ABFER 13th Annual Conference | 18-21 May 2026 Singapore

Writing Quality and Soft Information in the GenAI Age: Evidence from Online Credit Markets

Lin William Cong (Cornell University, ABFER, IC3, & NBER)
Yanhong Guo (Dalian University of Technology)
Xin Zhao (Dalian University of Technology)
Wenjun Zhou (University of Tennessee Knoxville)

May 20th, 2026 Behavioural, Experimental Economics and Finance

Key Findings

1. Generative AI (GenAI) can *significantly* enhance **writing quality** and **perceived loan quality** based on text descriptions.
2. GenAI is particularly helpful for **high-credit** borrowers, who treat good credit and quality writing as substitutes, whereas they are complements for receiving funding, as suggested by historical data.
3. Due to the convergence in writing, GenAI **decreases soft information** conveyed, which also leads to **misallocation** (~15% reduction in lender ROI).
4. If lenders adjust their lending decisions considering GenAI adoption, they recapture some soft information, mitigating potential credit misallocation.
5. We apply a novel framework to analyze **counterfactual equilibrium** LLM adoption and online lending, in a **data-driven** manner.

Literature and Contributions

- **Does it pay to write well? And do LLMs help?** We verify that indeed LLM could have helped **improve the writing** of applications, which in turn helps with **application outcomes**.
- To talk about lender responses to the use of LLMs and its potential impact, e.g., on misallocation, we have to talk about counterfactuals. Observations of loan approvals are selected, so we introduce **Heckman for a multimodal transformer** (a potential methodological contribution).
- First paper to analyze a **counterfactual equilibrium** in a data-driven manner that is closer to reality.

Data and Sample

Prosper.com – first and largest P2P lending platform in the U.S.

All loan requests posted between April 2007 and October 2008.

Features:

- **Loan request** (amount, max interest rate, etc.)
- **Credit profile** (credit level, debt-income ratio, etc.)
- **Description** in free text

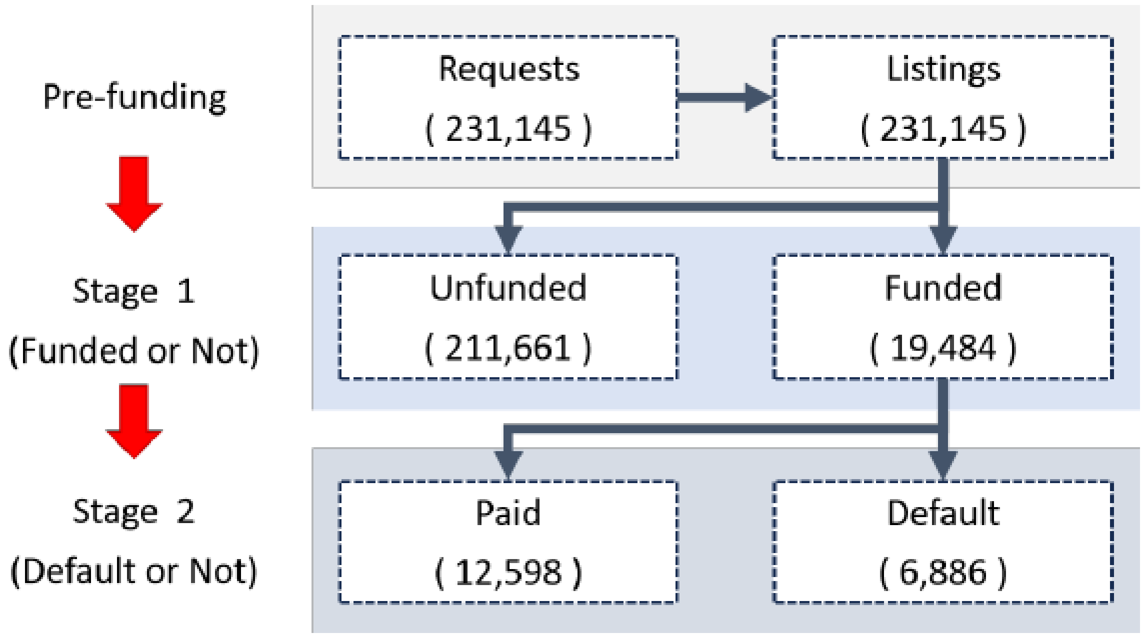

Figure 1: Life Cycle of Online Lending

Example Listing

Loan request

Credit profile

Description
(Free Text)

