardized risk tolerances. Those systems provide a natural channel through which provisions, expected-loss estimates, capital effects, and portfolio constraints can influence approval decisions. Smaller institutions can rely more heavily on local knowledge and borrower-specific relationships, which may weaken the transmission from a bank-level accounting signal to an individual mortgage decision.

Research on automated and model-based lending reinforces this point. Fuster, Goldsmith-Pinkham, Ramadorai and Walther (2022) show that predictive modeling can change credit-market disparities because model inputs and flexibility affect approval patterns. Howell, Kuchler, Snitkof, Stroebel and Wong (2024) show that automation can reduce some racial disparities by limiting human discretion, while also emphasizing that outcomes depend on the data and rules embedded in the system. Sirignano, Sadhwani and Giesecke (2018) demonstrate the complexity of mortgage-risk prediction in large-scale models. The implication is not that formal systems necessarily increase disparities. It is that system design can transmit common risk signals broadly and can produce heterogeneous outcomes without requiring race-specific decision rules.

Prediction 2. The association between post-CECL provision pressure and mortgage-credit allocation is stronger at large banks, where centralized and formalized credit-risk systems are more likely to transmit expected-loss signals into approval decisions.

2.3 Mortgage-Credit Channels

Mortgage credit is not a single decision. Home-purchase loans determine entry into homeownership, while refinancing and home-improvement credit affect households after a home has already been acquired. Refinancing can change borrowing costs, loan terms,

liquidity, and payment risk. Home-improvement credit can finance maintenance and investment that preserve property quality and collateral value. For this reason, a bank-level tightening mechanism need not affect all mortgage purposes equally.

The distinction also matters for household wealth. Kermani and Wong (2026) show that distressed sales reduce realized housing returns for Black and Hispanic homeowners and emphasize the role of liquidity and income stability. Ghent, Hernández-Murillo and Owyang (2014) document racial differences in subprime mortgage pricing across neighborhoods, and Avenancio-León and Howard (2022) show that racial inequities in property taxation impose different effective tax burdens on otherwise comparable properties. These studies underscore that the financial return to homeownership depends on continued access to credit and liquidity, not only on the initial purchase transaction.

Accordingly, we treat loan purpose as part of the economic mechanism rather than as an after-the-fact subsample exercise. We compare home-purchase lending with home-improvement, refinancing, cash-out refinancing, and other mortgage-purpose channels. The theory does not require every nonpurchase category to respond identically. Instead, it predicts that the effect of provision pressure can vary across products because the underlying borrower needs, collateral positions, and bank portfolio considerations differ.

2.4 Distributional Consequences of Credit Allocation

Mortgage markets also provide a setting in which a common change in credit standards can have unequal consequences across borrower groups. A long literature documents persistent racial and ethnic differences in mortgage outcomes. Munnell, Tootell, Browne and McEneaney (1996) provide foundational evidence on Black and Hispanic mortgage denials, while Ross and Yinger (2002) emphasize the difficulty of separating differences driven by underwriting risk from differential treatment. More recent studies continue to document disparities in richer-data and modern-lending settings, including mortgage pricing, lender segmentation, broker and borrower race, and denial outcomes (Bayer, Ferreira and Ross 2018; Ambrose, Conklin and Lopez 2021; Bartlett, Morse, Stanton and Wallace 2022; Lewis-Faupel and Tenev 2024; Bhutta, Hizmo and Ringo 2025).

That literature motivates our borrower-group tests, but it also defines their boundary. Public HMDA data provide broad national coverage and important borrower and loan characteristics, but they do not reveal every underwriting input or internal lender decision rule (Avery, Brevoort and Canner 2007). Public data also do not disclose borrower credit scores. We therefore do not use racial differences in denial outcomes as a direct measure of discrimination. The relevant question is whether the association between bank-level provision pressure and mortgage-credit allocation differs across borrower groups after accounting for the observable information available in the public data.

A heterogeneous response does not require a race-specific rule. A common increase in risk thresholds can affect groups differently if their applications are distributed differently around approval margins. Differences in wealth, collateral, income stability, prior access to credit, and neighborhood conditions can influence where applicants fall relative to those margins. The housing and wealth literature shows why these differences can be persistent and economically important. Charles and Hurst (2002) connect homeownership transitions to the Black-White wealth gap. Aaronson, Hartley and Mazumder (2021) document lasting effects of historical redlining, Chetty, Hendren, Jones and Porter (2020) show persistent racial differences in economic opportunity, and Derenoncourt, Kim, Kuhn and Schularick (2024) trace the long-run persistence of the U.S. racial wealth gap.

This mechanism is especially relevant for continuing access to housing finance. If broad credit tightening changes access to refinancing or home-improvement credit, the consequences can extend to payment risk, property maintenance, liquidity, and the preservation of housing wealth. The paper therefore treats the Black-White comparison as an important distributional test of the allocation mechanism, while also reporting Hispanic and Asian results to determine whether the pattern is broad or group-specific.

Prediction 3. The association between post-CECL provision pressure and mortgage-credit allocation can differ across borrower groups. The primary distributional test examines whether the association is stronger for Black applicants than for White applicants in the settings where the overall provision-pressure relation is strongest.

2.5 Research Design Boundaries and Interpretation

The empirical predictions are intentionally associational. CECL implementation overlapped with COVID-19, unusual monetary and fiscal conditions, changing mortgage demand, and substantial housing-market volatility. Adoption timing, bank size, portfolio composition, and provision pressure are also not randomly assigned. More generally, recent econometric research shows that staggered difference-in-differences designs can be difficult to interpret when treatment effects vary across cohorts or over time (de Chaisemartin and D'Haultfoeuille 2020; Callaway and Sant'Anna 2021; Goodman-Bacon 2021; Sun and Abraham 2021; Baker, Larcker and Wang 2022; Gardner 2022; Borusyak, Jaravel and Spiess 2024).

For these reasons, we do not treat CECL timing as a clean causal shock. The post-CECL indicator identifies the period in which expected-loss measurement is more likely to be embedded in bank provisioning and risk-management systems, while provision pressure captures

continuous bank-year variation in recognized credit-loss pressure. The empirical tests ask whether that pressure is associated with mortgage-credit allocation, whether the association is stronger where the proposed transmission mechanism should be more important, and whether the association is distributed differently across borrowers.

This framing separates the paper's claims from stronger conclusions the design cannot support. A positive Black-White interaction does not establish intentional discrimination, and a large-bank interaction does not show that size itself causes the result. Likewise, the paper does not require provision pressure to be a pure exogenous treatment. The contribution is narrower: to identify whether an economically meaningful bank provisioning signal is systematically associated with the allocation of mortgage credit in the period after expected-loss accounting became operative, and to test the institutional and distributional patterns implied by the proposed mechanism.

3. Data and Sample Construction

The empirical design requires three types of information that are not available in a single public source: mortgage application outcomes and borrower characteristics, a reliable link between mortgage lenders and regulated banks, and bank-year measures of credit-loss provisioning and financial condition. We combine Home Mortgage Disclosure Act (HMDA) application data, the Federal Reserve Bank of Philadelphia HMDA Lender File, and FDIC/FFIEC Call Report data. The linked structure allows us to connect applicant-level mortgage outcomes to the bank-year provisioning environment in which those applications were evaluated.

3.1 Data Sources

The mortgage application data come from public HMDA records (Consumer Financial Protection Bureau 2024; Federal Financial Institutions Examination Council 2024b). HMDA reports application outcomes and lender identifiers together with borrower race and ethnicity, loan purpose, loan type, loan amount, applicant income, property geography, tract characteristics, and other loan attributes used in the analysis. The main outcome is mortgage denial. We code an application as denied when the reported action indicates a denial and as zero when the application is approved or originated within the model sample. Applications that are withdrawn, incomplete, or otherwise outside the approved/originated-versus-denied comparison are excluded from the model denominator.

