

---

**| RESEARCH ARTICLE**

**From Utility Bills to Retrofit Finance: An AI Framework for Energy-Burden-Aware Underwriting in Residential Decarbonization**

**Shah Farhan Rabbani**

North South University

Email: [shah.rabbani@northsouth.edu](mailto:shah.rabbani@northsouth.edu)

ORCID: <https://orcid.org/0009-0001-8434-223X>

**Yusuf Oli Rahat**

Brac University

Email: [yusuf.rahat@gmail.com](mailto:yusuf.rahat@gmail.com)

ORCID: <https://orcid.org/0009-0001-8834-7759>

**Md Kamrul Islam**

University of New Haven, Business Analytics

Email: [misla22@unh.newhaven.edu](mailto:misla22@unh.newhaven.edu)

ORCID: <https://orcid.org/0009-0001-8906-630X>

**Corresponding Author:** YUVACHANDRA MARASANI, **E-mail:** [yuvachandramarasani@gmail.com](mailto:yuvachandramarasani@gmail.com)

---

**| ABSTRACT**

This study develops a literature-driven conceptual framework for energy-burden-aware underwriting within the context of U.S. residential decarbonization finance. It addresses a critical and growing disconnect between the policy-driven urgency to decarbonize residential buildings and the financial constraints faced by households most in need of retrofit investments. Notably, those households that stand to benefit the most from energy-efficiency improvements are often underserved by traditional underwriting models. Conventional lending frameworks prioritize stable income streams, established credit histories, and property-based collateral signals, whereas energy-burdened households—characterized by high utility costs, inefficient housing conditions, and constrained liquidity—exhibit affordability dynamics that are poorly captured by these models but are highly relevant for socially beneficial investment. Drawing upon interdisciplinary literature spanning energy burden, building decarbonization, alternative data underwriting, explainable artificial intelligence, and algorithmic accountability, this paper proposes an AI-enabled underwriting framework that integrates utility-bill histories, housing attributes, retrofit cost-benefit dynamics, and consumer protection considerations. Importantly, the framework emphasizes transparency and interpretability, avoiding reliance on opaque risk-scoring mechanisms. Methodologically, the study employs a structured integrative review of academic and policy-oriented sources, synthesizing insights into a multimodal framework encompassing data governance, feature engineering, model development, explainability outputs, fairness diagnostics, and financing pathway design. The contributions of this research are threefold. First, it reconceptualizes utility payment data from a signal of financial distress into a meaningful input for transition finance. Second, it establishes a direct linkage between household energy burden and underwriting design principles. Third, it outlines governance and accountability conditions under which AI-assisted retrofit financing can advance decarbonization objectives without reinforcing systemic exclusion. Given the absence of original empirical data, the findings are presented as a conceptual architecture, evaluation framework, and forward-looking research agenda rather than as empirical validation. The proposed framework is intended to guide future implementation and testing by financial institutions, utilities, public green banks, and housing policy stakeholders. paper develops a literature-based conceptual framework for energy-burden-aware underwriting in U.S. residential decarbonization finance. The study addresses a widening disconnect between the climate-policy imperative to decarbonize homes and the financial reality that many households most likely to benefit from retrofit investment are also the least well served by conventional underwriting. Existing lending models

privilege stable income, established credit files, and property value signals, whereas the households facing high utility costs, inefficient housing, and limited liquidity often present precisely the affordability patterns that make retrofit investment socially valuable. Drawing on scholarship on energy burden, building decarbonization, alternative-data underwriting, explainable artificial intelligence, and algorithmic accountability, the paper proposes an AI framework that uses utility-bill histories, housing characteristics, retrofit economics, and consumer-protection rules to assess payment feasibility without reducing underwriting to opaque risk scoring. Methodologically, the paper uses a structured integrative review academic and policy sources and synthesizes them into a multimodal framework comprising data governance, feature design, model development, explanation outputs, fairness diagnostics, and financing-channel selection. The paper's contribution is threefold: it reframes utility bills from a distress signal into a transition-finance input; it links household energy burden to underwriting design; and it specifies governance conditions under which AI-assisted retrofit finance may advance decarbonization without reproducing exclusion. Because no original dataset is analyzed, the results are presented as a conceptual model, evaluation logic, and research agenda rather than empirical claims. The framework is intended to support future testing by lenders, utilities, public green banks, and housing-policy practitioners.

## | KEYWORDS

Fintech; artificial intelligence; residential decarbonization; retrofit finance; energy burden; alternative data; algorithmic fairness; explainable AI; underwriting; climate finance

## | ARTICLE INFORMATION

**ACCEPTED:** 01 November 2024

**PUBLISHED:** 25 December 2024

**DOI:** 10.32996/fcsai.2024.3.2.10

## 1. Introduction

Residential decarbonization has moved from a peripheral sustainability concern to a central question in U.S. climate policy, housing policy, and consumer finance. Buildings are a major source of U.S. greenhouse-gas emissions and household energy costs, and the federal 2024 blueprint for the buildings sector frames retrofit and electrification as indispensable to national decarbonization goals (Langevin et al., 2024; U.S. Department of Energy, 2024). Yet the transition from climate ambition to household action remains constrained by a stubborn financing problem. Deep energy retrofits, weatherization, envelope improvements, and electrification upgrades frequently require capital at the point of decision, while the economic benefits often accrue over time through lower operating costs, improved comfort, and reduced exposure to volatile fuel prices. This temporal mismatch is particularly acute for households already experiencing high energy burdens, limited liquidity, or thin credit files.

The financial barrier is not simply that retrofits are expensive. It is that conventional underwriting methods are poorly designed to evaluate the kind of future affordability gains that retrofits are meant to create. Standard consumer underwriting tends to rely on backward-looking credit histories, debt-to-income ratios, and collateral proxies. These measures are familiar, scalable, and regulatorily legible, but they are not well suited to assessing retrofit-enabled cash-flow improvement in homes with inefficient shells, outdated appliances, or dependence on costly fossil-based heating and cooling. As a result, households with the strongest social case for decarbonization finance may appear unattractive in the very lending systems that would need to support their transition.

This tension is intensified by the structure of household energy insecurity in the United States. Energy burden is typically defined as the share of household income spent on home energy, and a large literature now shows that its distribution is shaped by more than income alone. Housing condition, tenure, race, geography, and community vulnerability influence both exposure to high bills and the capacity to respond through efficiency investments (Brown, Kale & Anthony, 2023; Shen et al., 2023; Graff, 2024). Recent U.S. policy briefs likewise emphasize that low-income households continue to face disproportionately high burdens and frequently exceed the conventional 6% affordability threshold (ACEEE, 2024). In this context, utility bills are not just evidence of expense. They are an informative record of thermal inefficiency, seasonal strain, coping pressure, and possible retrofit opportunity.

Fintech and artificial intelligence create a plausible pathway for translating that record into more nuanced financial decisions. The broader fintech literature has shown that machine learning and alternative data can improve decision speed, extend credit evaluation beyond traditional bureau data, and potentially widen access for credit-invisible or thin-file applicants (Tierno, 2024; Schneider, 2024). At the same time, work in fair lending and responsible AI warns that algorithmic tools can reproduce historical exclusion if they are built from biased data, evaluated only on predictive performance, or deployed without meaningful explanation and governance (Kozodoi, Jacob & Lessmann, 2022; Garcia, Pinto Garcia & Rigobon, 2024; Ratzan & Rahman, 2024).

In other words, AI can improve underwriting, but not automatically and not neutrally. In climate-linked household finance, the normative stakes are even higher because lending decisions shape who can participate in the energy transition and who remains locked into inefficient housing and elevated utility costs.

A further challenge is that the literatures most relevant to this problem rarely speak directly to one another. Studies of residential decarbonization focus on retrofit pathways, energy savings, and policy design (Berrill et al., 2022; York et al., 2022; Supple, 2024). Research on energy burden examines vulnerability, health, insecurity, and affordability (Brown, Kale & Anthony, 2023; Hatch & Graff, 2024; Graff, 2024). Fintech scholarship considers alternative data, credit modeling, explainability, and governance (Černevičienė & Kabašinskas, 2024; Schneider, 2024). Yet there is comparatively little work that asks how AI-enabled underwriting should be redesigned when the financing objective is not generic consumer credit, but equitable residential decarbonization under conditions of high energy burden. This gap matters because decarbonization finance cannot be treated as a simple extension of ordinary lending. It involves projected energy savings, public incentives, housing quality, utility-payment behavior, and justice concerns that exceed the logic of conventional scorecards.