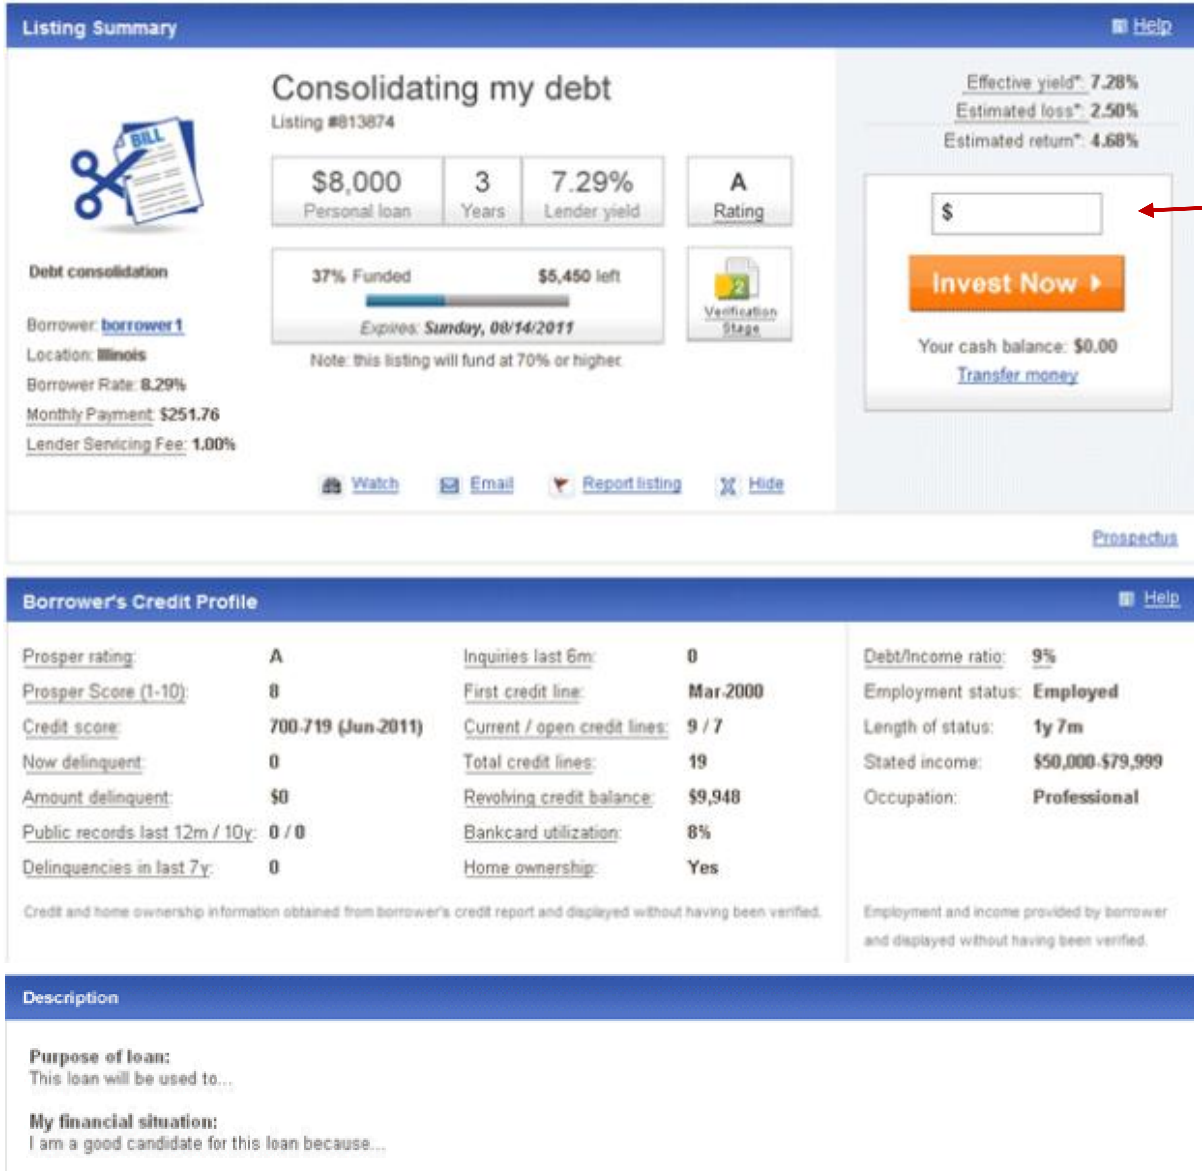

Lender Decision

A Basic ChatGPT Prompt

You are a practitioner who has been in the lending business for many years and is well versed in the matters of applying for loans.

Role Setting

Please help me rewrite the following narrative in a loan application.

Task Setting

'''
Purpose of loan: This loan will be used to pay off an existing high interest loan and my student loan
My financial situation: I am a good candidate for this loan because I do make my monthly payments on time, and all my accounts are current. My utility accounts are in good standing. I am not in a financial crisis, I am just trying to pave the way for an easier future
Monthly net income: $1920
Monthly expenses: $1550
'''

Original Text

Make the reason compelling, and make myself sound trustworthy, so lenders are more willing to fund it. Please try to use advanced vocabulary, use a combination of long and short sentences, correct grammatical errors and spelling errors, and keep the word count between 200-400.

Rule Setting

Human Assessment of Benefits of Using LLM

- Prolific: 10 adults in the United States
- Each evaluated 48 pieces of texts (24x2) shown in random order
  - funding status (unfunded, funded)
  - credit level (high, medium, low)
  - WQI (divided into four quantiles)
- Rate perceived writing quality or willingness to fund (1-10)

	DV: Writing Quality	DV: Willingness to Lend
Rewrite	**2.342***** (0.223)	**1.450***** (0.249)
Constant	5.283*** (0.158)	4.500*** (0.176)
Listing Fixed Effect	Yes	Yes
Rater Fixed Effect	Yes	Yes
n	240	240
R^2	0.730	0.717

LLM Modifications of Textual Features

Table 2: Textual Features Before and After Rewrite

	Original		Rewrite		Change	
	Mean	(SD)	Mean	(SD)	DIF[a]	(%)
Text Length	202.5	(132.2)	281.1	(99.44)	78.6***	(+38.8%)
Spelling Errors	0.035	(0.021)	0.009	(0.007)	-0.026***	(-74.3%)
Grammatical Errors	0.064	(0.041)	0.028	(0.023)	-0.036***	(-56.3%)
Formality	0.239	(0.058)	0.367	(0.035)	0.128***	(+53.6%)
Readability	10.45	(3.775)	15.75	(1.684)	5.30***	(+50.7%)
Tone	0.745	(0.157)	0.754	(0.144)	0.009***	(+ 1.2%)

[a] Sig. levels: * $p < .05$, ** $p < .01$, *** $p < .001$.

Composite Metric using Semantic Similarity

$$CosineSim(\boldsymbol{x}_1, \boldsymbol{x}_2) = \frac{\boldsymbol{x}_1 \cdot \boldsymbol{x}_2}{\|\boldsymbol{x}_1\| \|\boldsymbol{x}_2\|}$$

	"The dog plays in the garden"	"The new movie is so great"	"A woman watches TV"
"The new movie is awesome"	0.05	0.89	-0.05
"The cat sits outside"	0.28	0.00	0.13
"A man is playing guitar"	0.22	-0.01	-0.03

$$WQI(\boldsymbol{x}) = CosineSim(\underbrace{\boldsymbol{e}(\boldsymbol{x})}_{\text{Embedding of the original text } \boldsymbol{x}}, \underbrace{\boldsymbol{e}(\boldsymbol{x}')}_{\text{Embedding of the rewrite } \boldsymbol{x}'}) = \frac{\boldsymbol{e}(\boldsymbol{x}) \cdot \boldsymbol{e}(\boldsymbol{x}')}{\|\boldsymbol{e}(\boldsymbol{x})\| \|\boldsymbol{e}(\boldsymbol{x}')\|}$$

High-Credit Borrowers Historically Wrote Worse but Seem to Potentially Benefit More from LLMs