We classify applicants into four primary groups: White, Black, Hispanic, and Asian. White applicants are the reference group in the borrower-group comparisons. HMDA also identifies mortgage purpose. We distinguish home-purchase, home-improvement, refinancing, cash-out refinancing, and other-purpose applications because these products serve different financial functions and are central to the channel predictions developed in Section 2. We also retain loan-type information, including conventional, FHA-insured, VA-guaranteed, and USDA/RHS/FSA-guaranteed lending, for validation and robustness analyses.

Bank accounting and financial-condition data come from quarterly FDIC/FFIEC Call Reports (Federal Deposit Insurance Corporation 2024; Federal Financial Institutions Examination Council 2024a). These reports provide the credit-loss, capital, profitability, and bank-size measures used to characterize the lender-year environment. Because HMDA outcomes are organized by application year, we construct annual bank-year accounting measures from the quarterly Call Report data and align them with mortgage applications from the same year.

Table 1
HMDA Variable Definitions and Coding

Variable/code	Source field	Coding	Use in analysis
Denied	HMDA action taken / derived denial indicator	1 for denied applications; 0 for approved/originated applications in the model sample	Main dependent variable
Borrower group	HMDA race/ethnicity fields	White, Black, Hispanic, Asian	Main borrower-group comparisons; White is the reference group
Loan purpose = 1	HMDA loan purpose	Home purchase	Used to separate initial home-purchase credit
Loan purpose = 2	HMDA loan purpose	Home improvement	Key nonpurchase channel
Loan purpose = 31	HMDA loan purpose	Refinancing	Key nonpurchase channel
Loan purpose = 32	HMDA loan purpose	Cash-out refinancing	Secondary nonpurchase channel; evidence is more mixed
Loan purpose = 4	HMDA loan purpose	Other purpose	Supportive nonpurchase channel
Loan type = 1	HMDA loan type	Conventional	Key validation channel
Loan type = 2	HMDA loan type	FHA-insured	Used to assess standardized-underwriting channel
Loan type = 3	HMDA loan type	VA-guaranteed	Supplemental loan-type validation
Loan type = 4	HMDA loan type	USDA/RHS/FSA-guaranteed	Supplemental loan-type validation
<i>Note. This table defines the main HMDA variables used to construct the mortgage outcome, borrower groups, loan-purpose channels, and loan-type validation tests. Loan-purpose and loan-type codes follow the coding used in the analysis files.</i>			

Table 2
FDIC Call Report Variable Definitions

Variable	Construction	Interpretation
Provision pressure	Loan-loss provision expense / gross loans	Primary accounting-pressure measure; income-statement recognition of expected credit losses
Allowance ratio	Allowance for credit losses / gross loans	Balance-sheet reserve measure; related but more stock-like than provision expense
Net charge-offs	Net charge-offs / gross loans	Backward-looking realized credit-loss measure
Past-due loans	30 + or 90 + past-due loans / gross loans	Credit deterioration measure
Total assets	FDIC Call Report total assets	Used for bank-size partitions and CECL timing proxy
Equity-to-assets	Equity / total assets	Capital-position control or diagnostic
Return on assets	Net income / assets	Profitability control or diagnostic
<i>Note. Provision pressure is the primary accounting-pressure variable because it captures income-statement recognition of expected credit losses and is most directly connected to the CECL mechanism. Related variables are used to test whether the result reflects a broader credit-loss pressure environment.</i>		

3.2 Lender Matching

HMDA and Call Reports do not share a common lender identifier, so the analysis requires a crosswalk between mortgage reporting institutions and regulated banks. We use the Federal Reserve Bank of Philadelphia HMDA Lender File (Federal Reserve Bank of Philadelphia 2025) to bridge the two systems. The file links HMDA reporting institutions to parent organizations and bank regulatory identifiers, allowing us to identify HMDA lenders associated with FDIC-insured banks.

The matching process proceeds in two stages. We first connect HMDA lender identifiers to the Philadelphia Fed lender file. We then use the bank identifiers supplied by that file, with FDIC certificate number as the principal merge key, to connect the matched lender to Call Report

information. The resulting lender-year link allows bank accounting variables to be attached to individual mortgage applications by year. This step is essential to the research question: HMDA alone identifies mortgage outcomes but not bank provision pressure, while Call Reports identify provision pressure but not the borrower-level mortgage applications associated with it.

3.3 Main Variables

The primary bank-level variable is provision pressure, defined as loan-loss provision expense divided by gross loans. Provision pressure is an income-statement measure of recognized credit-loss expense relative to the size of the loan portfolio and is the measure most directly connected to the expected-loss mechanism developed in Section 2. We also retain the allowance-to-gross-loans ratio, net charge-offs scaled by gross loans, past-due loan measures, total assets, equity-to-assets, and return on assets. These related variables help distinguish the provision-pressure result from broader credit deterioration, realized losses, capital position, or general bank distress.

The main borrower-level outcome is the denial indicator described above. Borrower group and loan purpose provide the two principal dimensions of heterogeneity. Bank size is also central because the theory predicts stronger transmission of expected-loss signals through the formalized credit-risk systems of large banks. The empirical models additionally use observable borrower, loan, geographic, tract, and bank characteristics described in Section 4. Table 1 summarizes the HMDA coding used for the dependent variable, borrower groups, loan-purpose channels, and loan-type validation tests; Table 2 summarizes the principal Call Report measures.

3.4 Sample Construction and Final Analysis Sample

We construct the linked sample in four steps. First, we process the national HMDA files and retain applications with usable information for the mortgage outcome and the characteristics required by the analysis. Second, we connect HMDA lenders to FDIC-insured banks using the Philadelphia Fed lender file. Third, we construct annual bank-year accounting measures from the quarterly Call Reports. Fourth, we merge those bank-year measures to the matched HMDA applications. The final application-level records therefore combine the mortgage outcome and borrower and loan characteristics with the accounting and financial condition of the associated bank-year.

Table 3 reports the sample funnel. The national compacted HMDA build contains 122,961,061 application records. Matching HMDA lenders to FDIC-insured banks yields 43,599,577 application rows. After applying the model-sample requirements, the main raw analysis universe contains 33,020,189 mortgage applications. The principal weighted-cell models represent 31,644,205 applications in 453,873 regression cells after minimum-cell filtering. The placebo timing analysis separately represents 9,992,643 pre-CECL applications from 2018–2019 with future provision pressure attached for the falsification tests.

The difference between the application count and the regression-cell count reflects the aggregation used for the main weighted models rather than a loss of application-scale information. As described in Section 4, applications are grouped into comparable bank-year, borrower-group, loan, and geographic cells, and each cell is weighted by the number of applications it represents. This structure makes the national analysis computationally feasible while keeping the estimates tied to the underlying mortgage applications.

Table 4 provides descriptive evidence on the model-ready sample by borrower group, including denial rates, loan amounts, income, loan-to-value ratios, tract characteristics, and the number of banks represented. These descriptive differences reinforce the need for the observable controls and risk-oriented diagnostics developed in the next section. They are not interpreted as causal group differences; rather, they describe the borrower and loan composition of the public-data sample used in the empirical tests.

