Did the Rate Pay for the Risk? Consumer Loan Coupons vs Realised Returns, 2007–2015

Gaurav Gupta · Research note · 27 October 2026 · 618,731 LendingClub 36-month loans ($7.7bn), issued Jun 2007–Dec 2015, all run off by the Mar 2019 data cut

7.0% vs 2.4%
Realised return a year, grade B vs grade G (coupon 10.8% vs 25.3%)

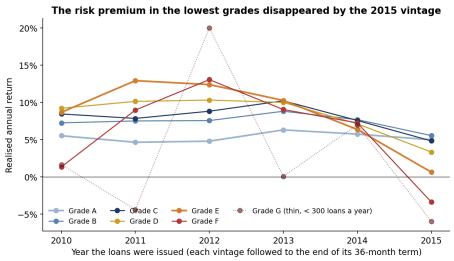
−3.4%
Grade F loans issued in 2015; 9.0% a year in 2010–2013 vintages

0.7 pts
Average miss on lifetime default from a month-18 read (pricing assumption: 2.0)

+39 bp
Out-of-sample return from skipping riskiest 20% in each sub-grade (2015)

Grade	Loans	Coupon	$ defaulted	Net loss	Realised return	2015 vintage
A	142,925	7.2%	5.2%	2.4%	5.3%	5.0%
B	211,692	10.8%	11.0%	5.3%	7.0%	5.6%
C	161,843	13.9%	17.9%	9.3%	6.8%	4.9%
D	76,074	17.2%	24.0%	13.4%	6.3%	3.4%
E	21,410	19.8%	29.8%	17.5%	4.8%	0.7%
F	4,265	23.5%	35.5%	21.9%	4.1%	−3.4%
G	522	25.3%	38.4%	24.7%	2.4%	−6.0% (248 loans)

Realised return = IRR of pooled monthly cash flows after a 1% servicing fee; recoveries net of collection fees six months after last payment. 2010–2013 vintages: D 10.1%, E 11.0%. Pricing assumption (sub-grade mix × run-off history) for 2014Q4–2015Q4: 11.2% default; actual 14.4%; month-12 read 13.6%, month-18 read 14.8%.



Realised annual return by grade and year of issue: the premium in the lowest grades disappeared by 2015.

SOURCES

- LendingClub issued loans 2007–2018Q4, Hugging Face copy codesignal/lending-club-loan-accepted (CC0), sha256-pinned
- 3-year Treasury: US Treasury via defeatbeta/yahoo-finance-data (ODC-BY)
- Emekter et al. (2015); Serrano-Cinca et al. (2015); Breeden (2007); Mack (1993); Ke et al. (2017); Siddiqi (2006)

METHOD

- Complete = 36-month term ended before the data cut; default = charged off, default or 31–120 days late
- Vintage curves: share of $ that stopped paying by month on book; early read = defaults at month k × development factor from vintages fully run off by that date
- Screen: gradient boosting on application fields only (tested), train 2010–13, early-stop 2014, test once on 2015; random-skip control; 200-draw bootstrap