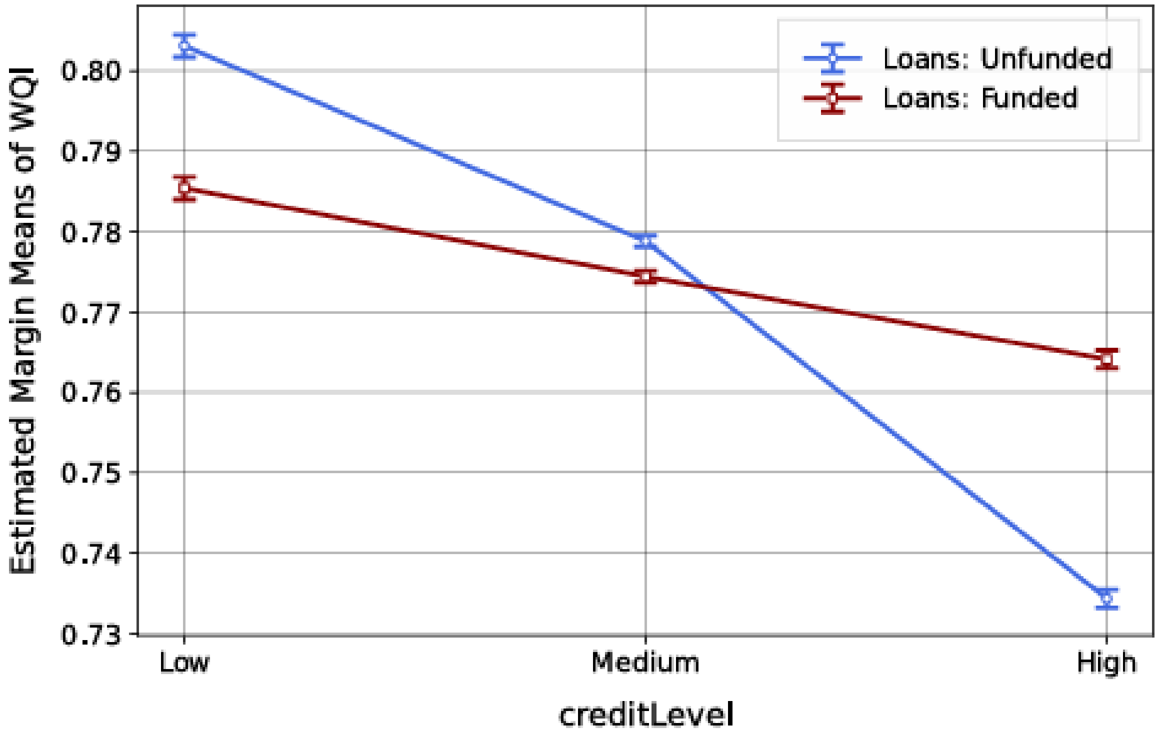

Figure 3: WQI by Credit Level Category and Funding Outcome

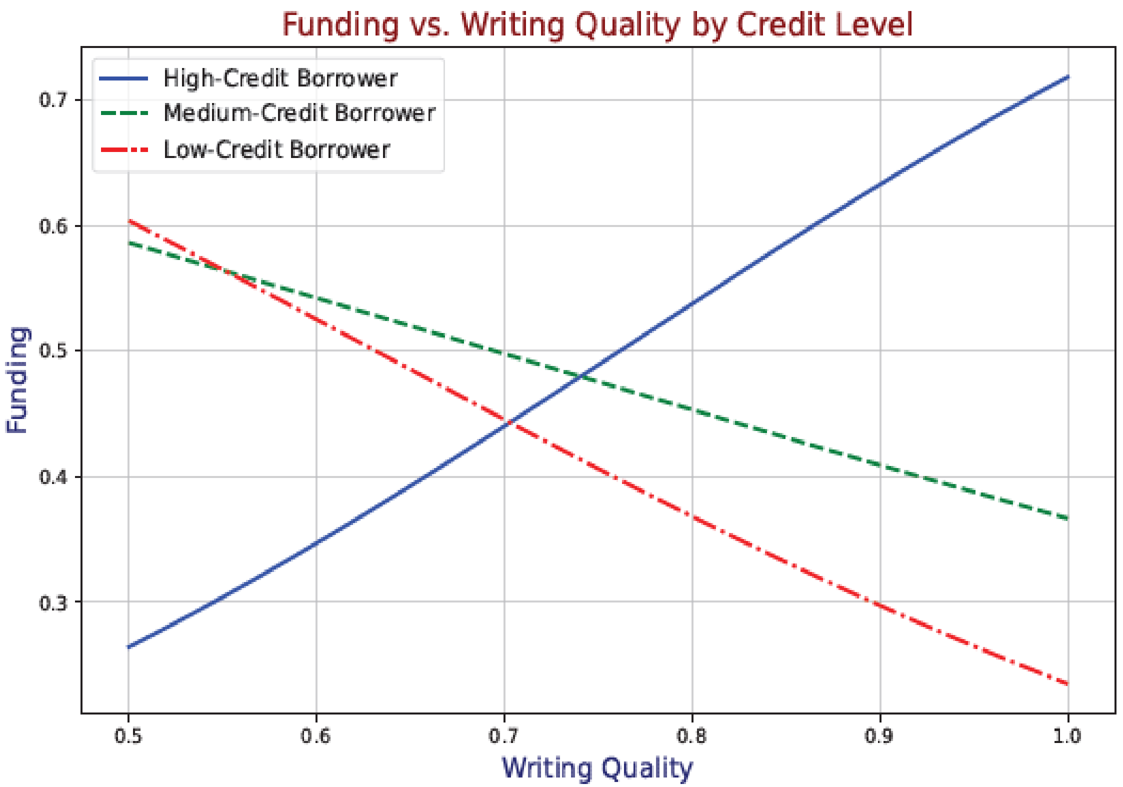

Figure 4: Heterogeneous Incremental Effect of WQI on Fundability

LLM Appears to Improve Fundability

Higher predicted probability of being funded for most loans.

Caveat: Lenders can update their decision model, especially when adoption is widespread.