Table 3
Sample Construction and Linked Data Scale

Item	Count	Description
National compacted HMDA application build	122,961,061	Application records compiled from the national HMDA build before the FDIC-bank match.
FDIC-bank matched HMDA application rows	43,599,577	Applications linked to FDIC-insured banks through the Philadelphia Fed HMDA Lender File and FDIC bank identifiers.
Main raw model universe	33,020,189	All-loans model sample before weighted-cell aggregation and minimum-cell filtering.
Main weighted-cell application count	31,644,205	Applications represented by the weighted least squares cells after requiring a minimum number of applications in each cell.
Main weighted-cell regression cells	453,873	Cell-level regression observations used in the weighted models. Each cell groups applications by bank, year, borrower group, and loan characteristics.
Placebo preperiod application count	9,992,643	Pre-CECL 2018–2019 applications with future provision pressure attached for placebo timing tests.
<i>Note. This table summarizes the scale of the linked HMDA-FDIC data and the model-ready analysis samples. Counts come from the sample-construction outputs and summarize the application-level and cell-level samples used in the main and placebo analyses.</i>		

Table 4
Observable Risk Proxy Summary by Borrower Group

Borrower group	Applications	Denial rate	Avg. loan amount	Avg. income (\$000s)	Avg. LTV (%)	Avg. tract income (% MSA)	Avg. tract minority %	Banks
Asian	2,578,300	19.5%	\$428,767	203.7	168.8	128.4	47.5	2,591
Black	2,235,851	32.1%	\$191,845	124.0	154.3	97.3	56.7	2,568
Hispanic	2,999,310	29.9%	\$211,552	202.5	125.3	103.7	53.7	2,767
White	25,206,728	16.5%	\$260,911	238.8	104.9	121.1	22.1	3,058
<i>Note. This table reports descriptive statistics from the model-ready public-HMDA sample. Applicant income is reported in thousands of dollars, tract income is expressed relative to the metropolitan-area median, and LTV is shown as the raw numeric public-HMDA measure without winsorization in this descriptive table. Public HMDA does not disclose borrower credit scores. Debt-to-income ratio is available in privacy-modified form and is incorporated in the targeted common-sample DTI robustness tests reported in Table 8. The statistics shown here describe baseline borrower, loan, and tract composition and are not interpreted as underwriting cutoffs.</i>								

Section 4 develops the empirical design used to test whether post-CECL provision pressure is associated with mortgage-credit allocation, whether the association is stronger in the institutional and product settings predicted by the theory, and whether the association differs across borrower groups.

4. Research Design

4.1 Empirical Architecture and CECL Timing

The empirical design asks whether bank-year provision pressure is associated with mortgage-credit allocation after expected-loss accounting became operative, whether the association is stronger in the institutional and product settings predicted by the theory, and whether the association differs across borrower groups. The analysis uses HMDA applications from 2018 through 2024. The pre-CECL years provide the basis for timing and placebo diagnostics, while the post-CECL period identifies the setting in which expected-loss provisioning is more likely to be embedded in bank risk-management and lending systems.

Following the existing design, we use an asset-size-based proxy for CECL timing. Banks with at least \$1 billion in assets are classified as post-CECL beginning in 2020, while smaller banks are classified as post-CECL beginning in 2023. This timing indicator is not treated as a clean causal treatment. CECL implementation overlapped with the COVID-19 period, unusual monetary and fiscal conditions, changing mortgage demand, and substantial housing-market volatility. The post indicator therefore identifies the period in which expected-loss accounting is more likely to be operative in bank provisioning systems; the focal explanatory variation comes from continuous bank-year provision pressure.

The \$10 billion asset threshold serves a different purpose. It identifies the large-bank segment in which centralized underwriting, formal risk models, portfolio-level controls, and scalable approval systems should make the transmission mechanism developed in Section 2 most visible. We therefore treat the \$10 billion partition as a mechanism test rather than as an alternative CECL adoption rule.

4.2 Main Empirical Specification

The main outcome is the mortgage denial rate. In the weighted-cell specifications, applications are grouped by bank, year, borrower group, and loan and geographic characteristics. The core model interacts borrower group with the post-CECL period and bank provision pressure. White applicants are the reference group.

The primary coefficient is the Black x Post-CECL x Provision Pressure interaction. A positive coefficient means that, in the post-CECL period, higher bank provision pressure is associated with a larger Black-White difference in mortgage denial rates. Parallel Hispanic and Asian interactions provide the corresponding comparisons with White applicants. The specification also includes all relevant lower-order terms, continuous observable controls, and categorical controls.

$$\begin{aligned} DenialRate_{bgyc} = & \alpha + \beta_1(Black_g \times PostCECL_{by} \times ProvisionPressure_{by}) \\ & + \beta_2(Hispanic_g \times PostCECL_{by} \times ProvisionPressure_{by}) + \beta_3(Asian_g \times PostCECL_{by} \times ProvisionPressure_{by}) \\ & + lower - order\ terms + X_{bgyc}\Gamma + C_{bgyc} + \epsilon_{bgyc}. \quad (1) \end{aligned}$$

In Eq. (1), X contains the continuous borrower, loan, tract, and bank controls described below, while C contains categorical indicators for activity year, loan type, loan purpose, lien status, occupancy status, and state when those dimensions vary in the relevant sample. Lower-order terms include the corresponding main effects and two-way interactions.

The empirical sequence follows the theory. We first examine the provision-pressure association across the linked sample and then ask whether it is concentrated in the large-bank segment and in particular mortgage-purpose channels. The borrower-group interactions are interpreted within those institutional and product settings rather than as stand-alone tests of discrimination. Because provision pressure, bank size, borrower composition, and CECL timing are not randomly assigned, all interaction coefficients are interpreted as conditional associations.

4.3 Observable Controls and Underwriting Risk

The main models control for observable borrower, loan, neighborhood, geographic, and bank characteristics that can affect mortgage approval decisions and may also be correlated with the banks to which applicants apply. Borrower and loan controls include loan amount, applicant income, loan-to-value ratio, loan type, loan purpose, lien status, and occupancy status. Geographic and neighborhood controls include state, tract income relative to the metropolitan area, and tract minority population share. Bank-level variables are used to characterize size and financial condition and to distinguish provision pressure from related credit-loss measures.

These controls are intended to reduce observable compositional differences, not to recreate a lender's complete underwriting file. Credit score remains an important unobserved underwriting variable because it is not available in the public loan-level HMDA data used here. Debt-to-income ratio (DTI), however, is present in the public HMDA data used by the project and is incorporated in a targeted underwriting-risk robustness test.

The underwriting-risk analysis is cumulative. In addition to the full observable-control specification, we examine conventional versus FHA-insured loans, observable-risk partitions, within-bank-year models, and the DTI tests described next. None of these tests eliminates all unobserved underwriting information. Together, they ask whether the focal provision-pressure association survives progressively stronger adjustments for observable borrower risk and lender composition.

4.4 Debt-to-Income Ratio and the Common-Sample Test

The HMDA data used in this study contain DTI in a privacy-modified form. We map the observed public-HMDA values into six ordered categories: less than 20 percent; 20 to less than 30 percent; 30 to less than 36 percent; 36 to less than 50 percent; 50 to 60 percent; and greater than 60 percent. Applications for which DTI is missing, not applicable, or exempt are retained as a separate category in the all-application analysis. This approach uses the underwriting information in DTI without imposing false continuous precision on privacy-modified values.

The principal DTI test is designed to separate the effect of controlling for DTI from a change in sample composition. For each focal channel, we first rebuild the weighted cells with DTI category included in the cell definition. We then estimate two models on exactly the

same cell rows, weights, and bank clusters. Model A uses the full set of existing controls but omits DTI-category indicators. Model B adds DTI-category indicators. Because the estimation sample is identical across the pair, the difference between Models A and B reflects the incremental DTI adjustment rather than a change in which applications survive the minimum-cell requirement.

We implement this paired design for three focal samples: banks with at least \$10 billion in assets across all loan purposes, home-improvement applications at those banks, and refinancing applications at those banks. In the all-application common samples, the paired models use 426,363 cells representing 18,195,322 applications for large banks overall; 74,400 cells representing 2,996,744 home-improvement applications; and 89,806 cells representing 4,768,279 refinancing applications. The two models within each pair have identical cell counts, application weights, and bank clusters.