This paper argues that retrofit finance should be assessed through payment feasibility rather than through conventional risk proxies alone, and that an explainable, fairness-aware AI architecture can operationalize that shift. The paper therefore develops a conceptual framework for energy-burden-aware underwriting that integrates utility bills, housing stock information, retrofit economics, public incentive structures, and responsible AI governance. The contribution is not an empirical test with fabricated results; no original dataset is used. Instead, the paper offers a structured review and conceptual synthesis suitable for a December 2024 publication context. It reframes utility bills as transition-finance inputs, identifies the research gap at the intersection of fintech and residential decarbonization, and proposes a model architecture that can guide future empirical testing by lenders, utilities, public green banks, and policymakers. The remainder of the paper reviews the literature, states the research gap and objectives, outlines the review-based methodology, presents the proposed framework, and discusses its theoretical and practical implications.

## **2. Background / Literature Review**

### *Residential decarbonization and the financing problem*

Residential decarbonization depends not only on technology diffusion but on financing architectures capable of reaching heterogeneous households. Berrill et al. (2022) demonstrate that low-emission pathways for the U.S. residential sector require extensive renovation of existing homes, rapid electrification of end uses, and cleaner electricity supply. That finding is important because it shifts attention from new construction alone to the much harder problem of upgrading an aging and unevenly efficient housing stock. The federal 2024 buildings blueprint follows a similar logic, positioning building decarbonization as a long-horizon national strategy with implications for consumer energy costs, public health, and climate mitigation (Langevin et al., 2024; U.S. Department of Energy, 2024). However, transition pathways at the sector level do not tell households how to finance retrofits, nor do they solve the problem that the homes most in need of improvement are often occupied by financially constrained residents.

Research and practice reports increasingly emphasize that retrofit uptake is shaped by capital constraints, fragmented incentives, and the inability of many program designs to serve low- and moderate-income households at scale. York et al. (2022) note that affordable housing decarbonization requires combining efficiency and electrification rather than treating them as separate policy silos, but they also show that the sector faces barriers related to upfront cost, building ownership structure, split incentives, and administrative complexity. Bowen (2023) similarly documents the need for financing mechanisms that address upfront costs in affordable housing contexts. NREL work by Charan et al. (2024) pushes the conversation further by arguing that inclusive utility program design can reduce affordability barriers when decarbonization measures are structured around customer realities rather than idealized borrowers. The key lesson from this literature is that financing cannot be an afterthought to decarbonization; it is one of the central determinants of whether retrofit policy reaches households equitably.

At the household level, the retrofit decision is usually evaluated through a difficult mix of present sacrifice and uncertain future gain. Even when projected bill savings are positive over time, households may still confront credit constraints, distrust of contractors, uncertainty about realized savings, and limited ability to navigate rebates or stacked incentives. Supple (2024) shows that equitable building-efficiency and electrification strategies require attention to local housing characteristics, renter-owner differences, and community-centered program design. The implication for underwriting is straightforward but underappreciated: retrofit loans should not be evaluated as generic unsecured consumption credit. They are transition-finance instruments whose repayment feasibility depends in part on future operating-cost reductions, equipment performance, and public support structures. That makes them analytically closer to project-based affordability improvement than to ordinary discretionary borrowing.

*Energy burden, vulnerability, and the justice dimension*

The literature on energy burden and energy insecurity provides the social context that retrofit finance must confront. Brown, Kale and Anthony (2023) argue that annualized burden measures can obscure substantial seasonal and situational hardship, especially in the U.S. Southeast, where representative households in highly burdened tracts may spend an exceptionally large share of income on energy during peak periods. Shen et al. (2023) further demonstrate that community vulnerability, not income alone, is a key determinant of diverse energy burdens across U.S. counties. Their analysis underlines the spatial heterogeneity of burden and suggests that vulnerability is produced through layered socioeconomic and environmental conditions rather than a single financial variable. For underwriting, this means that a household's payment strain is partly embedded in place-based and housing-based characteristics that standard credit files often miss.

Graff (2024) extends this discussion by showing that energy insecurity in the United States remains widespread and that assistance programs, while essential, do not fully resolve structural affordability problems. Energy insecurity is not reducible to arrears or disconnection notices; it also includes harmful coping strategies, unmet thermal comfort, deferred spending on essentials, and persistent financial stress. Hatch and Graff (2024) add a housing-stability perspective, showing that neighborhood energy burden is associated with eviction filing rates. This matters because it links utility-cost stress to broader forms of precarity, making clear that energy affordability is not merely an energy-policy issue but part of a wider system of housing insecurity. In practical terms, an underwriting model that ignores utility stress may misread financial vulnerability, while one that recognizes it only as negative risk may foreclose precisely the investments that could reduce the burden.

The justice-oriented building literature reinforces this point. Flores et al. (2024), in a study of a low-income Southern California community, find that weatherization, efficiency, and electrification interventions can reduce energy costs and emissions, though not all interventions affect household bills in the same way. Their work is valuable for a finance-oriented framework because it demonstrates that retrofit interventions generate multidimensional outcomes: lower utility costs, lower emissions, different load shapes, and potential trade-offs. Hence, finance design should account for both burden relief and technical performance. The problem is not simply to lend more. It is to lend in ways that improve affordability without shifting transition risk onto households least able to absorb underperformance.

*Alternative data, fintech underwriting, and the role of utility payments*

Fintech scholarship provides a second pillar for the present paper: the use of alternative data in underwriting. Recent policy reviews show that alternative data can include bank-transaction records, rent payments, utility payments, employment signals, and other behavioral indicators that augment traditional bureau files (Schneider, 2024). The appeal of such data lies in their capacity to improve visibility for thin-file consumers and to support more context-sensitive assessments of cash-flow behavior. The Congressional Research Service also notes that alternative data raise legal and policy questions related to privacy, fairness, explainability, and compliance. This dual character—greater inclusion potential alongside greater governance complexity—is central to any proposal for energy-burden-aware underwriting.

The AI and financial-services literature offers both optimism and caution. Tierno (2024) notes that AI and machine learning are increasingly used across financial services to improve automation, pattern recognition, and decision support. In adjacent domains, recent finance studies show the versatility of predictive analytics in anti-money laundering, market surveillance, and financial reporting integrity (Hasan et al., 2023; Ibrahim et al., 2024a; Jahan et al., 2024; Pritty et al., 2024). These studies are not about retrofit finance, but they do show how financial institutions increasingly treat AI as a tool for early warning, anomaly detection, and scalable risk management. Likewise, Rasel et al. (2023) show how multimodal modeling may expand credit evaluation beyond traditional scores in mortgage contexts. Taken together, this adjacent literature suggests that multimodal underwriting architectures are technically feasible and increasingly normal within fintech; the unresolved question is how they should be designed for socially consequential climate-linked lending.

Utility-payment data occupy a distinctive place within this broader alternative-data conversation. Unlike some controversial nonfinancial data sources, utility bills have a direct conceptual relationship to residential decarbonization finance: they capture payment regularity, usage patterns, seasonal stress, and the economic case for efficiency improvements. They can also help identify whether a household's affordability problem is driven by low income alone, unusually high consumption, poor thermal performance, or combinations of these factors. Yet the literature has not fully resolved how utility data should be interpreted in credit decisions. A high bill could signal distress, but it could also signal strong retrofit potential. A missed payment could indicate fragile cash flow, but it may also reflect a housing system already imposing excessive operating costs. The underwriting challenge is therefore interpretive, not merely predictive.

Although outside consumer lending, work on AI-enabled operational decision systems in large retail environments also illustrates the broader institutional normalization of data-rich automation in economically consequential settings (Arman &

Fahim, 2023). This adjacent evidence does not validate retrofit underwriting by itself, but it reinforces the plausibility of deploying structured AI systems in settings where operational efficiency, risk management, and accountability must coexist.