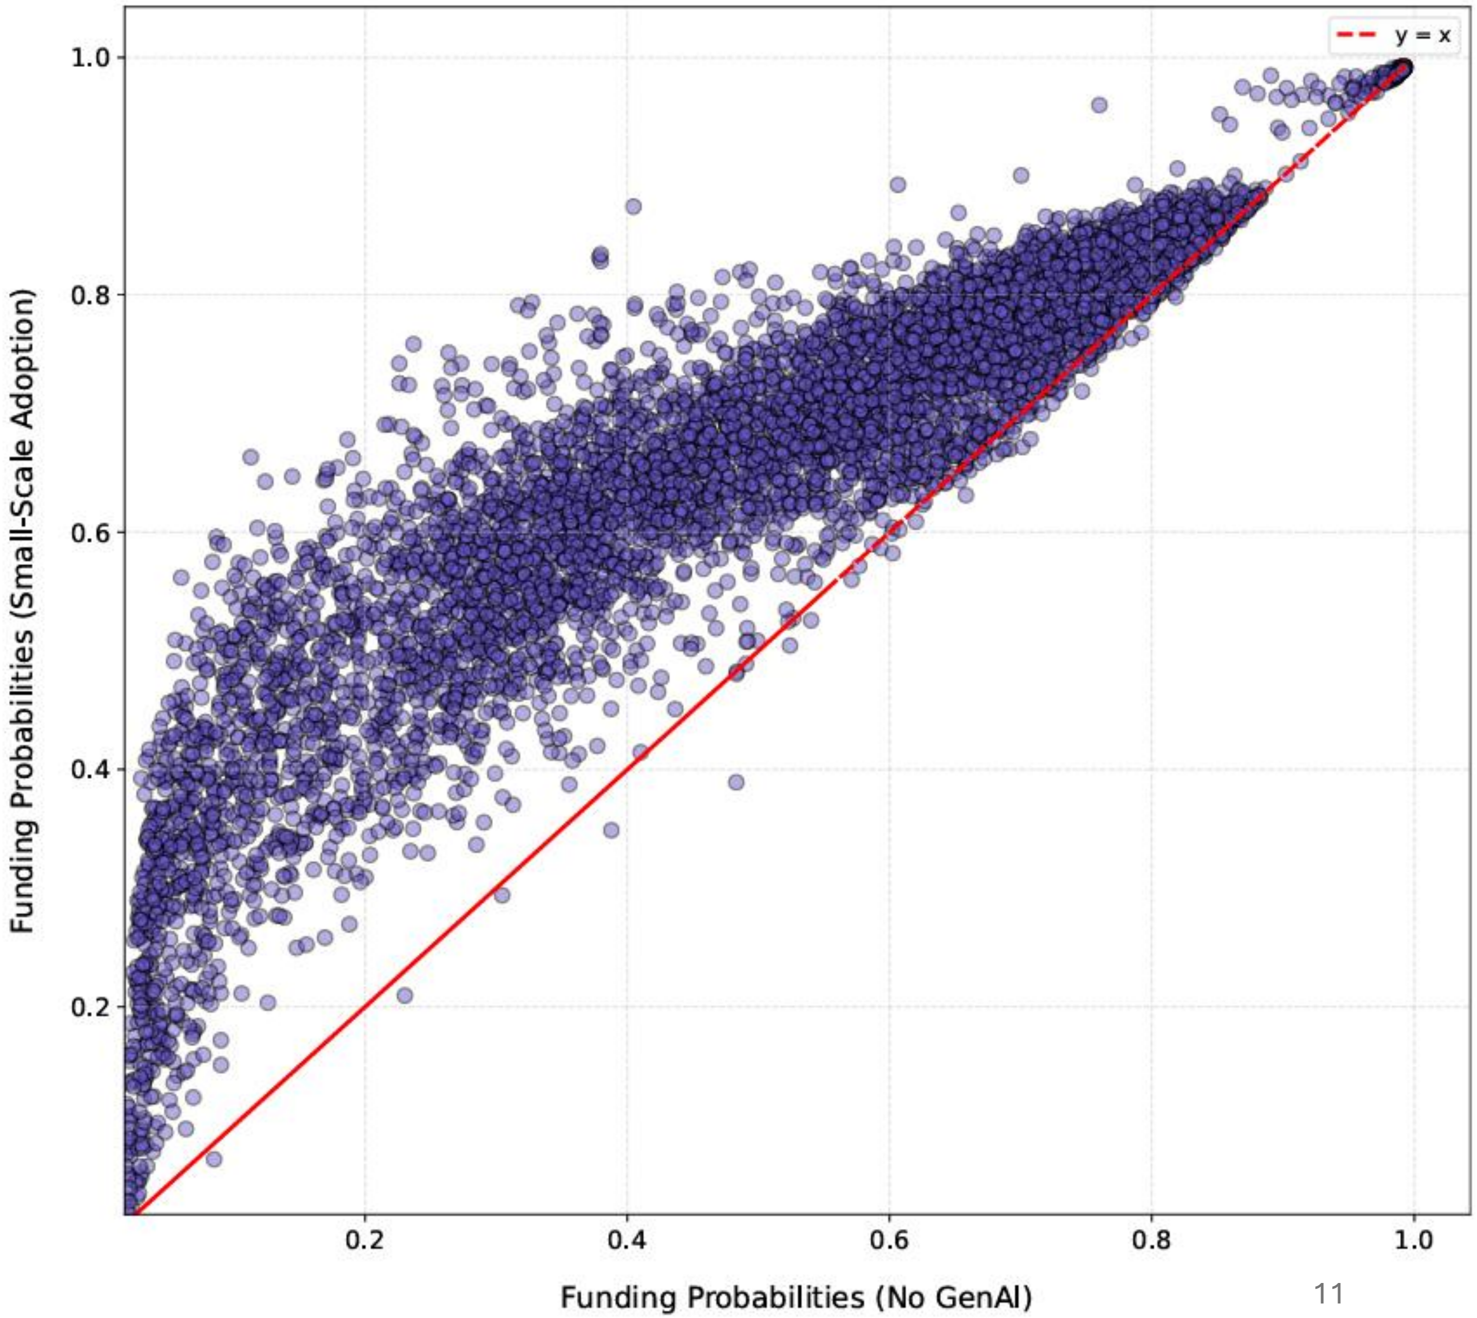

LLM Adoption & Response Scenarios

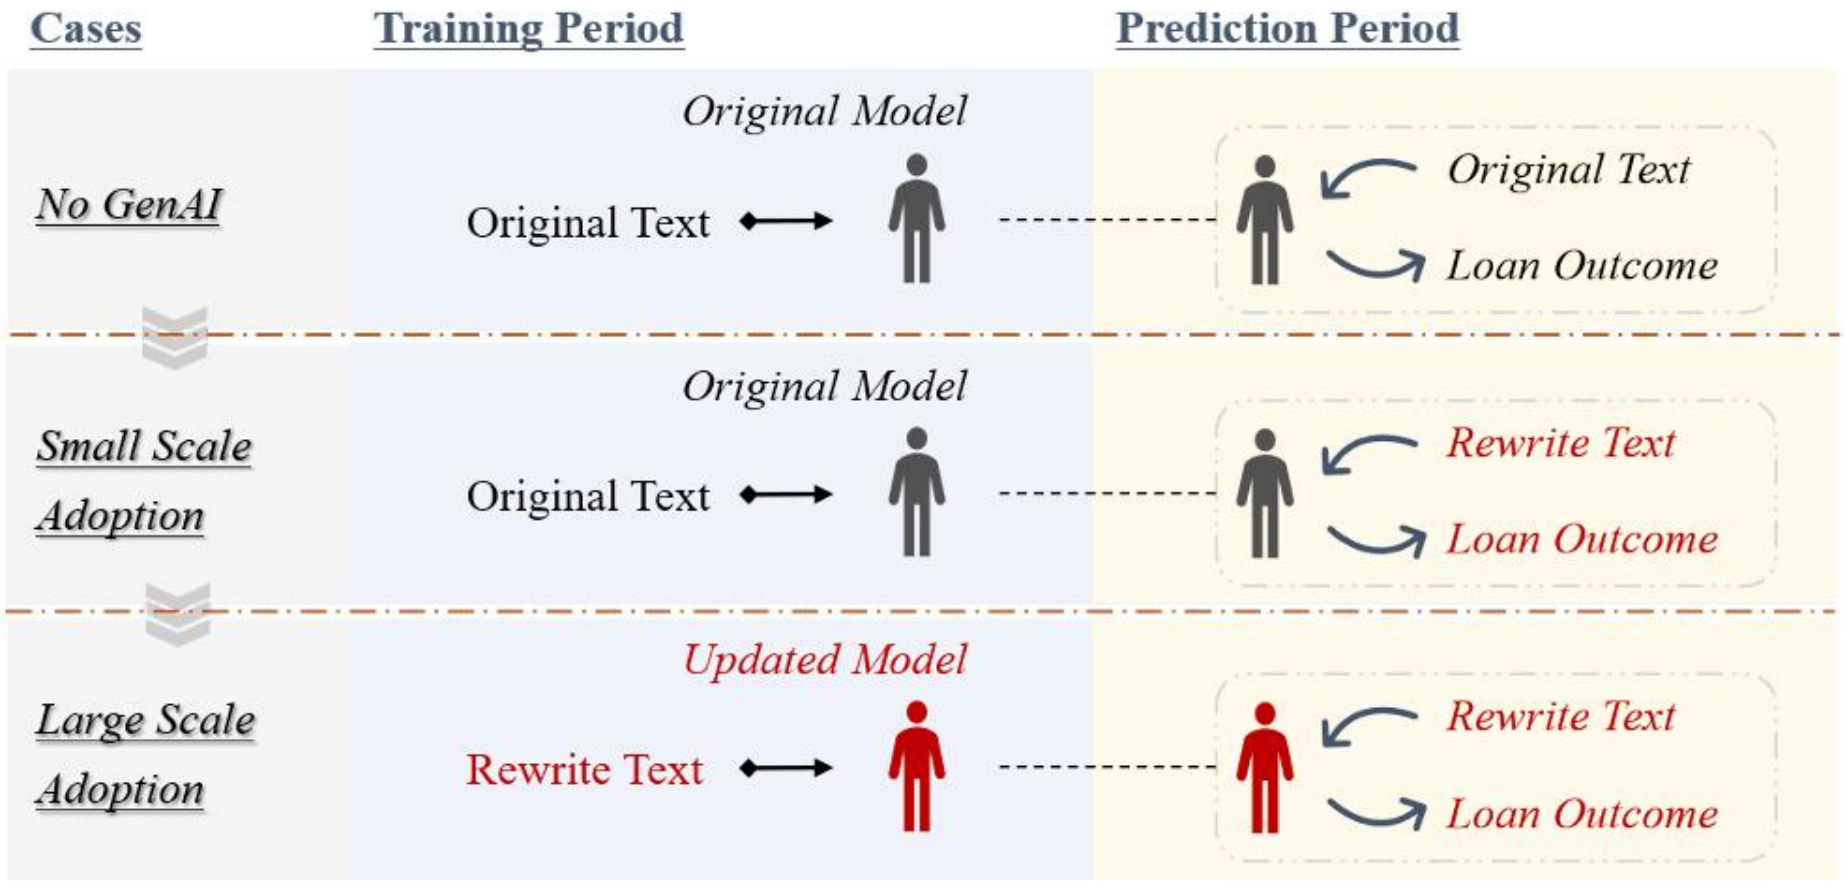

Multimodal Transformer Architecture

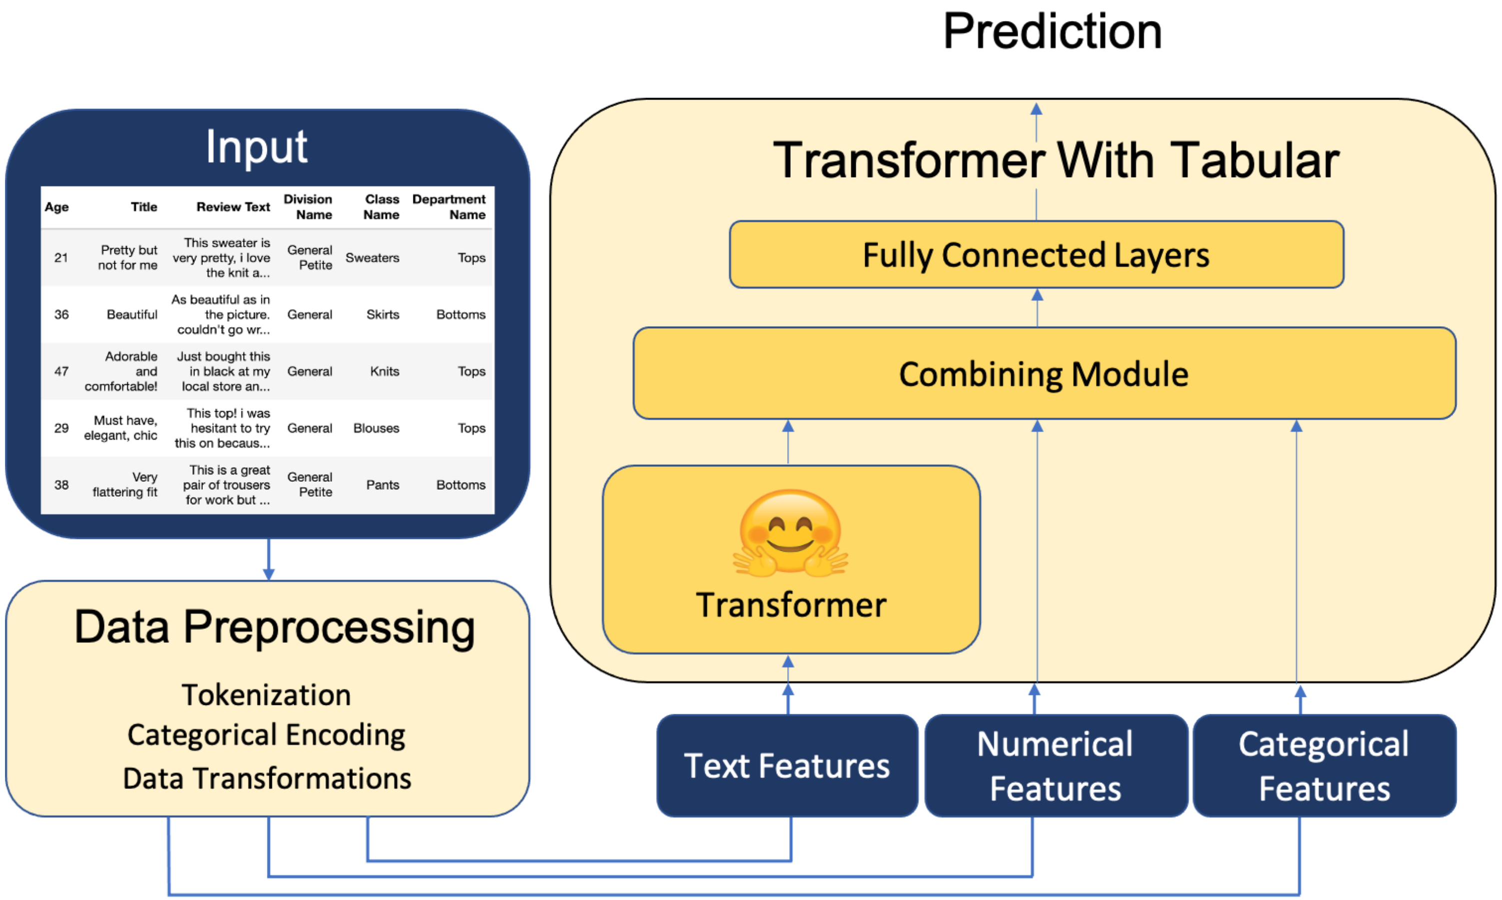

https://github.com/georgian-io/Multimodal-Toolkit

Loan Quality: Deep Heckman Correction for Unbiased Estimate for Unfunded Loans

- **Truncated observations**: if a listing is not funded, we would not know whether the loan would have been repaid if funded
- **Heckman-type correction** to the rescue
  - Stage 1: Fund ~ Listing Features,
    obtain the Inverse Mills Ratio (IMR)
    – using all listings