We repeat the same paired comparison after restricting the data to applications with reported DTI. The reported-DTI common samples contain 402,760 cells representing 15,693,287 applications for large banks overall; 70,367 cells representing 2,793,304 home-improvement applications; and 83,158 cells representing 3,887,511 refinancing applications. This second comparison addresses the possibility that a missing-DTI category drives the all-application result. Section 5 reports the coefficients from both common-sample comparisons.

4.5 Weighted Cell Estimation, Categorical Controls, and Inference

The principal models use weighted least squares on aggregated application cells. The baseline cells group applications by bank, year, borrower group, loan type, loan purpose, lien status, occupancy status, state, and CECL timing status. Each cell is weighted by the number of applications it represents. DTI is added as an additional cell dimension only for the targeted DTI robustness tests described above. A minimum of five applications per cell is required to reduce instability from sparse cells.

The baseline categorical controls include activity-year, loan-type, loan-purpose, lien-status, occupancy-status, and state indicators when those dimensions vary in the relevant sample. We separately estimate within-bank-year specifications as a stronger lender-composition diagnostic. Those models ask whether the distributional pattern remains when borrower-group outcomes are compared within the same bank-year environment rather than across banks with different business models or financial conditions.

Inference emphasizes dependence within lenders. We therefore use bank-clustered standard errors for the principal reported tests and compare them with alternative clustering structures, including bank-year and state clustering where feasible (Bertrand, Duflo and Mullainathan 2004; Cameron, Gelbach and Miller 2011; Cameron and Miller 2015). The purpose is to determine whether the substantive conclusion is sensitive to reasonable assumptions about residual dependence rather than to select the standard-error structure that produces the smallest p-value.

4.6 Alternative Explanations and Diagnostic Tests

The robustness design is organized around specific alternative explanations. Lender concentration is addressed with top-bank exclusions and leave-one-top-bank-out tests. Extreme provisioning observations are addressed with alternative winsorization rules and exclusions. Persistent differences across lender types are evaluated with pre-period placebo tests. Provision pressure is compared with allowance, charge-off, and past-due measures to assess whether the result simply reflects general bank distress or realized credit deterioration.

Underwriting-risk concerns are addressed with the full observable controls, the DTI common-sample tests, loan-type validation, observable-risk partitions, and within-bank-year specifications. These tests are complementary. No single public-data control establishes that otherwise identical applicants received different treatment. The relevant question is whether the provision-pressure association survives a sequence of increasingly demanding comparisons that target the most plausible observable alternatives.

This design preserves the paper's claim boundary. The analysis tests whether expected-loss provision pressure is associated with mortgage-credit allocation and whether that association is concentrated where the proposed banking mechanism predicts. It does not treat the interaction coefficients as proof that CECL caused racial disparities, that banks intentionally discriminated, or that all underwriting risk has been observed.

5. Results

5.1 Overview of Results

The results follow the empirical sequence developed in Sections 2 and 4. The evidence is not uniform across all banks, loan purposes, or borrower groups. The strongest and most consistent provision-pressure associations appear among banks with at least \$10 billion in assets and in nonpurchase mortgage-credit channels, especially home improvement and refinancing. Within those settings, the distributional association is strongest for Black applicants relative to White applicants.

Three features of the evidence are especially important. First, the large-bank concentration is consistent with the prediction that expected-loss signals are more likely to be transmitted through centralized credit systems. Second, the loan-purpose results show that the association is concentrated in continuing financial-services channels of homeownership rather than only in initial home purchase. Third, the Black-White interaction remains positive under richer underwriting-risk controls, including the DTI common-sample tests. The combined evidence therefore supports a credit-allocation interpretation while remaining observational.

5.2 Large Banks and Loan-Purpose Channels

Table 5 reports the main large-bank and loan-purpose estimates. Using bank-clustered inference, the Black x Post-CECL x Provision Pressure coefficient is 2.900 in the sample of banks with at least \$10 billion in assets (standard error 0.928, $p = 0.002$), representing 22,940,229 mortgage applications across 149 bank clusters. The positive coefficient means that, after CECL, higher provision pressure is associated with a larger Black-White denial difference in the large-bank segment. Alternative state, bank-year, and year clustering produce the same substantive conclusion.

The loan-purpose tests sharpen this result. For home-improvement applications at large banks, the Black interaction is 2.796 (standard error 0.560, $p < 0.001$) across 3,522,274 applications and 132 bank clusters. For refinancing, the coefficient is 4.109 (standard error 1.577, $p = 0.009$) across 5,775,253 applications and 148 bank clusters. Other-purpose lending also produces a positive coefficient of 2.845 ($p < 0.001$). Cash-out refinancing has a positive point estimate of 5.742, but its bank-clustered p -value is 0.082, so we treat that channel as suggestive rather than as part of the strongest evidence.

The concentration in home improvement and refinancing is economically meaningful. Home-improvement credit can finance property maintenance and collateral preservation, while refinancing can reduce borrowing costs, change loan terms, and help households manage payment risk and liquidity. The findings therefore concern the continuing financial value of homeownership as well as access to purchase credit. This pattern is consistent with the channel prediction developed before the results rather than with a uniform tightening across every mortgage product.

Table 5
Main Large-Bank and Loan-Purpose Results

Channel	Black coefficient	SE	p-value	Weighted applications	Bank clusters
Banks $\geq$ \$10B, all loan purposes	2.900	0.928	0.002	22,940,229	149
Home improvement, banks $\geq$ \$10B	2.796	0.560	< 0.001	3,522,274	132
Refinancing, banks $\geq$ \$10B	4.109	1.577	0.009	5,775,253	148
Other purpose, banks $\geq$ \$10B	2.845	0.317	< 0.001	3,327,991	140
Cash-out refinancing, banks $\geq$ \$10B	5.742	3.306	0.082	2,735,120	141

Note. This table reports the primary Black x Post-CECL x Provision Pressure coefficient using bank-clustered standard errors for the focal large-bank and loan-purpose channels. White applicants are the reference group. A positive coefficient indicates that higher provision pressure after CECL is associated with a larger Black-White denial difference. Cash-out refinancing is treated as suggestive because its bank-clustered p -value exceeds 0.05.

5.3 Economic Magnitudes

Table 6 translates the main interaction coefficients into changes in denial-rate differences over economically observed changes in provision pressure. A move from the 25th to the 75th percentile of provision pressure is associated with a 1.36 percentage-point larger Black-White denial difference for large banks overall, a 1.17 percentage-point difference in home-improvement lending, and a 2.20 percentage-point difference in refinancing. The corresponding estimate for other-purpose lending is 1.27 percentage points.

These magnitudes should not be interpreted as the total Black-White mortgage denial gap or as a causal effect of CECL. They translate the estimated conditional association into denial-rate units. Even changes of one to two percentage points can be economically meaningful in national mortgage markets because the estimates apply to millions of applications and to credit channels that affect household liquidity, borrowing costs, and property investment.