*Explainability, fairness, and algorithmic accountability*

Because retrofit finance implicates household welfare and fair-access concerns, any AI-based system must be evaluated not only for predictive performance but also for transparency and accountability. The financial explainability literature has grown rapidly, with Černevičienė and Kabašinskas (2024) showing that explainable AI in finance remains methodologically diverse and still faces unresolved challenges around user trust, regulatory alignment, and practical deployment. Their review suggests that explainability in finance cannot be treated as a decorative add-on to black-box prediction. It must be linked to the actual decisions stakeholders need to understand—approval, pricing, product recommendation, exception handling, and model monitoring.

Fair lending research offers a sharper warning. Kozodoi, Jacob and Lessmann (2022) argue that fairness criteria in credit scoring require contextual interpretation and explicit trade-off management rather than generic technical optimization. Garcia, Pinto Garcia and Rigobon (2024) likewise review algorithmic discrimination in credit and show that data selection, proxy variables, and feedback loops can perpetuate exclusion even when protected variables are omitted. Ratzan and Rahman (2024) translate these concerns into an organizational governance lens, proposing that responsible AI in banking requires measurable governance dimensions rather than abstract ethical commitments. The relevance to retrofit finance is direct. A model that classifies high-burden households as universally high risk could systematically exclude the same communities targeted by public decarbonization policy. Conversely, a model that overcorrects without robust repayment logic could undermine lender confidence and long-term program viability.

The work by Fahim et al. (2023) is especially relevant here because it focuses on algorithmic accountability in U.S. consumer fintech. Their emphasis on governance mechanisms for credit risk and fair lending aligns with the present paper’s argument that an AI system for retrofit finance must be auditable, contestable, and policy-aware. In addition, Ibrahim et al. (2022) connect predictive analytics to climate-related financial risk and capital allocation, reminding us that climate-linked finance should not be treated as a niche ethical concern separate from mainstream financial stability and allocation logic. The present paper extends that insight from macro-level climate finance into household-level transition finance.

*From consumer credit logic to payment-feasibility logic*

When the foregoing strands of literature are considered together, a central conceptual shift becomes visible. Conventional consumer credit models ask whether a borrower resembles those who repaid similar obligations in the past. Energy-burden-aware retrofit underwriting asks whether a borrower can reasonably sustain payments when the financed intervention is expected to alter the household’s cost structure, comfort, and risk exposure. This is a different question. It requires combining retrospective indicators such as payment behavior with prospective indicators such as expected bill reduction, incentive capture, and dwelling-specific retrofit need. It also requires explicit recognition that some financing channels may be better suited than others depending on the borrower’s circumstances. On-bill structures, public guarantee layers, grant-blended products, and staged retrofit packages may sometimes be more appropriate than standard unsecured loans.

The literature therefore points toward a multimodal and governance-intensive underwriting architecture, but it does not yet provide a mature framework for one. That is the space this paper addresses. It synthesizes the retrofit, energy-burden, fintech, and responsible-AI literatures into a conceptual model centered on payment feasibility, explainability, and equity-aware product matching. Rather than assuming that better prediction alone will solve access problems, the framework treats underwriting as a socio-technical design problem in which model outputs, financing instruments, public incentives, and consumer protections must be aligned.

*Table 1. Summary of key literature informing the framework*

Source	Theme	Core insight	Implication for this paper
Berrill et al. (2022)	Residential decarbonization pathways	Shows that rapid renovation of existing homes and electrification are central to U.S. residential decarbonization pathways.	Retrofit finance is structurally important to decarbonization, not peripheral.
York et al. (2022)	Affordable housing decarbonization	Identifies cost, ownership structure, split incentives, and program design as barriers in affordable housing.	Financing products must account for affordability and implementation frictions.
Brown, Kale and	Energy burden	Demonstrates that annual averages can	Utility-bill data should be

Source	Theme	Core insight	Implication for this paper
Anthony (2023)	measurement	understate seasonal and situational burden.	modeled dynamically rather than as a static annual ratio.
Shen et al. (2023)	Community vulnerability and burden	Finds that community vulnerability strongly shapes energy burdens across U.S. counties.	Contextual variables are needed to interpret payment strain fairly.
Graff (2024)	Energy insecurity policy	Shows that assistance programs remain important but incomplete responses to household energy insecurity.	Underwriting should complement, not substitute for, broader affordability policy.
Hatch and Graff (2024)	Energy burden and eviction	Links neighborhood energy burden to eviction filing rates.	Energy affordability is connected to wider housing instability and should be treated as such.
Černevičienė and Kabašinskas (2024)	Explainable AI in finance	Reviews XAI methods and unresolved deployment challenges in finance.	Explainability must be tailored to decision context and user audience.
Kozodoi, Jacob and Lessmann (2022)	Fairness in credit scoring	Explains the trade-offs among fairness criteria and profit objectives.	Fairness must be monitored explicitly in credit-related models.
Charan et al. (2024)	Inclusive utility program design	Shows that inclusive program design can reduce upfront barriers for disadvantaged communities.	Product choice is part of underwriting, not a separate administrative afterthought.
Schneider (2024); Tierno (2024)	Alternative data and AI in finance	Summarize how utility, rent, and transaction data are entering financial-services decision systems, together with legal and governance concerns.	Utility-bill data have inclusion potential but require strong governance.

Note: This table is selective rather than exhaustive. It summarizes the studies and reports that most directly shape the paper's conceptual framework.

### 3. Research Gap

The literature reveals four connected gaps. First, residential decarbonization studies identify the scale and urgency of retrofit investment but rarely specify how underwriting should be redesigned for households whose energy burdens make conventional consumer credit a poor fit. Second, research on energy burden and energy insecurity shows that utility costs are socially patterned and deeply consequential, yet it usually stops short of translating those insights into concrete financing architecture. Third, fintech research on alternative data and AI underwriting demonstrates that nontraditional signals can expand visibility, but it has not developed a climate-transition-specific model that treats utility-bill patterns and housing inefficiency as both risk indicators and opportunity indicators. Fourth, responsible-AI scholarship in finance has established the importance of explainability, fairness, and governance, but it has not been closely integrated with retrofit finance, where the social objective is not only repayment but equitable decarbonization. This paper addresses these gaps by proposing an AI framework that links energy burden, housing characteristics, and retrofit economics to payment-feasibility assessment under explicit fairness and accountability constraints.

Table 2. Research gap and contribution mapping

Observed gap	What prior work shows	What remains unresolved	How this paper responds
Decarbonization literature specifies retrofit need but not underwriting redesign	Sector-wide pathways and building policy reports emphasize retrofits and electrification.	How should lenders assess payment feasibility when retrofits change household operating costs?	The paper reframes retrofit underwriting around projected affordability change.

Observed gap	What prior work shows	What remains unresolved	How this paper responds
Energy-burden research diagnoses hardship but rarely translates it into finance architecture	Energy burden and insecurity are shown to be unevenly distributed and socially consequential.	Burden is usually discussed as a policy problem, not a lending-design variable.	The paper treats energy burden as a structured input for product matching and context-sensitive credit assessment.
Alternative-data underwriting literature is not climate-transition specific	Utility and nontraditional payment data may widen visibility for thin-file consumers.	The literature does not define how utility bills should be interpreted in retrofit lending.	The framework distinguishes distress, structural inefficiency, and retrofit opportunity.
Responsible-AI scholarship in finance is under-connected to retrofit finance	Fairness, explainability, and governance are well recognized in finance AI.	Those principles are rarely operationalized in home decarbonization finance.	The paper embeds fairness diagnostics, explanation layers, and appeals processes into the proposed model.

**4. Research Objectives / Research Questions / Propositions**

This paper has three principal objectives. First, it synthesizes the literature on residential decarbonization, energy burden, alternative-data underwriting, and responsible AI to clarify the state of knowledge at their intersection. Second, it develops a conceptual underwriting framework in which utility-bill histories, housing characteristics, retrofit economics, and public incentives are used to assess payment feasibility more appropriately than conventional credit proxies alone. Third, it specifies a governance architecture for using AI in retrofit finance, including explainability outputs, fairness diagnostics, escalation rules, and product-selection logic.