  - Stage 2: Repay ~ Listing Features + IMR,
    obtain adjusted model
    – using funded loans
  - Adjusted model generalizes to all listings,
    funded or not

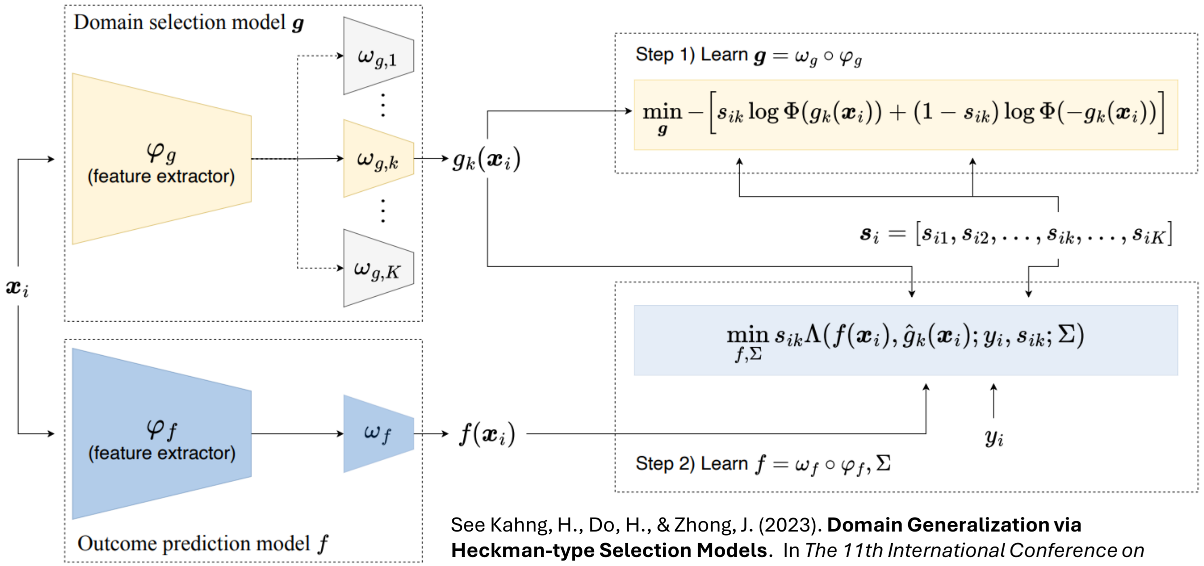

See Kahng, H., Do, H., & Zhong, J. (2023). **Domain Generalization via Heckman-type Selection Models**. In *The 11th International Conference on Learning Representations*.

- **Challenge**: Heckman Model not well-defined for deep learning
- A novel **Deep Heckman Correction + Multimodal** framework

LLMs Reduce Soft Information

$$\text{Soft Information} = U(\text{text} + \text{structured}) - U(\text{structured})$$

Table 5: Measuring Soft Information Using AUC and ROI

	AUC			ROI		
Scenario	Model 1A (Structured)[a]	Model 1B (Structured + Text)	Diff.[b] (Soft Info.)	Model 2A (Structured)[a]	Model 2B (Structured + Text)	Diff.[b] (Soft Info.)
No GenAI		74.7%	2.2%***		18.63%	2.61%***
Small-Scale Adoption	72.5%	73.8%	1.3%***	16.02%	15.68%	-0.34
Large-Scale Adoption		74.0%	1.5%***		16.78%	0.76

[a] The performance metric does not change by LLM adoption since the model does not use text.
[b] Sig. levels: * $p < .05$, ** $p < .01$, *** $p < .001$. All AUCs and ROIs are averaged across 15 replications of the ten-fold mean.