Table 6
Economic Magnitude Translations

Channel Label	Cov Type	Delta Type	Effect Percentage Points	Effect Formatted	Coef	Pvalue
Banks >= \$10B, all loan purposes	cluster_bank	p75_minus_p25	1.361	1.36 pp	2.900	0.002
Banks >= \$10B, all loan purposes	cluster_bank	p90_minus_p50	2.521	2.52 pp	2.900	0.002
Banks >= \$10B, all loan purposes	cluster_bank	one_sd	2.246	2.25 pp	2.900	0.002
Home improvement, banks >= \$10B	cluster_bank	p75_minus_p25	1.166	1.17 pp	2.796	< 0.001
Home improvement, banks >= \$10B	cluster_bank	p90_minus_p50	2.682	2.68 pp	2.796	< 0.001
Home improvement, banks >= \$10B	cluster_bank	one_sd	2.647	2.65 pp	2.796	< 0.001
Refinancing, banks >= \$10B	cluster_bank	p75_minus_p25	2.196	2.20 pp	4.109	0.009
Refinancing, banks >= \$10B	cluster_bank	p90_minus_p50	4.037	4.04 pp	4.109	0.009
Refinancing, banks >= \$10B	cluster_bank	one_sd	3.493	3.49 pp	4.109	0.009
Other purpose, banks >= \$10B	cluster_bank	p75_minus_p25	1.274	1.27 pp	2.845	< 0.001
Other purpose, banks >= \$10B	cluster_bank	p90_minus_p50	2.809	2.81 pp	2.845	< 0.001
Other purpose, banks >= \$10B	cluster_bank	one_sd	2.719	2.72 pp	2.845	< 0.001
Cash-out refinancing, banks >= \$10B	cluster_bank	p75_minus_p25	2.705	2.70 pp	5.742	0.082
Cash-out refinancing, banks >= \$10B	cluster_bank	p90_minus_p50	4.829	4.83 pp	5.742	0.082
Cash-out refinancing, banks >= \$10B	cluster_bank	one_sd	3.090	3.09 pp	5.742	0.082
<i>Note. This table translates selected coefficients into percentage-point changes in the Black-White denial difference for changes in provision pressure. Magnitudes help interpret the economic size of the estimates beyond statistical significance.</i>						

5.4 Underwriting Risk, DTI, and Within-Bank Comparisons

The observable-control tests show that the main Black interaction is not eliminated as borrower, loan, tract, and categorical controls are added. In the full-control specification, the large-bank coefficient is 2.900 ($p = 0.002$), the home-improvement coefficient is 2.796 ($p < 0.001$), and the refinancing coefficient is 4.109 ($p = 0.009$). The conventional-loan validation produces the same directional pattern: 2.627 for large banks overall, 2.797 for home improvement, and 3.306 for refinancing, all statistically significant. FHA results are weaker or distinct, which is consistent with the main evidence being concentrated in the more bank-controlled conventional channel, although the FHA subsamples differ substantially in size and composition.

Observable-risk partitions provide a second check. In large-bank nonpurchase lending, the Black interaction remains positive in both higher- and lower-LTV groups, both higher- and lower-income groups, higher- and lower-tract-income areas, and higher- and lower-minority-share tracts. Within-bank-year specifications also remain positive: the coefficient is 2.567 for the combined large-bank nonpurchase sample, 2.502 for home improvement, and 2.636 for refinancing, with $p < 0.001$ in each case. These tests reduce the concern that the result is driven only by broad bank composition or one observable borrower-risk segment. Table 7 reports the within-bank-year specifications.

Table 7
Within-Bank-Year Specifications

Channel	Black coefficient	p-value	Weighted applications	Bank clusters	WLS cells
Banks >= \$10B, home improvement/refinance/other purpose	2.567	< 0.001	12,616,904	145	112,503
Home improvement, banks >= \$10B	2.502	< 0.001	3,509,954	118	32,086
Refinancing, banks >= \$10B	2.636	< 0.001	5,761,486	144	45,425
<i>Note. This table reports within-bank-year-style specifications for the main large-bank nonpurchase channels. These models help address the concern that the evidence is driven solely by stable bank-level composition or broad year effects.</i>					

The DTI common-sample test directly addresses a central underwriting variable. In the all-application DTI-defined common sample for large banks, Model A, which omits DTI indicators, produces a Black interaction of 3.595 ($p = 0.001$). Model B adds DTI-category indicators on the exact same 426,363 cells and 18,195,322 weighted applications; the coefficient is 3.638 ($p < 0.001$). Thus the large-bank result does not depend on either omitting DTI or changing the estimation sample. Table 8 reports the paired common-sample estimates.

The same paired design yields a more informative attenuation pattern in home-improvement lending. On the same 74,400 cells representing 2,996,744 applications, the coefficient falls from 2.611 ($p = 0.018$) without DTI indicators to 2.008 ($p = 0.002$) with DTI indicators, a reduction of about 23 percent. DTI therefore explains part of the association in this channel, but it does not eliminate it. In refinancing, the coefficient rises from 4.468 ($p = 0.012$) to 6.510 ($p < 0.001$) on the same 89,806 cells representing 4,768,279 applications.

Restricting the analysis to applications with reported DTI produces the same conclusion and removes reliance on a missing-DTI category. In the reported-DTI common sample, the large-bank coefficient changes from 4.384 to 4.151 after adding DTI controls, a decline of about 5 percent; the home-improvement coefficient changes from 4.632 to 3.324, a decline of about 28 percent; and the refinancing coefficient changes from 4.476 to 4.136, a decline of about 8 percent. All three DTI-controlled coefficients remain statistically significant, with the refinancing p-value equal to 0.024 and the other two below 0.001.

The DTI evidence materially improves the underwriting-risk assessment because the paired models hold the cells, weights, and bank clusters fixed. The results therefore cannot be attributed to DTI creating finer cells that change which applications enter the regression. At the same time, DTI does not resolve the remaining credit-score limitation or provide access to internal lender risk models. We therefore interpret the DTI evidence as a strong observable-risk robustness result, not as proof that all underwriting differences have been removed.

Table 8
DTI Common-Sample Robustness

Sample	Channel	No DTI control β (p-value)	With DTI control β (p-value)	Change	WLS cells	Weighted applications	Bank clusters
All applications	Banks >= \$10B, all purposes	3.595 (0.001)	3.638 (< 0.001)	+ 1.2%	426,363	18,195,322	148
All applications	Home improvement, banks >= \$10B	2.611 (0.018)	2.008 (0.002)	-23.1%	74,400	2,996,744	120
All applications	Refinancing, banks >= \$10B	4.468 (0.012)	6.510 (< 0.001)	+ 45.7%	89,806	4,768,279	146
Reported DTI only	Banks >= \$10B, all purposes	4.384 (< 0.001)	4.151 (< 0.001)	-5.3%	402,760	15,693,287	147
Reported DTI only	Home improvement, banks >= \$10B	4.632 (< 0.001)	3.324 (< 0.001)	-28.2%	70,367	2,793,304	117
Reported DTI only	Refinancing, banks >= \$10B	4.476 (0.021)	4.136 (0.024)	-7.6%	83,158	3,887,511	145

Note. Each row compares two bank-clustered weighted least squares models estimated on exactly the same DTI-defined cells, application weights, and bank clusters. The reported coefficient is Black x Post-CECL x Provision Pressure. The no-DTI model omits DTI-category indicators; the with-DTI model adds them. DTI categories are < 20%, 20%-<30%, 30%-<36%, 36%-<50%, 50%-60%, and > 60%; the all-applications sample also retains missing/not-applicable/exempt DTI as a separate category. The reported-DTI-only sample excludes those observations. Change is the percentage change in the Black coefficient after adding DTI controls. Identical cell and application counts within each row ensure that the comparison isolates DTI adjustment rather than sample-composition changes.

5.5 Additional Robustness and Alternative Explanations

The principal results are not sensitive to reasonable alternative inference structures. The large-bank, home-improvement, and refinancing coefficients remain statistically significant with bank-year clustering, and state-clustered results point in the same direction. Provision-pressure outlier tests also remain supportive. Alternative winsorization thresholds, exclusion of negative provision observations, and exclusion of the highest provision observations leave the large-bank, home-improvement, and refinancing coefficients positive and statistically significant in the principal variants.

Lender-concentration tests require a more qualified interpretation. Excluding the five or ten largest lenders preserves positive and significant results in the main large-bank, home-improvement, and refinancing tests. Removing each of the ten largest banks one at a time also leaves the large-bank coefficient between 2.727 and 4.829, the home-improvement coefficient between 2.657 and 7.922, and the refinancing coefficient between 3.889 and 6.870. When much larger portions of the large-bank lending market are removed at once, however, some coefficients lose significance or change sign. This pattern indicates that the evidence is not driven by one dominant lender, but it is genuinely concentrated in the large-bank sector rather than being invariant to removing much of that sector.