These objectives generate the following research questions:

- RQ1. How can utility-bill and housing-performance information be incorporated into retrofit underwriting without reducing the model to a black-box risk score?
- RQ2. What fairness and explainability requirements are necessary for AI-assisted retrofit finance in high-stakes household decisions?
- RQ3. Under what conditions can energy-burden-aware underwriting improve access to residential decarbonization finance while preserving prudent risk management?

Because this is a conceptual paper, formal empirical hypotheses are not tested. Instead, the framework is organized around five propositions:

- P1. Underwriting models that incorporate utility-bill histories and dwelling-level energy context will assess retrofit payment feasibility more accurately than models relying on conventional credit variables alone in households with thin or noisy credit files.
- P2. Energy burden should be modeled as a contextual affordability indicator rather than interpreted only as a negative risk proxy.
- P3. Explainability mechanisms that connect approval, pricing, and product recommendations to understandable household-level drivers will improve contestability and governance quality.
- P4. Fairness-aware constraints and subgroup monitoring are necessary because unadjusted machine-learning models may reproduce housing and income inequalities embedded in historical payment data.
- P5. Product matching—such as grants, guarantees, tariffed on-bill structures, or blended-capital loans—should be treated as part of underwriting rather than as a separate downstream administrative choice.

**5. Methodology**

This paper uses a structured integrative review and conceptual design methodology. The aim is not to estimate treatment effects or report original model performance, but to develop a defensible framework from existing scholarship and policy material.

Integrative review is appropriate because the problem spans multiple domains—energy policy, housing, consumer finance, machine learning, and AI governance—and the relevant evidence includes peer-reviewed articles, technical reports, and high-authority public documents. The review prioritizes studies published before December 2024, with emphasis on 2020–2024 literature and selected earlier seminal works where necessary for conceptual grounding.

The source base includes peer-reviewed journal articles, National Institute of Standards and Technology publications, U.S. Department of Energy and National Renewable Energy Laboratory reports, American Council for an Energy-Efficient Economy briefs and reports, and Congressional Research Service materials. Searches were conducted using combinations of terms such as “energy burden”, “energy insecurity”, “residential decarbonization”, “retrofit finance”, “building electrification”, “alternative data”, “utility payments”, “credit underwriting”, “fair lending”, “explainable AI”, and “responsible AI in finance”. Database and search-string details are summarized in Appendix A. Sources were included when they addressed at least one of four themes: (1) household energy burden and insecurity; (2) retrofit, electrification, or decarbonization finance; (3) AI, alternative data, or credit underwriting in finance; and (4) explainability, fairness, or governance in algorithmic decision-making. Sources that fell outside the paper’s U.S.-relevant conceptual scope or were published after November 2024 were excluded.

The review proceeded in three steps. First, the literature was organized into thematic clusters and coded for problem definition, variables or constructs used, treatment of equity and affordability, and implications for finance design. Second, the themes were compared to identify where the literatures overlapped and where they left unresolved tensions. Third, these tensions were synthesized into a conceptual framework composed of six layers: data intake, feature construction, model selection, explanation interface, governance controls, and financing-channel recommendation. This process is interpretive but systematic. It makes no claim that the resulting framework is the only possible design; rather, it is presented as a rigorous and policy-aware architecture supported by the best available evidence.

No empirical model was trained for this paper, and no proprietary or household-level dataset was used. Accordingly, the “results” section below should be read as an analytical framework and expected-results logic rather than as reported field outcomes. To make this explicit, the paper also specifies a future empirical testing plan. That plan includes candidate outcome variables, fairness metrics, consumer-explanation requirements, and validation procedures that could be used in later work by lenders, public green banks, utilities, or researchers. Validity in the present study therefore rests on conceptual coherence, breadth and quality of sources, consistency with established responsible-AI principles, and transparency about the paper’s non-empirical scope.

Ethically, the framework assumes that household retrofit finance is a high-stakes decision context. As such, it should be subject to minimization of unnecessary data collection, informed-consent safeguards where applicable, strict limits on sensitive proxy use, robust documentation, adverse-action explanation procedures, and an option for human review in borderline or contested cases. These design commitments are derived from the responsible-AI and fair-lending literature rather than from a completed institutional compliance review. They are therefore proposed governance requirements, not evidence that any specific lender currently meets them. Terminology around trustworthiness, explainability, and risk management follows the NIST AI glossary and AI Risk Management Framework (Atherton et al., 2023; Tabassi, 2023).

## 6. Results or Analytical Framework

This section presents the paper’s central output: a proposed AI framework for energy-burden-aware underwriting in residential decarbonization finance. The framework begins from a simple but important premise: payment feasibility for retrofit finance should be estimated using information about how the retrofit changes the household’s energy-cost structure, not solely how the household performed under pre-retrofit conditions. Conventional underwriting may view a high utility burden as evidence of fragility. The proposed framework views it as an ambiguous but economically meaningful signal. A household with high summer cooling bills in an inefficient home may indeed face repayment pressure, but it may also be one of the cases in which a well-designed retrofit produces the most immediate affordability gain. The underwriting task is therefore to disentangle burden due to structural inefficiency from burden due to non-payment propensity, and then to align financing accordingly.

The framework is multimodal because no single data source can support that judgment. Utility-bill histories provide recurring evidence of payment timing, arrears, seasonal volatility, and consumption intensity. Housing and building data provide information about dwelling age, size, envelope quality, climate zone, heating fuel, appliance stock, and likely retrofit potential. Applicant data provide conventional credit indicators, income information, tenure status, and any evidence of income volatility. Programmatic data provide information on rebates, tax credits, weatherization grants, green-bank guarantees, and utility incentives. Finally, retrofit-package data specify the expected intervention, its projected capital cost, and its estimated effects on energy consumption and monthly bills. The literature supports the inclusion of each of these domains separately, but their joint use is where the proposed contribution lies.

At the data-intake stage, the framework distinguishes between core variables and context variables. Core variables are those directly relevant to financing and retrofit feasibility: recent utility-payment regularity, estimated pre-retrofit energy burden, projected monthly savings, applicant income, and financing amount. Context variables include housing condition, local weather severity, dwelling fuel mix, neighborhood vulnerability markers, and the presence of stacked incentives. This distinction is important because context variables should not automatically drive adverse decisions. Their role is to improve interpretation of core affordability signals. For example, late utility payments may have different meaning in a drafty all-electric dwelling with unusually high winter loads than in a relatively efficient home with similar income but discretionary spending stress. Context-aware interpretation is one of the places where AI can outperform static scorecards, provided that the model is auditable.

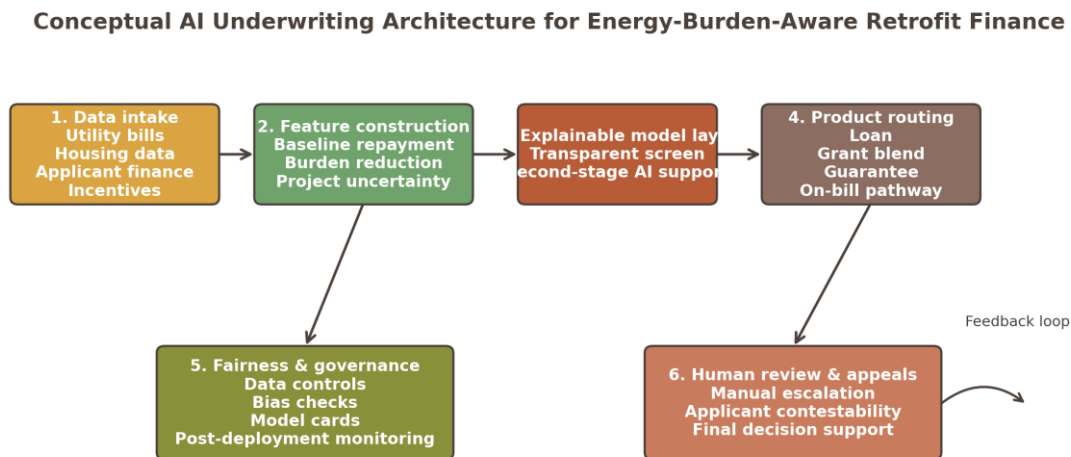
Feature construction is the heart of the proposed framework. Rather than creating a single undifferentiated risk score, the model should generate three analytically separate latent constructs: baseline repayment capacity, retrofit-induced affordability improvement, and implementation or performance uncertainty. Baseline repayment capacity uses traditional features such as payment history, debt load, income stability, and existing obligations. Retrofit-induced affordability improvement uses estimated bill savings, burden reduction potential, and net monthly payment after incentives. Implementation or performance uncertainty captures the possibility that the retrofit will not generate expected savings because of contractor quality, household behavior, housing defects, or technology mismatch. Separating these constructs prevents a common failure mode in automated lending: folding all uncertainty into borrower-level risk and thereby treating structural housing inefficiency as if it were identical to borrower unreliability.