Antecedent and Consequence of Reduced Soft Information

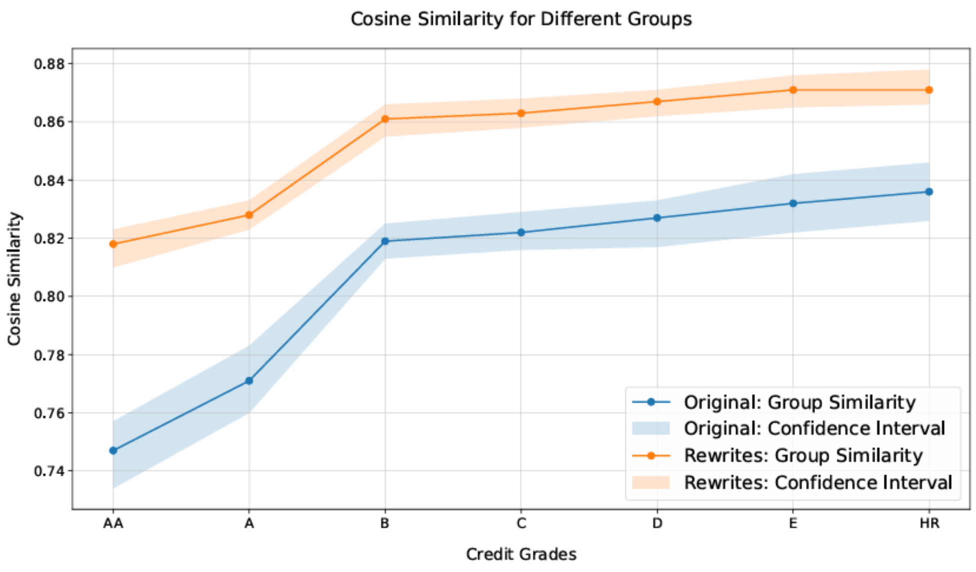

Figure 9: Text Similarity by Credit Grades Before and After Rewriting

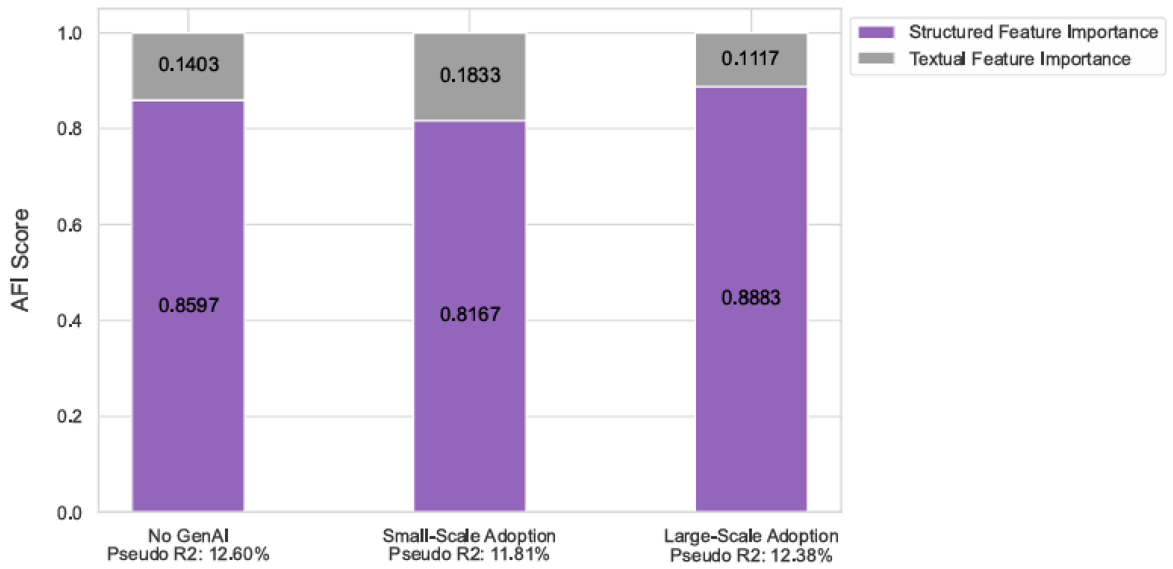

Figure 10: Aggregated Feature Importance (AFI) Scores in Different Scenarios

Impacts on Credit (Mis)allocation

Table 6: Misallocation in Different GenAI Adoption Scenarios

Scenario	FPR	DIFF.[a]	FNR	DIFF.[a]
No GenAI	17.50%		20.86%	
Small-Scale Adoption	23.12%	**+5.62%*** **	23.28%	**+2.42%*** **
Large-Scale Adoption	20.52%	**+3.02%*** **	20.28%	- 0.58%

[a] The differences are relative to the "No GenAI" scenario. Significance tests are based on 15 repetitions of the experiments.
[b] Sig. levels: * $p < .05$, ** $p < .01$, *** $p < .001$.

Why This Counterfactual Equilibrium Is Hard

GenAI adoption changes borrower behavior, lender inference, and the data-generating process.

Frozen counterfactual

"What if text changes, but the lender model stays fixed?"

Historical loan data

↓

Learned lender model θ

↓

Counterfactual prediction

Lucas critique: θ is no longer the right environment.

Equilibrium counterfactual

Solve jointly for behavior, data, and the learned environment.

Borrowers adapt

Applications change

Fixed point

Payoffs update

Lenders retrain

Counterfactual analysis becomes a search for a self-consistent market state.

Source: Cong (2026), Economic World Models and Data-Driven Generative Equilibria.

Economic World Models → DDGE (Cong, 2025)

EWM: a learned economic system

Agents, environments, and interactions can be theory-specified, learned, or hybrid.

Agents
policies, objectives, beliefs

Environment
states, transitions, constraints

Interactions
markets, institutions, information

Economics adds meaning:
interventions, feasibility, welfare.

DDGE: equilibrium for data-driven EWMs

$$\theta^* = L(D(\pi^*, \mu^*, \theta^*; i))$$

Learned environment = retraining on equilibrium-generated data

Strategy π^*

Beliefs μ^*

Closure

Learning L

Data D

Core idea: equilibrium discipline is applied to the learning loop itself.

Our DDGE: Generative Borrowers + Discriminative Lenders

1 Initialize adoption + θ^0 → **2 Generate** LLM applications → **3 Retrain** lender model → **4 Update** benefits/costs → **5 Stop when** residual ≈ 0

Generative side: borrowers

High-level action: use GenAI or not

Credit type + cost → LLM rewrite module → New text data

generated market data

Marginal borrower adopts when expected fundability gain exceeds GenAI cost.

Discriminative side: lenders

Decision rule adapts to the new data environment

Text + credit profile → Multimodal Heckman model → Funding + ROI

Equilibrium object: borrower adoption rule and lender model are mutually consistent.

Contribution: first empirical, data-driven counterfactual equilibrium analysis of GenAI adoption in online credit markets.

Source: Cong (2026), Economic World Models and Data-Driven Generative Equilibria.

Data-Driven Generative Equilibrium (Cong, 2025)

- Borrowers as counterfactual data generator

$$\max_{a_i^H} \mathbb{E}\left[P_{\text{funding}}(a_i^H, a_i^L, \sigma_j^*)\right]$$

- Lenders as data-driven decision-makers

$$\sigma_j^* = \arg\max_{\sigma_j} \mathbb{E}_{\mu_j}[u_j(\sigma_j(D^C), a^H, a^L, \omega)]$$

- Marginal borrower's adoption condition:

$$\mathbb{E}[\text{Fundability}(\text{ChatGPT}, \sigma^*)] - \mathbb{E}[\text{Fundability}(\text{No ChatGPT}, \sigma^*)] = C_{\text{ChatGPT}}$$