The placebo tests do not show the positive pre-period relation that would be expected if later provision pressure merely identified banks with persistently larger Black-White denial differences. The large-bank placebo coefficient is -1.610, and the home-improvement, refinancing, other-purpose, and cash-out placebo coefficients are also nonpositive. Some are statistically different from zero in the negative direction, but none provides the positive pre-period pattern that would mirror the main post-CECL evidence.

Alternative accounting-pressure variables provide a final diagnostic. Allowance ratios and net charge-offs produce positive associations in several of the main channels, while past-due measures are substantially less consistent. Provision pressure remains the primary construct because it is the current income-statement recognition measure tied most directly to the expected-loss mechanism and produces the most stable pattern across the focal specifications. The related variables nevertheless show that the result sits within a broader credit-loss environment rather than being mechanically unique to a single accounting ratio.

Taken together, the results support a deliberately bounded conclusion. Post-CECL provision pressure is associated with mortgage-credit allocation in the settings where the proposed transmission mechanism should be strongest: large banks and selected nonpurchase mortgage channels. The association is stronger for Black applicants relative to White applicants, while Hispanic and Asian results are less consistent. The DTI, within-bank-year, loan-type, risk-partition, placebo, inference, and outlier tests reduce several important alternative explanations, but they do not convert the design into a causal test of CECL or intentional discrimination.

Additional robustness and diagnostic results are reported in Online Resource 1 (Tables S1-S10).

6. Discussion and Implications

6.1 What the Results Tell Us About Mortgage-Credit Allocation

The evidence extends the consequences of expected-loss provisioning beyond aggregate credit supply and loan pricing. The central result is that bank-year provision pressure is associated with how mortgage credit is allocated across applications after CECL. The association is not uniform. It is concentrated in large banks and in selected nonpurchase mortgage channels, and the distributional association is strongest for Black applicants relative to White applicants. This pattern is more consistent with a credit-allocation channel than with a simple claim that high-provision banks deny more applications across the board.

The large scale of the linked HMDA and Call Report data is useful because the allocation question cannot be answered from either data source alone. HMDA reveals applicant-level outcomes but not the bank-year accounting environment, while Call Reports reveal provision pressure but not which mortgage applications are approved or denied. Linking the two makes it possible to observe whether accounting pressure and borrower-level credit outcomes move together within specific institutional and product settings.

The underwriting-risk tests sharpen that interpretation. Adding DTI does not eliminate the focal Black-White association in any of the three principal common-sample tests. The coefficient attenuates in home-improvement lending, which indicates that DTI explains part of the observed difference in that channel, but the association remains positive and statistically significant. In the large-bank and refinancing samples, the DTI-adjusted results remain at least as strong as the corresponding same-sample estimates without DTI controls. The evidence therefore does not support an interpretation in which observable debt burden alone explains the main findings.

At the same time, the results do not establish that provision pressure is an exogenous treatment or that otherwise identical applicants receive different decisions because of CECL. The empirical contribution is narrower: expected-loss provision pressure is systematically associated with mortgage-credit allocation in the settings where the proposed transmission mechanism should be most visible, and that association survives several direct tests of observable underwriting risk and lender composition.

6.2 Why Large Banks and Nonpurchase Financial Services Matter

The concentration of the results among banks with at least \$10 billion in assets is important for evaluating the proposed mechanism. Large banks are more likely to rely on centralized underwriting, formal risk models, portfolio-level controls, and standardized approval processes. Those organizational features create a setting in which a bank-year accounting signal can be incorporated into lending rules at scale. The evidence is therefore consistent with the idea that provision pressure matters not only as a reported accounting outcome but also as an input into broader credit-risk management systems. These tests identify a mechanism-consistent concentration of the association; they do not directly observe changes in underwriting rules or internal risk models.

This interpretation does not imply that bank size itself causes the result. Large institutions differ from smaller banks in technology, product mix, geography, regulatory scrutiny, borrower composition, and business strategy. The size partition is useful because it identifies the organizational setting in which the hypothesized transmission channel is most plausible. The top-bank and leave-one-out diagnostics reinforce that distinction: the result is not attributable to one dominant lender, but it does weaken when a substantial portion of the large-bank sector is removed. The evidence is therefore best described as large-bank concentrated rather than universal across institutions.

The product concentration is equally informative. The strongest evidence appears in home-improvement and refinancing rather than solely in home-purchase lending. These products represent continuing financial services associated with homeownership. Refinancing can reduce borrowing costs, alter loan terms, and help households manage liquidity and payment risk. Home-improvement credit can support property maintenance, investment, and collateral preservation. A provision-pressure association in these channels therefore has implications for the financial benefits households can obtain after acquiring a home, not only for whether they can purchase one initially.

This product pattern also makes the paper more clearly a financial-services study. The findings concern how bank risk systems allocate access to household credit products under an expected-loss regime. They therefore connect accounting measurement to the provision of financial services through the institutional processes by which banks translate portfolio-level risk signals into application-level decisions.

6.3 Distributional Consequences and Claim Boundaries

The distributional evidence is strongest for Black applicants relative to White applicants. A useful interpretation is that a common tightening in risk tolerances can have unequal consequences even without an explicitly race-specific rule. If borrower groups differ in their distribution around underwriting margins because of differences in wealth, income stability, collateral, home equity, neighborhood conditions, or prior access to credit, then a common change in approval thresholds can affect those groups differently.

The DTI results make this interpretation more credible but also more precise. DTI is an important underwriting input, and the analysis controls for it directly. On identical DTI-defined samples, the focal Black interaction remains positive and statistically significant after DTI controls are added. In home-improvement lending the attenuation is meaningful, which shows that observable debt burden contributes to the pattern. The remaining association, however, indicates that DTI alone does not account for the result. Credit score and banks' internal underwriting information remain unobserved, so the analysis cannot claim that all borrower-risk differences have been removed.

The Hispanic and Asian estimates are less consistent than the Black-White results. That heterogeneity is informative and should not be smoothed into a broad minority-versus-White conclusion. Borrower groups differ across markets, lenders, loan purposes, income distributions, collateral conditions, and product choices. The evidence therefore supports a specific distributional result for the Black-White comparison in the settings where the provision-pressure association is strongest, rather than a claim of uniform effects across all racial and ethnic groups.

These boundaries are central to the paper. The analysis does not identify intentional discrimination, does not observe loan officers' motives, and does not estimate a clean causal effect of CECL. Provision pressure itself may reflect underlying economic conditions that also affect lending. The contribution is to show that an accounting measure tied to expected credit losses is associated with the distribution of mortgage-credit outcomes after CECL, even after extensive observable-risk, timing, lender-composition, and within-bank diagnostics.

6.4 Implications for Financial Institutions, Regulators, and Standard Setters

The results have implications for how the real effects of expected-loss accounting are evaluated. CECL is designed to improve the timeliness of credit-loss recognition, not to prescribe which households receive mortgage credit. Yet when expected-loss measures enter

capital planning, portfolio controls, pricing systems, and underwriting processes, their effects can extend beyond reported financial statements. A reporting rule can therefore influence financial services indirectly through the decision systems that use the resulting accounting information.

For financial institutions, the evidence suggests that model governance should consider distributional outcomes when portfolio-level risk signals are incorporated into standardized credit systems. The relevant issue is not whether banks should ignore expected losses. Rather, institutions may benefit from evaluating whether broad risk adjustments have different effects across products or applicant populations, particularly when approval thresholds are changed at scale.

For bank regulators, the findings suggest that the consequences of provisioning regimes may not be fully visible in aggregate loan growth. Banks can adjust credit allocation across applications and products even when aggregate lending changes are modest. Monitoring how credit adjustments are distributed across refinancing, home improvement, purchase credit, and borrower groups can therefore complement traditional measures of total credit supply and bank resilience.