*Table 3. Proposed variable domains for energy-burden-aware underwriting*

<b>Variable domain</b>	<b>Illustrative inputs</b>	<b>Analytical role</b>	<b>Governance consideration</b>
Utility-bill and payment history	Monthly bills, payment timing, arrears episodes, seasonal volatility, shutoff notices where lawful	Captures stress, regularity, and seasonality of household energy costs	High: directly relevant but must be normalized for climate and billing anomalies
Housing and building characteristics	Dwelling age, floor area, tenure, fuel type, HVAC type, envelope condition, climate zone	Identifies structural inefficiency and retrofit potential	High: avoid stale or overly coarse property proxies
Applicant financial indicators	Income, debt obligations, liquid reserves where available, conventional credit file, employment stability	Anchors baseline repayment capacity	High: required, but should not dominate when retrofit materially changes affordability
Retrofit package attributes	Intervention type, capital cost, projected savings, contractor quality flags, pre-retrofit repairs	Estimates transition benefit and implementation uncertainty	High: separates borrower risk from project-performance risk
Public-support variables	Rebates, tax credits, guarantees, weatherization eligibility, green-bank support	Captures the real financing environment rather than a private-credit-only view	High: requires current program data
Community and vulnerability context	Neighborhood burden, energy vulnerability, exposure to extreme temperatures, housing stress indicators	Provides interpretive context for affordability and fairness monitoring	Medium to high: should guide review and monitoring, not become blunt exclusion triggers

Note: Variables listed here are illustrative and should be implemented subject to applicable consumer-protection, privacy, consent, and fair-lending requirements.

Figure 1. Conceptual architecture of the proposed AI underwriting framework



Source: author-designed conceptual framework informed by the reviewed literature.

An especially important design choice is how to represent energy burden. The proposed framework does not use energy burden as a standalone penalty variable. Instead, it uses burden in interaction with dwelling characteristics and projected retrofit effects. High burden combined with low projected savings may indicate that financing should not proceed without grants or alternative products. High burden combined with strong projected savings and high payment regularity may indicate a compelling case for transition finance. Moderate burden combined with irregular payment timing may call for a smaller staged retrofit or an on-bill structure rather than a conventional loan. In short, burden becomes a stratification and product-matching variable rather than a blunt exclusion rule. This reinterpretation is theoretically grounded in the energy-justice literature and practically necessary if retrofit finance is to avoid reinforcing existing inequities.

Model selection should also reflect the paper’s governance commitments. The framework favors a tiered architecture rather than a single opaque model. A transparent baseline model—such as regularized logistic regression or gradient-boosted trees with strong explanation tooling—can be used for primary screening, while more complex multimodal models can be reserved for recommendation support or second-stage review. The purpose of this design is not to reject technical sophistication; it is to preserve tractability in a high-stakes domain. As Černevičienė and Kabašinskas (2024) note, explainability in finance is inseparable from context. A model may be highly predictive and still unsuitable if its decision logic cannot be meaningfully communicated to applicants, compliance staff, and program administrators. For that reason, the framework assumes that model complexity should be justified against explanation burden.

Explanation outputs in the proposed framework are multi-audience by design. Applicants should receive plain-language explanations centered on the main drivers of approval, pricing, or product recommendation—for example, high seasonal bills, stable utility payment behavior, limited debt capacity, or strong projected savings after rebates. Underwriters and program managers should receive richer diagnostic explanations such as feature-attribution summaries, threshold triggers, sensitivity to incentive assumptions, and reasons for escalation to manual review. Compliance and audit teams should receive model cards, subgroup performance reports, and records of feature governance. These explanation layers are necessary because “explainability” means different things to different users. A household deciding whether to take on retrofit debt does not need a Shapley tutorial; it needs a credible account of how the financing decision was made and what assumptions may affect affordability.

The product-selection component is what most clearly differentiates the proposed framework from ordinary credit scoring. Underwriting here is not limited to a binary approve/deny output. Instead, the model recommends a financing pathway conditional on household circumstances. Possible pathways include a standard retrofit loan, a grant-blended loan, an interest-

rate buydown, a public guarantee structure, an on-bill repayment pathway, or a deferral pending a lower-cost staged package. This is justified by the literature on inclusive utility programs and affordable-housing decarbonization, which repeatedly shows that one-size-fits-all financing mechanisms are poorly matched to heterogeneous households (York et al., 2022; Charan et al., 2024; Supple, 2024). A household should not be labeled “too risky” when the real issue is that the wrong financial instrument has been selected.

*Table 4. Governance and accountability matrix for deployment*

<b>Governance layer</b>	<b>Operational requirement</b>	<b>Illustrative evidence or control</b>
Data governance	Document data provenance, refresh cycles, normalization rules, and permissible uses	Model cards, data dictionaries, consent records where applicable
Fairness review	Test subgroup performance and proxy risks; compare multiple fairness metrics	Approval-rate monitoring, error-rate differentials, pricing audits
Explainability	Provide plain-language applicant explanations and richer internal diagnostics	Adverse-action notices, explanation templates, model-attribution reports
Human oversight	Escalate edge cases, contested cases, or low-confidence cases to manual review	Escalation logs, override rationale, appeals outcomes
Product suitability	Treat product selection as part of the underwriting decision	Routing rules for loans, grant blends, guarantees, or utility-linked pathways
Post-deployment monitoring	Track realized affordability, arrears, and retrofit completion over time	Outcome dashboards, drift detection, periodic model recalibration

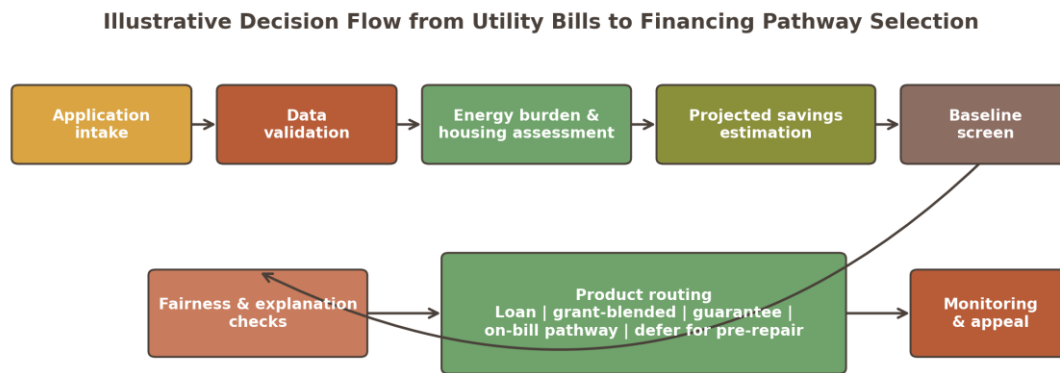
The framework also incorporates a retrofit-readiness screen. Some homes require pre-weatherization repairs, wiring upgrades, moisture remediation, or landlord coordination before efficiency or electrification investments will generate reliable savings. If these conditions are ignored, the underwriting model may overestimate benefit and understate household exposure. The framework therefore treats technical readiness as a separate module linked to implementation risk. In practice, this means that the financing recommendation may route some households to grants, public programs, or pre-retrofit assistance rather than debt finance. Such routing is consistent with the energy-insecurity literature, which cautions against solutions that assume all households can convert long-term savings into near-term payment capacity without structural support.