KEY CALCULATIONS

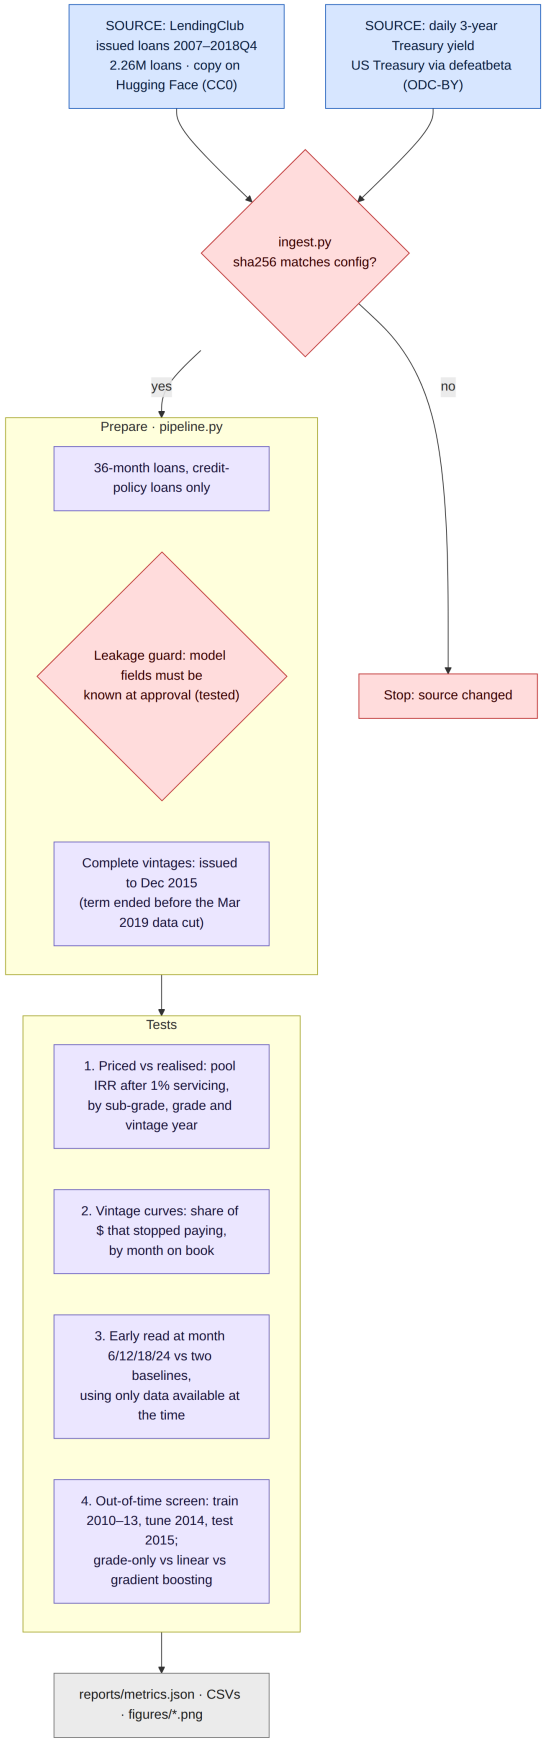
- Early-read miss: last run-off 2.8, pricing mix 2.0, month 6 3.7, 12 1.9, 18 0.7, 24 0.3 pts
- 2015 test AUC 0.678 / 0.678 / 0.685 (grade / linear / boosted); skip 20%: 4.79% → 5.05% / 5.19%; random 4.80%
- Predicted 12.7% vs actual 14.9% default; stability index 0.006
- ≈ $17M extra net loss per $1bn, 2014Q4–15Q4 (illustrative)

RECOMMENDATIONS

- Set grade limits on realised, risk-adjusted return; cap the lowest grades
- Read every vintage at months 12 and 18 vs plan; review price or limits if projected loss runs >1 pt above
- Within-grade screening as an overlay (~40 bp), not a substitute for price
- Monitor calibration against outcomes monthly; stability indices miss outcome drift

Independent research on public data; not investment or lending advice; not deployed. Code: github.com/guptagaurav-info/consumer-loan-pricing · Full write-up: guptagaurav.info/articles/did-the-rate-pay-for-the-risk

Pipeline


SOURCE: LendingClub
issued loans 2007–2018Q4
2.26M loans · copy on
Hugging Face (CC0)
SOURCE: daily 3-year
Treasury yield
US Treasury via defeatbeta
(ODC-BY)
ingest.py
sha256 matches config?
yes
no
Prepare · pipeline.py
36-month loans, credit-
policy loans only
Leakage guard: model
fields must be
known at approval (tested)
Complete vintages: issued
to Dec 2015
(term ended before the Mar
2019 data cut)
Stop: source changed
Tests
1. Priced vs realised: pool
IRR after 1% servicing,
by sub-grade, grade and
vintage year
2. Vintage curves: share of
$ that stopped paying,
by month on book
3. Early read at month
6/12/18/24 vs two
baselines,
using only data available at
the time
4. Out-of-time screen: train
2010–13, tune 2014, test
2015;
grade-only vs linear vs
gradient boosting
reports/metrics.json · CSVs
· figures/*.png

How a credit committee would run the purchase programme