For accounting standard setters, the evidence illustrates a broader real-effects consideration. Expected-loss standards can improve recognition and measurement while also becoming inputs into organizational decision systems. The paper does not imply that these indirect effects outweigh the benefits of timely loss recognition or that CECL should be evaluated primarily through distributional outcomes. It does suggest, however, that the consequences of accounting standards may extend to the allocation of household financial services when accounting measures are embedded in bank risk-management processes.

Taken together, the findings point to a broader view of provisioning systems. Expected-loss accounting is part of a financial institution's information and control architecture. Once expected-loss measures are incorporated into that architecture, their consequences may appear not only in reserves, earnings, and aggregate lending, but also in which credit channels expand or contract and which applications are most affected.

7. Conclusion

This paper asks whether expected-loss provision pressure is associated with how banks allocate mortgage credit after CECL. Using linked national HMDA application data, the Philadelphia Fed HMDA Lender File, and FDIC/FFIEC Call Report data, we connect applicant-level mortgage outcomes to the bank-year accounting environment in which lending decisions are made. The central finding is that post-CECL provision pressure is associated with mortgage-credit allocation, and that the relation is concentrated in the institutional and product settings where formal bank risk systems are most likely to transmit expected-loss signals.

The evidence is strongest among banks with at least \$10 billion in assets and in nonpurchase mortgage channels, especially home improvement and refinancing. Within those settings, higher post-CECL provision pressure is associated with a larger Black-White denial difference. The result survives alternative inference approaches, lender and outlier diagnostics, placebo timing tests, observable-risk partitions, loan-type validation, within-bank-year comparisons, and the DTI common-sample tests. Importantly, the DTI analysis shows that observable debt burden explains part of the home-improvement association but does not eliminate the principal result.

These findings should be interpreted within the limits of the research design. Provision pressure is not randomly assigned, CECL implementation overlaps with important macroeconomic changes, and public HMDA does not provide borrower credit scores or banks' internal underwriting information. The analysis therefore does not identify a clean causal effect of CECL and does not establish intentional discrimination. What the data establish is narrower: an accounting measure tied to expected credit losses is systematically associated with the distribution of mortgage-credit outcomes after CECL, even after substantial adjustment for observable underwriting risk and lender composition.

The broader contribution is to show that expected-loss accounting may matter not only for when banks recognize losses, how much credit they supply, or how they price loans, but also for how credit is allocated across applications and financial-service channels. When provisioning measures enter centralized underwriting, portfolio controls, and other bank decision systems, the consequences of accounting can extend from financial statements to household credit access. That link between expected-loss provisioning and mortgage-credit allocation is the main fact this study adds to the banking and financial-services literature.

Declarations

Author Contribution

David H. Olsen performed the data retrieval and empirical analysis and prepared the results summaries. Cindy Greenman conducted the literature review and drafted the introduction, literature review, and theory development. Both authors jointly developed the theory, revised the manuscript, and approved the final version.

Data Availability

The data used in this study are publicly available. Home Mortgage Disclosure Act (HMDA) mortgage application data are available from the Consumer Financial Protection Bureau and the Federal Financial Institutions Examination Council. The HMDA Lender File is publicly available from the Federal Reserve Bank of Philadelphia. FDIC/FFIEC Call Report data are publicly available from the Federal Deposit Insurance Corporation and the Federal Financial Institutions Examination Council. Replication code and documentation are available from the authors upon reasonable request. Processed analysis files can be made available subject to file-size limitations and the terms of use of the underlying public data sources.

References

1. Aaronson D, Hartley D, Mazumder B (2021) The effects of the 1930s HOLC redlining maps. *Am Economic Journal: Economic Policy* 13(4):355–392. <https://doi.org/10.1257/pol.20190414>
2. Ambrose BW, Conklin JN, Lopez LA (2021) Does borrower and broker race affect the cost of mortgage credit? *Rev Financial Stud* 34(2):790–826. <https://doi.org/10.1093/rfs/hhaa087>
3. Avery RB, Brevoort KP, Canner GB (2007) Opportunities and issues in using HMDA data. *J Real Estate Res* 29(4):351–380. <https://www.federalreserve.gov/pubs/bulletin/2007/pdf/hmda.pdf>
4. Avenancio-León CF, Howard T (2022) The assessment gap: Racial inequalities in property taxation. *Q J Econ* 137(3):1383–1434. <https://doi.org/10.1093/qje/qjac009>
5. Baker AC, Larcker DF, Wang CCY (2022) How much should we trust staggered difference-in-differences estimates? *J Financ Econ* 144(2):370–395. <https://doi.org/10.1016/j.jfineco.2022.01.004>
6. Bank for International Settlements, Basel Committee on Banking Supervision (2015) Guidance on credit risk and accounting for expected credit losses. <https://www.bis.org/bcbs/publ/d350.htm>
7. Bartlett R, Morse A, Stanton R, Wallace N (2022) Consumer-lending discrimination in the FinTech era. *J Financ Econ* 143(1):30–56. <https://doi.org/10.1016/j.jfineco.2021.05.047>
8. Bayer P, Ferreira F, Ross SL (2018) What drives racial and ethnic differences in high-cost mortgages? The role of high-risk lenders. *Rev Financial Stud* 31(1):175–205. <https://doi.org/10.1093/rfs/hhx035>
9. Beatty A, Liao S (2011) Do delays in expected loss recognition affect banks' willingness to lend? *J Account Econ* 52(1):1–20. <https://doi.org/10.1016/j.jacceco.2011.02.002>
10. Beatty A, Liao S (2014) Financial accounting in the banking industry: A review of the empirical literature. *J Account Econ* 58(2–3):339–383. <https://doi.org/10.1016/j.jacceco.2014.08.009>
11. Berger AN, Udell GF (2002) Small business credit availability and relationship lending: The importance of bank organisational structure. *Econ J* 112(477):F32–F53. <https://doi.org/10.1111/1468-0297.00682>
12. Berger AN, Miller NH, Petersen MA, Rajan RG, Stein JC (2005) Does function follow organizational form? Evidence from the lending practices of large and small banks. *J Financ Econ* 76(2):237–269. <https://doi.org/10.1016/j.jfineco.2004.06.003>
13. Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *Q J Econ* 119(1):249–275. <https://doi.org/10.1162/003355304772839588>
14. Bhutta N, Hizmo A, Ringo D (2025) How much does racial bias affect mortgage lending? Evidence from human and algorithmic credit decisions. *J Finance* 80(3):1463–1496. <https://doi.org/10.1111/jofi.13444>
15. Bikker JA, Metzmakers PAJ (2005) Bank provisioning behaviour and procyclicality. *J Int Financ Mark Inst Money* 15(2):141–157. <https://doi.org/10.1016/j.intfin.2004.03.004>
16. Borusyak K, Jaravel X, Spiess J (2024) Revisiting event-study designs: Robust and efficient estimation. *Rev Econ Stud* 91(6):3253–3285. <https://doi.org/10.1093/restud/rdae007>
17. Bushman RM, Williams CD (2012) Accounting discretion, loan loss provisioning, and discipline of banks' risk-taking. *J Account Econ* 54(1):1–18. <https://doi.org/10.1016/j.jacceco.2012.04.002>