Fairness governance is embedded across the model lifecycle rather than appended at the end. At the development stage, the framework requires feature review for proxy discrimination, subgroup error analysis, and comparison across multiple fairness criteria rather than reliance on a single score. At the deployment stage, it requires monitoring of approval rates, pricing outcomes, escalation patterns, and realized delinquency across protected or policy-relevant groups where legally and ethically appropriate. At the post-decision stage, it requires an appeals or adverse-action process through which households can contest errors in utility records, housing data, income assumptions, or incentive calculations. This design follows the broader literature’s insistence that fairness is not a one-time pre-processing exercise but an ongoing governance obligation (Kozodoi, Jacob & Lessmann, 2022; Garcia, Pinto Garcia & Rigobon, 2024; Ratzan & Rahman, 2024).

An important conceptual implication follows from this architecture: underwriting quality should be evaluated using more than default prediction. A model could reduce short-term credit losses by declining precisely the applicants facing the greatest transition barriers, yet such a model would fail the social and policy objective of equitable decarbonization. The proposed framework therefore uses a balanced outcome logic. Relevant performance dimensions include predictive accuracy, burden reduction potential, approval parity conditional on relevant characteristics, explanation quality, alignment of product type to household need, and realized affordability after retrofit. This multidimensional evaluation approach is essential because the public value of retrofit finance lies partly in what happens after disbursement: whether bills fall, whether comfort improves, whether energy insecurity declines, and whether households avoid harmful debt stress.

The framework can be summarized as a movement from descriptive data to normative decision support. Utility bills, housing conditions, and credit information are used not merely to classify households, but to recommend the least exclusionary viable path into decarbonization finance. That recommendation is then bounded by governance rules that constrain unacceptable trade-offs. Conceptually, this is where the paper's AI contribution lies. The framework is not "AI for AI's sake"; it is a proposal for how computational finance can be reoriented toward household transition feasibility under explicit social constraints.

Figure 2. Illustrative decision flow from utility bills to financing pathway selection



Source: author-designed flow based on the paper's proposed underwriting logic.

Because no empirical dataset is analyzed, the appropriate way to present "results" is as a structured expected-results framework derived from the literature. The central expectation is not that AI will universally expand credit access, but that an energy-burden-aware system will improve decision quality along multiple margins relative to conventional underwriting. First, it should identify households for whom projected burden reduction materially changes feasible payment capacity. Second, it should reduce false negatives among households with thin credit files but stable utility behavior and strong retrofit economics. Third, it should improve alignment between financing product and household context, especially where grants, guarantees, or utility-linked repayment structures can address constraints that ordinary loans cannot. Fourth, it should provide better documentation for human underwriters, compliance staff, and applicants about why a financing outcome occurred.

The literature also suggests that implementation quality will determine whether these expected gains materialize. If projected savings are generated from crude assumptions, then even a fair and transparent model may produce poor recommendations. If utility data are incomplete or not normalized for climate, occupancy, and billing anomalies, the model may mistake measurement noise for affordability stress. If public incentive availability is not updated in real time, the system may understate feasible burden reduction. Therefore, one of the framework's implicit results is organizational rather than algorithmic: institutions attempting this form of underwriting will need stronger data governance, retrofit estimation protocols, and cross-functional coordination than standard consumer-lending products require.

A future empirical implementation should evaluate performance across at least five dimensions. The first is predictive validity, including delinquency, default, prepayment, and realized payment-to-income outcomes. The second is transition effectiveness, such as estimated and realized bill savings, burden reduction, and retrofit completion rates. The third is fairness, including subgroup differences in approvals, pricing, adverse-action explanations, and realized error rates. The fourth is explanation quality, measured through user testing with applicants and staff. The fifth is program fit, including the proportion of applications routed to an appropriate financing channel rather than denied under a default product. The point of this multi-criteria evaluation is to avoid the narrow benchmarking practice that rewards models for conservative rejection without considering whether they enable socially valuable and economically workable retrofits.

This expected-results logic also clarifies the role of human judgment. The proposed framework does not envision a fully autonomous system. Instead, it treats AI as a decision-support layer embedded within a larger governance structure. Automated recommendations should be particularly helpful in cases where conventional scorecards underperform: mixed evidence on affordability, thin files, high but explainable utility burdens, or complex incentive stacks. Human review remains essential where the data are incomplete, where housing-condition uncertainty is high, or where applicants contest explanation outputs. In this respect, the proposed system is best understood as a hybrid governance architecture rather than as a replacement for institutional discretion.

If tested successfully, the framework would have implications beyond retrofit lending. It would suggest that fintech models in other socially consequential domains should be designed around mechanism-specific feasibility rather than generic repayment resemblance. Energy-burden-aware underwriting is therefore both a substantive proposal for residential decarbonization and a methodological example of how AI can be made more purpose-specific in financial decision-making. That broader implication reinforces the value of studying retrofit finance not as a niche green product, but as a demanding test case for responsible fintech design.

*Table 5. Proposed evaluation and expected-results framework for future empirical testing*

<b>Evaluation dimension</b>	<b>Illustrative metrics</b>	<b>Expected pattern if the framework works as intended</b>
Predictive validity	Delinquency, default, prepayment, realized payment-to-income ratio	At least comparable or better risk performance than conventional scorecards in target segments
Affordability improvement	Estimated versus realized utility-bill reduction; change in energy burden	Meaningful burden reduction after financing and retrofit completion
Equity and fairness	Approval, pricing, override, and adverse-action patterns across subgroups	Reduced unjustified disparities and documented monitoring process
Explanation quality	Applicant comprehension, staff usability, contestability of decisions	Higher procedural clarity and lower unresolved disputes
Program fit	Correct routing to loan, guarantee, grant-blended, or utility-linked option	Lower denial rates caused by product mismatch rather than true infeasibility

Note: This table is a template for future empirical validation. It does not report observed results from a completed study.

**7. Discussion**

The proposed framework contributes to the literature by reclassifying utility-bill information from a peripheral affordability measure into a central underwriting input. Existing scholarship often treats bills in one of two ways: as evidence of household distress in the energy-poverty literature or as one item among many forms of alternative data in fintech. The present framework bridges these interpretations. It argues that utility-bill histories should be read through the lens of housing performance and retrofit potential, not as simple analogues of missed credit-card payments. This is a meaningful conceptual shift because it recognizes that some payment stress is structurally produced by inefficient homes and therefore may be partially remediable through the financed intervention itself.

The framework also extends current debates on explainable AI in finance by locating explanation within a concrete policy use case. Much of the XAI literature asks how models can be made more interpretable in general financial settings. Retrofit finance clarifies why this matters. Households need to understand not only why they were approved or denied but also why one financing channel was recommended over another, how estimated savings were used, and what assumptions would have to change for the decision to differ. This decision context therefore pushes beyond model-centric explanation toward action-oriented explanation. In that sense, the paper adds a practical governance lens to the XAI literature: explanations are valuable not only because they improve trust, but because they make transition finance contestable and administratively usable.

Another contribution concerns the relationship between inclusion and prudence. In consumer-finance debates, inclusion is sometimes discussed as if it were necessarily in tension with rigorous underwriting. The present framework suggests a different possibility. A model that accounts for projected burden reduction, incentive support, and retrofit readiness may make more prudent decisions precisely because it better understands the economics of the financed intervention. Ignoring those variables is

not neutral; it may systematically misprice or reject viable households by relying on pre-retrofit distress signals alone. Thus, fairness and prudence need not be treated as opposing poles. Under the right design conditions, they can be partially aligned through richer modeling of household cost structures.

At the same time, the framework does not assume that more data automatically produce better or fairer outcomes. Utility-payment data can be noisy, incomplete, or shaped by billing-system idiosyncrasies. Housing-condition data may be stale or inferred from weak proxies. Savings estimates may be optimistic. And any model trained on historical repayment patterns may import the exclusionary logic of prior markets. For these reasons, the paper insists on product matching, manual-review pathways, and post-deployment monitoring. These are not peripheral safeguards; they are part of what makes the framework intellectually coherent. Without them, a climate-finance model could simply become another instrument for rationing credit under the language of innovation.