Search and Verification of Equilibrium

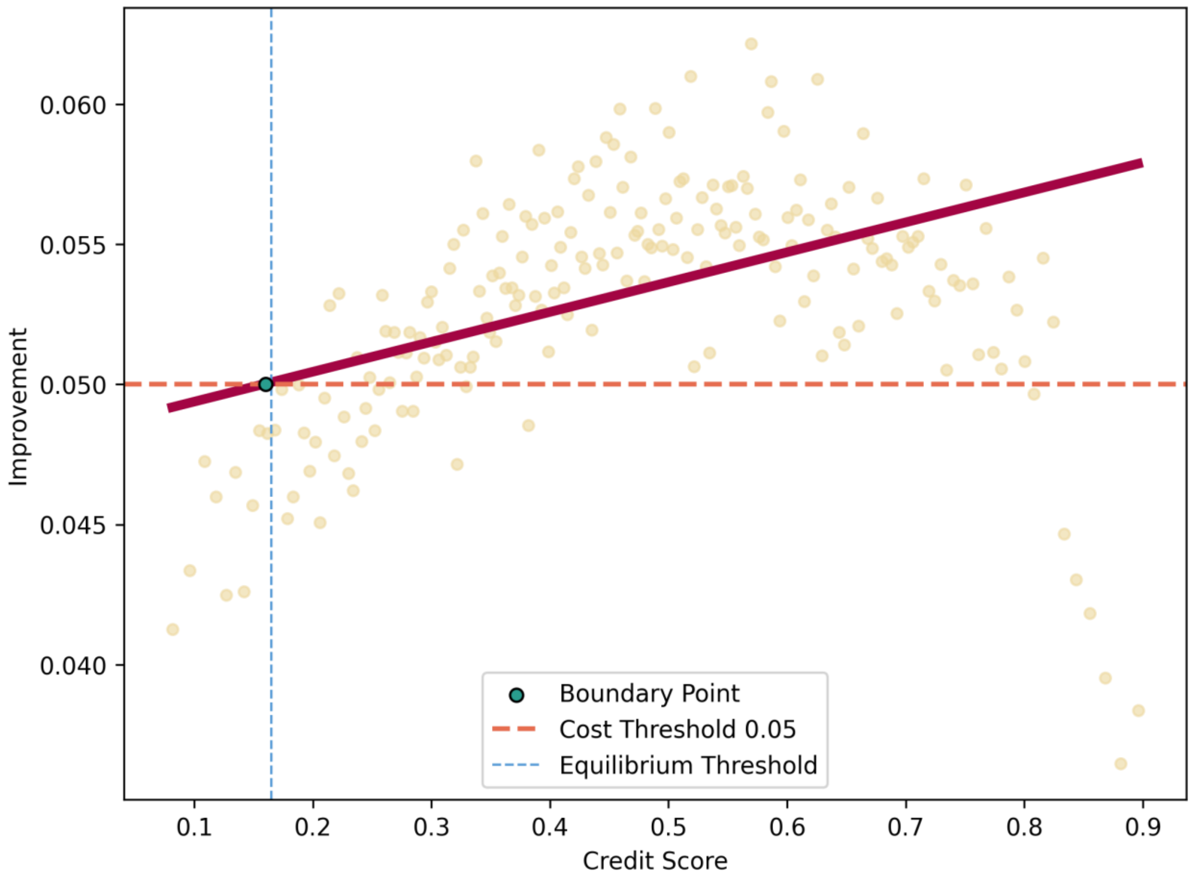

- A linear trend is fitted between credit score and improvement.
- Results under Adoption Rate: 0.95
- Top 95% credit-ranked borrowers and only them obtain positive net benefits
- The boundary borrower satisfies:

  Improvement ≈ Cost Threshold

Takeaways

- LLM can potentially help borrowers (writing quality, funding probability), as suggested by historical observations
- LLM adoption reduces soft information → greater misallocation
- Lenders may respond by focusing more on hard information → misallocation mitigated but not eliminated
- Counterfactual equilibrium analyses: low adoption cost → unique linear equilibrium with full adoption

Additional Work Handling Selection Bias in Deep Credit Risk Prediction: Heckman-Style Framework

Deep Heckman Framework

Shared-Branch:

- shared factors drive both funding and repayment ;
- branch-specific factors capture unique process effects.

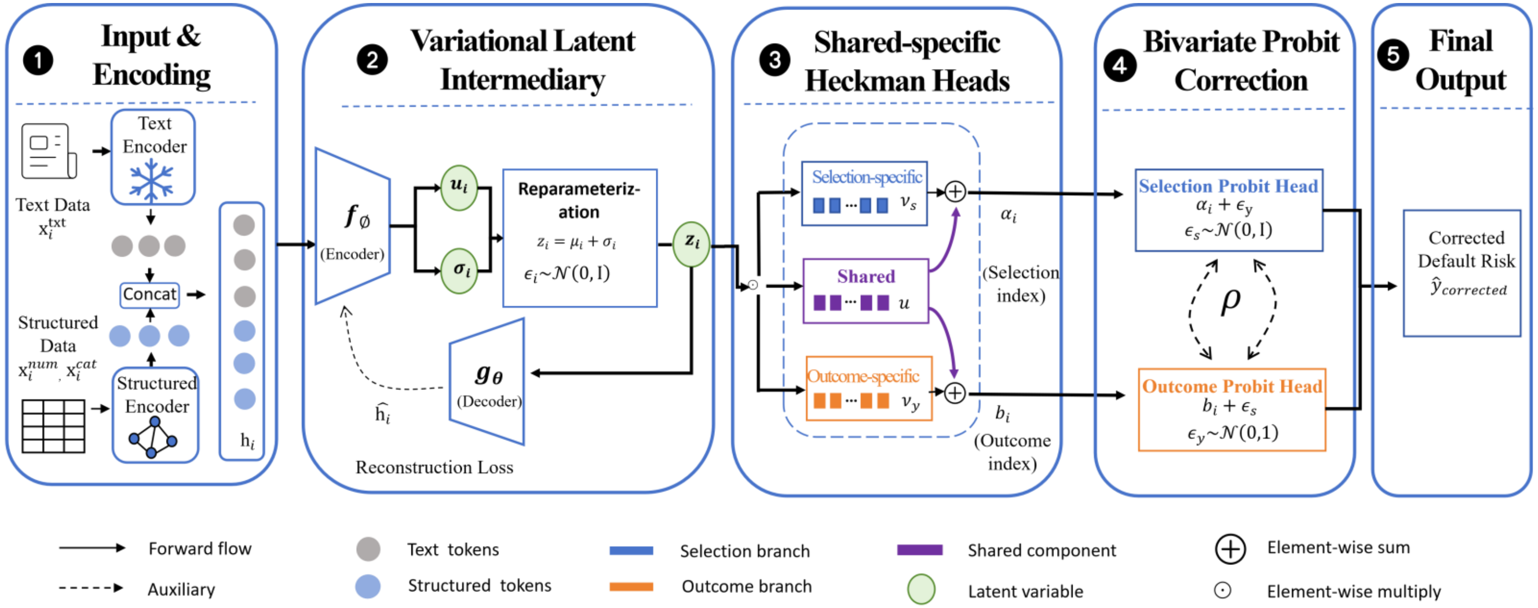

Observational states usage:

- **Unfunded applications**: inform selection
- **Completed loans**: inform joint selection–outcome likelihood
- **Ongoing loans**: improve selection calibration via Maturity-Aware Posterior Supervision (MAPS) (soft targets instead of hard labels)

Jiaxin Dou, Lin William Cong, Wenjun Zhou, Yanhong Guo. Handling Selection Bias in Deep Credit Risk Prediction: A Heckman-Style Framework.