18. Callaway B, Sant'Anna PHC (2021) Difference-in-differences with multiple time periods. *J Econ* 225(2):200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
19. Cameron AC, Miller DL (2015) A practitioner's guide to cluster-robust inference. *J Hum Resour* 50(2):317–372. <https://doi.org/10.3368/jhr.50.2.317>
20. Cameron AC, Gelbach JB, Miller DL (2011) Robust inference with multiway clustering. *J Bus Economic Stat* 29(2):238–249. <https://doi.org/10.1198/jbes.2010.07136>
21. Charles KK, Hurst E (2002) The transition to home ownership and the Black–White wealth gap. *Rev Econ Stat* 84(2):281–297. <https://doi.org/10.1162/003465302317411532>
22. Chen C, Huang D (2026) A tale of two banks: When credit loss models meet economic crises. *J Acc Res Adv* online publication. <https://doi.org/10.1111/1475-679X.70030>
23. Chetty R, Hendren N, Jones MR, Porter SR (2020) Race and economic opportunity in the United States: An intergenerational perspective. *Q J Econ* 135(2):711–783. <https://doi.org/10.1093/qje/qjz042>
24. Consumer Financial Protection Bureau (2024) Home Mortgage Disclosure Act (HMDA) data. <https://www.consumerfinance.gov/data-research/hmda/>
25. de Chaisemartin C, D'Haultfoeulle X (2020) Two-way fixed effects estimators with heterogeneous treatment effects. *Am Econ Rev* 110(9):2964–2996. <https://doi.org/10.1257/aer.20181169>
26. Derenoncourt E, Kim CH, Kuhn M, Schularick M (2024) Wealth of two nations: The U.S. racial wealth gap, 1860–2020. *Q J Econ* 139(2):693–750. <https://doi.org/10.1093/qje/qjad044>
27. Federal Deposit Insurance Corporation (2024) BankFind Suite and financial data downloads. <https://banks.data.fdic.gov/>
28. Federal Financial Institutions Examination Council (2024a) Central Data Repository's Public Data Distribution: Call Report data. <https://cdr.ffiec.gov/public/>
29. Federal Financial Institutions Examination Council (2024b) HMDA data publication platform. <https://ffiec.cfbp.gov/data-publication/>
30. Federal Reserve Bank of Philadelphia (2025) Home Mortgage Disclosure Act (HMDA) Lender File. <https://www.philadelphiafed.org/surveys-and-data/consumer-finance-data/home-mortgage-disclosure-act-lender-file>
31. Financial Accounting Standards Board (2016) Accounting Standards Update No. 2016-13: Financial Instruments—Credit Losses (Topic 326). <https://www.fasb.org/page/PageContent?pagelId=/projects/recentlycompleted/accounting-for-financial-instrumentscredit-losses.html>
32. Fonseca AR, González F (2008) Cross-country determinants of bank income smoothing by managing loan-loss provisions. *J Banking Finance* 32(2):217–228. <https://doi.org/10.1016/j.jbankfin.2007.02.012>
33. Fuster A, Goldsmith-Pinkham P, Ramadorai T, Walther A (2022) Predictably unequal? The effects of machine learning on credit markets. *J Finance* 77(1):5–47. <https://doi.org/10.1111/jofi.13090>
34. Gardner J (2022) Two-stage differences in differences. <https://doi.org/10.48550/arXiv.2207.05943>. arXiv:2207.05943
35. Ghent AC, Hernández-Murillo R, Owyang MT (2014) Differences in subprime loan pricing across races and neighborhoods. *Reg Sci Urban Econ* 48:199–215. <https://doi.org/10.1016/j.regsciurbeco.2014.07.007>
36. Goodman-Bacon A (2021) Difference-in-differences with variation in treatment timing. *J Econ* 225(2):254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>
37. Granja J, Nagel F (2025) Current expected credit losses and consumer loans. *J Account Econ* 80(1). <https://doi.org/10.1016/j.jacceco.2025.101781>. Article 101781
38. Howell ST, Kuchler T, Snitkof D, Stroebel J, Wong J (2024) Lender automation and racial disparities in credit access. *J Finance* 79(2):1457–1512. <https://doi.org/10.1111/jofi.13303>
39. Kanagaretnam K, Lobo GJ, Mathieu R (2003) Managerial incentives for income smoothing through bank loan loss provisions. *Rev Quant Financ Acc* 20(1):63–80. <https://doi.org/10.1023/A:1022187622780>
40. Kanodia C, Sapra H (2016) A real effects perspective to accounting measurement and disclosure: Implications and insights for future research. *J Accounting Res* 54(2):623–676. <https://doi.org/10.1111/1475-679X.12109>
41. Kermani A, Wong F (2026) Racial disparities in housing returns. *Am Econ Rev* 116(1):287–331. <https://doi.org/10.1257/aer.20240327>
42. Kim S (2022) Delays in banks' loan loss provisioning and economic downturns: Evidence from the U.S. housing market. *J Accounting Res* 60(3):711–754. <https://doi.org/10.1111/1475-679X.12415>
43. Laeven L, Majnoni G (2003) Loan loss provisioning and economic slowdowns: Too much, too late? *J Financ Intermediation* 12(2):178–197. [https://doi.org/10.1016/S1042-9573\(03\)00016-0](https://doi.org/10.1016/S1042-9573(03)00016-0)

44. Lewis-Faupel S, Tenev N (2024) Racial and ethnic disparities in mortgage lending: New evidence from expanded HMDA data. <https://doi.org/10.48550/arXiv.2405.00895>. arXiv:2405.00895
45. Liberti JM, Petersen MA (2019) Information: Hard and soft. *Rev Corp Finance Stud* 8(1):1–41. <https://doi.org/10.1093/rcfs/cfy009>
46. Liu C-C, Ryan SG (2006) Income smoothing over the business cycle: Changes in banks' coordinated management of provisions for loan losses and loan charge-offs from the pre-1990 bust to the 1990s boom. *Acc Rev* 81(2):421–441. <https://doi.org/10.2308/accr.2006.81.2.421>
47. López-Espinosa G, Ormazabal G, Sakasai Y (2021) Switching from incurred to expected loan loss provisioning: Early evidence. *J Accounting Res* 59(3):757–804. <https://doi.org/10.1111/1475-679X.12354>
48. Mahieux L, Sapra H, Zhang G (2023) CECL: Timely loan loss provisioning and bank regulation. *J Accounting Res* 61(1):3–46. <https://doi.org/10.1111/1475-679X.12463>
49. Munnell AH, Tootell GMB, Browne LE, McEneaney J (1996) Mortgage lending in Boston: Interpreting HMDA data. *Am Econ Rev* 86(1):25–53. <https://www.jstor.org/stable/2118269>
50. Petersen MA, Rajan RG (2002) Does distance still matter? The information revolution in small business lending. *J Finance* 57(6):2533–2570. <https://doi.org/10.1111/1540-6261.00505>
51. Ross SL, Yinger J (2002) The color of credit: Mortgage discrimination, research methodology, and fair-lending enforcement. MIT Press. <https://mitpress.mit.edu/9780262681264/the-color-of-credit/>
52. Roychowdhury S, Shroff N, Verdi RS (2019) The effects of financial reporting and disclosure on corporate investment: A review. *J Account Econ* 68(2–3). Article 101246. <https://doi.org/10.1016/j.jacceco.2019.101246>
53. Sirignano J, Sadhwani A, Giesecke K (2018) Deep learning for mortgage risk. <https://doi.org/10.48550/arXiv.1607.02470>. arXiv:1607.02470
54. Stein JC (2002) Information production and capital allocation: Decentralized versus hierarchical firms. *J Finance* 57(5):1891–1921. <https://doi.org/10.1111/1540-6261.00483>
55. Sun L, Abraham S (2021) Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J Econ* 225(2):175–199. <https://doi.org/10.1016/j.jeconom.2020.09.006>
56. U.S. Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and National Credit Union Administration (2016) Joint statement on the new accounting standard on financial instruments—credit losses. <https://www.fdic.gov/news/financial-institution-letters/2016/fil16039.html>
57. Yang H-C (2025) (Alex). How does loan loss accounting influence bank lending? Evidence from the Current Expected Credit Loss (CECL) model. *The Accounting Review*, 100(1), 465–490. <https://doi.org/10.2308/TAR-2022-0511>
58. Zhang J (2008) The contracting benefits of accounting conservatism to lenders and borrowers. *J Account Econ* 45(1):27–54. <https://doi.org/10.1016/j.jacceco.2007.06.002>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ESM1.pdf](#)