The citations become especially useful in this discussion because they illustrate how AI capabilities developed in adjacent financial domains can be repurposed but should not be transplanted uncritically. Ibrahim et al. (2024) and Jahan et al. (2024) demonstrate the value of predictive analytics and early warning systems in financial infrastructure monitoring, while Pritty et al. (2024) examine AI-assisted detection of narrative manipulation in disclosure settings. These applications show that finance increasingly relies on AI to identify weak signals and nonobvious patterns. Yet retrofit finance is qualitatively different because the objective is not only to detect risk but to structure a pathway toward socially valuable investment. The implication is that technical methods may transfer across domains, but the objective function and governance logic must change.

Finally, the framework underscores that residential decarbonization should be understood as a financial-inclusion problem as well as a climate problem. If capital flows primarily to households already advantaged by stable credit profiles, then decarbonization may lower emissions while leaving energy burdens and housing inequality largely intact. A transition framed only through aggregate emissions can therefore miss distributional failure. By centering payment feasibility, energy burden, and product selection, the paper offers a way to connect the climate-transition agenda to the everyday economics of household bills. That connection is essential if retrofit finance is to support both decarbonization and affordability rather than trading one objective against the other.

## 8. Theoretical Implications

The paper has three main theoretical implications. First, it connects climate-transition finance to the literature on household energy burden by arguing that operating-cost stress should be treated as a core underwriting variable rather than as an exogenous social condition. Second, it extends fair-ML scholarship in finance by proposing payment-feasibility logic as a domain-specific lens through which fairness criteria can be interpreted. Third, it broadens the concept of alternative data in fintech by emphasizing semantically grounded data—utility payments and housing characteristics that bear a direct relationship to the financed activity—over opportunistic expansion into weakly related behavioral proxies. In this sense, the paper contributes to an emerging theory of purpose-specific underwriting in which model design is shaped by the economic mechanism of the financed intervention.

## 9. Practical / Policy / Managerial Implications

For lenders and fintech firms, the framework suggests that residential decarbonization finance should be organized around modular decision support rather than monolithic scorecards. Institutions should separate borrower repayment capacity, projected retrofit savings, and implementation uncertainty, and then connect those assessments to product design. For utilities and public program administrators, the framework indicates that billing data and program eligibility data can be useful inputs to financing pathways, provided governance and consent requirements are satisfied. For public green banks and housing agencies, the framework provides a basis for blended-capital structures in which public support is targeted where private underwriting alone would misread retrofit feasibility. For regulators, the paper points to a governance agenda that includes model documentation, adverse-action explanation, proxy-risk review, and subgroup monitoring in climate-linked lending programs. For community organizations and housing counselors, the framework highlights the importance of applicant support, data correction processes, and education about retrofit assumptions so that households are not asked to trust opaque systems.

## 10. Limitations

This paper has several limitations. Most importantly, it is conceptual and literature-based. It does not estimate predictive performance, default rates, realized energy savings, or consumer responses to explanations. The framework therefore remains a design proposition rather than a validated production model. Second, the evidence base spans disciplines with different assumptions, methods, and units of analysis; this strengthens conceptual breadth but limits direct comparability. Third, several of the supporting reports are policy or technical documents rather than journal articles, although they were selected because they are authoritative and directly relevant to U.S. residential decarbonization. Fourth, the framework is U.S.-specific in its policy references and institutional assumptions. Finally, some references are adjacent rather than central to retrofit finance; they have

been integrated where they support transferable ideas about predictive analytics, multimodal modeling, and algorithmic accountability, but they should not be read as direct empirical evidence on residential decarbonization.

**11. Future Research Directions**

Future research should test the framework empirically with real program and household data. A first priority is model comparison: researchers should evaluate whether energy-burden-aware models improve prediction of repayment and post-retrofit affordability relative to traditional scorecards. A second priority is product-level experimentation: different households may be better served by standard loans, tariffed on-bill pathways, guarantee-backed products, or staged retrofit packages, and these choices should be analyzed as part of underwriting rather than after it. A third priority is explanation research focused on consumers. Experimental work could test which forms of explanation improve comprehension, acceptance, and perceived procedural fairness without generating information overload. A fourth priority is subgroup fairness analysis across tenure types, race-proxy concerns, neighborhood vulnerability, and housing conditions. Finally, longer-run studies should examine whether retrofit finance reduces not only monthly bills but also broader forms of insecurity, including arrears, utility shutoff risk, eviction exposure, and health-related stress.

**12. Conclusion**

Residential decarbonization requires more than subsidies, better technologies, or stronger climate targets. It also requires financial systems capable of recognizing that some households face high energy costs not because they are inherently poor credit risks, but because they live in inefficient homes and operate under structural affordability constraints. This paper has argued that an AI framework for retrofit finance should therefore be built around payment feasibility, not conventional credit logic alone. By integrating utility bills, housing characteristics, retrofit economics, and responsible-AI governance, the proposed framework offers a way to connect fintech innovation to equitable household decarbonization. Its central claim is modest but important: better underwriting for the energy transition will come not from more opaque prediction, but from more context-sensitive, explainable, and product-aware decision systems. The framework now requires empirical testing, but as a conceptual contribution it clarifies the design principles needed if retrofit finance is to expand decarbonization while reducing, rather than reproducing, household energy burden.

E. Appendices / supplementary attachments

*A. Appendix A. Search strategy and source-selection logic*

- Primary thematic strings included combinations of: 'energy burden', 'energy insecurity', 'residential decarbonization', 'building electrification', 'retrofit finance', 'alternative data', 'utility payments', 'credit underwriting', 'fair lending', 'explainable AI', 'responsible AI', and 'consumer finance'.
- Academic sources were prioritized when they were peer reviewed and directly relevant to at least one of the paper’s four thematic pillars: energy burden; retrofit and decarbonization; alternative-data underwriting; and AI governance.
- Policy and technical sources were included when produced by authoritative institutions with clear relevance to the U.S. buildings sector, financial-services regulation, or AI governance, such as NIST, DOE, NREL, ACEEE, and the Congressional Research Service.
- The search window was restricted to materials published before December 2024, with selected earlier seminal sources retained where necessary for conceptual grounding.

*B. Appendix B. Inclusion and exclusion criteria*

*Appendix Table B1. Inclusion and exclusion criteria*

Criterion	Operational rule
Included	Peer-reviewed or high-authority sources directly relevant to one or more framework components; U.S.-applicable policy and technical sources; multimodal modeling, or accountability.
Excluded	materials with unclear provenance; articles unrelated to residential energy, underwriting, or AI governance; unsupported empirical claims lacking sufficient metadata or credibility.

## C. Appendix C. Illustrative literature matrix

Cluster	Representative sources	Main takeaway for the framework
Energy burden	Brown, Kale and Anthony (2023); Shen et al. (2023); Graff (2024); Hatch and Graff (2024)	Energy burden is multidimensional, spatially uneven, and linked to broader precarity.
Residential decarbonization	Berrill et al. (2022); York et al. (2022); Langevin et al. (2024); Supple (2024); Charan et al. (2024)	Financing and program design are central to equitable retrofit deployment.
Fintech and alternative data	Schneider (2024); Tierno (2024); Rasel et al. (2023); Ibrahim et al. (2024); Jahan et al. (2024)	AI and nontraditional data can improve visibility and monitoring, but governance remains essential.
Responsible AI / fair lending	Černevičienė and Kabašinskas (2024); Kozodoi, Jacob and Lessmann (2022); Garcia et al. (2024); Ratzan and Rahman (2024); Fahim et al. (2023)	Explainability, fairness, and accountability must be built into the model lifecycle.

## D. Appendix D. Illustrative implementation notes

- This framework should be piloted first in a limited program setting with strong data-cleaning capacity and an established retrofit delivery network.
- Projected savings should be estimated conservatively and periodically back-tested against realized post-retrofit bills.
- Applicants should be able to request correction of utility records, occupancy assumptions, housing descriptors, or incentive assumptions before final adverse action.
- Program administrators should distinguish between loan denial due to true infeasibility and denial due to product mismatch or missing pre-retrofit support.

**Funding:** This research received no external funding.

**Conflicts of Interest:** The authors declare no conflict of interest.

**Publisher's Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

## References

- [1]. American Council for an Energy-Efficient Economy. (2024). Data update: City energy burdens.
- [2]. Arman, M., & Fahim, A. S. M. (2023). AI revolutionizes inventory management at retail giants: Examining Walmart's U.S. operations. *Journal of Business and Management Studies*, 5(6), 145–148. <https://doi.org/10.32996/jbms.2023.5.6.15>
- [3]. Atherton, D., Schwartz, R., Fontana, P., & Hall, P. (2023). The language of trustworthy AI: An in-depth glossary of terms. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-3>
- [4]. Berrill, P., Wilson, E. J. H., Reyna, J. L., Fontanini, A. D., & Hertwich, E. G. (2022). Decarbonization pathways for the residential sector in the United States. *Nature Climate Change*, 12(8), 712–718. <https://doi.org/10.1038/s41558-022-01429-y>
- [5]. Bowen, T. (2023). Louisville Communities LEAP engagement: Improving energy efficiency in affordable housing. National Renewable Energy Laboratory. <https://doi.org/10.2172/2331430>
- [6]. Brown, M. A., Kale, S., & Anthony, R. (2023). Rescaling energy burden: Using household surveys to examine vulnerabilities and consequences in the Southeastern United States. *Energy Research & Social Science*, 106, Article 103308. <https://doi.org/10.1016/j.erss.2023.103308>
- [7]. Charan, T., Langner, R., Stenger, K., Simmons, C., Fallahi, Z., Everett, M., Moran, C., Csinsi, N., & Alter, Z. (2024). Driving affordable decarbonization: Inclusive utility program designs for disadvantaged communities. National Renewable Energy Laboratory. <https://doi.org/10.2172/2447819>

- [8]. Černevičienė, J., & Kabašinskas, A. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. *Artificial Intelligence Review*, 57, Article 216. <https://doi.org/10.1007/s10462-024-10854-8>
- [9]. Fahim, A. S. M., Ibrahim, M., Pritty, A. A. A., & Tania, T. A. (2023). Algorithmic accountability in U.S. consumer FinTech: Governance mechanisms for credit risk, fair lending, and financial stability. *Journal of Economics, Finance and Accounting Studies*, 5(4), 80–93. <https://doi.org/10.32996/jefas.2023.5.4.8>
- [10]. Flores, R., Houssainy, S., Wang, W., Nguyen Cu, K., Nie, X., Woolfolk, N., Polly, B., Faramarzi, R., Maclay, J., Lee, J., & Brouwer, J. (2024). Addressing building related energy burden, air pollution, and carbon emissions of a low-income community in Southern California. *Advances in Applied Energy*, 14, Article 100169. <https://doi.org/10.1016/j.adapen.2024.100169>
- [11]. Garcia, A. C. B., Pinto Garcia, M. G., & Rigobon, R. (2024). Algorithmic discrimination in the credit domain: What do we know about it? *AI & Society*, 39, 2059–2098. <https://doi.org/10.1007/s00146-023-01676-3>
- [12]. Graff, M. (2024). Addressing energy insecurity: Policy considerations for enhancing energy assistance programs. *Heliyon*, 10(3), e24178. <https://doi.org/10.1016/j.heliyon.2024.e24178>
- [13]. Hatch, M. E., & Graff, M. (2024). Housing costs are not a monolith: The association between neighborhood energy burdens and eviction filing rates. *Cities*, 150, Article 104995. <https://doi.org/10.1016/j.cities.2024.104995>
- [14]. Hasan, M. N., Rasel, I. H., Arman, M., Ibrahim, M., & Jahan, N. (2023). Strengthening U.S. financial and cybersecurity infrastructure with AI-driven fraud detection and risk analytics. *Journal of Computational Analysis and Applications*, 31(2), 15–32. <https://www.eudoxuspress.com/index.php/pub/article/view/3823>
- [15]. Ibrahim, M., Huda Razib, M. N., Jahan, N., & Rahman, M. M. (2022). Climate risk, financial stability, and global capital allocation: A predictive analytics approach to assessing climate-related financial risk in international investment markets. *Journal of Business and Management Studies*, 4(4), 264–276. <https://doi.org/10.32996/jbms.2022.4.4.34>
- [16]. Ibrahim, M., Mahmud, S., Zadid, M. U., Jahan, N., Rahman, M. M., & Fahim, A. S. M. (2024). AI-driven predictive analytics framework for anti-money laundering risk management and financial infrastructure protection in U.S. banking systems. *Journal of Economics, Finance and Accounting Studies*, 6(1), 155–166. <https://doi.org/10.32996/jefas.2024.6.6.12>
- [17]. Jahan, N., Pritty, A. A., Ibrahim, M., Zadid, M. U., Fahim, A. S. M., & Mahmud, S. (2024). Machine learning-driven early warning analytics for identifying market manipulation, irregular trading activity, and suspicious market signals in U.S. stock markets. *Journal of Computer Science and Technology Studies*, 6(2), 257–283. <https://doi.org/10.32996/jcsts.2024.6.2.26>
- [18]. Kozodoi, N., Jacob, J., & Lessmann, S. (2022). Fairness in credit scoring: Assessment, implementation and profit implications. *European Journal of Operational Research*, 297(3), 1083–1094. <https://doi.org/10.1016/j.ejor.2021.06.023>
- [19]. Langevin, J., Wilson, E., Snyder, C., Narayanamurthy, R., Miller, J., Kaplan, K., Reiner, M., Risser, R., Mahoney, M., Geyer, J., & Ciraulo, R. (2024). Decarbonizing the U.S. economy by 2050: A national blueprint for the buildings sector. U.S. Department of Energy. <https://doi.org/10.2172/2338089>
- [20]. Pritty, A. A., Ibrahim, M., Fahim, A. S. M., & Zadid, M. U. (2024). Generative AI and U.S. financial reporting integrity: Detecting narrative manipulation, risk disclosure gaming, and fraud signals in 10-K filings. *Journal of Economics, Finance and Accounting Studies*, 6(4), 113–129. <https://doi.org/10.32996/jefas.2024.6.4.1>
- [21]. Rasel, I. H., Ibrahim, M., Pritty, A. A., Fahim, A. S. M., & Jahan, N. (2023). Beyond FICO: Enhancing mortgage default forecasting and inclusive lending via multimodal graph neural networks and urban mobility analytics. *Frontiers in Computer Science and Artificial Intelligence*, 2(2), 62–81. <https://doi.org/10.32996/fcsai.2023.2.2.5>
- [22]. Ratzan, J., & Rahman, N. (2024). Measuring responsible artificial intelligence (RAI) in banking: A valid and reliable instrument. *AI and Ethics*, 4(4), 1279–1297. <https://doi.org/10.1007/s43681-023-00321-5>
- [23]. Schneider, K. E. (2024). Alternative data in financial services (CRS In Focus IF11630). Congressional Research Service.
- [24]. Shen, Z., Chen, C.-f., Zhou, H., Fefferman, N., & Shrestha, S. (2023). Community vulnerability is the key determinant of diverse energy burdens in the United States. *Energy Research & Social Science*, 97, Article 102949. <https://doi.org/10.1016/j.erss.2023.102949>
- [25]. Supple, L. (2024). Equitable strategies for residential building energy efficiency and electrification in San José, California. National Renewable Energy Laboratory.
- [26]. Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- [27]. Tierno, P. (2024). Artificial intelligence and machine learning in financial services (CRS Report R47997). Congressional Research Service.
- [28]. U.S. Department of Energy. (2024, April 3). DOE releases first ever federal blueprint to decarbonize America's buildings sector.
- [29]. York, D., Cohn, C., Morales, D., & Tolentino, C. (2022). Building decarbonization solutions for the affordable housing sector. American Council for an Energy-Efficient Economy.