

DATING · BOOST-PRICING PERSONALISATION

Velvet

Boost-pricing personalisation

CUSTOMER

Series-B dating app · ~\$26M ARR

PERIOD

Mar – May 2026 (90 days)

SAMPLE

318,940 boost-eligible users

EXPERIMENT AUDITED

Boost price · fixed \$7.99 (control) vs fixed \$4.99 (variant)

DATING

Velvet · Boost-pricing personalisation

1 · WHAT THEIR TEAM REPORTED

METRIC	\$7.99 (CONTROL)	\$4.99 (VARIANT)
Boost-conversion rate	6.8%	7.8%
Boost revenue / eligible user	\$0.54	\$0.39

Team report: "lower price wins on conversion." Shipped \$4.99 globally. Boost revenue fell against forecast – attributed to 'low elasticity.'

2 · OUR RE-ANALYSIS · DOUBLY-ROBUST + PER-SEGMENT CATE

COHORT	DR CONV. LIFT	DR REV / USER	95% CI (REV)	ESS	VERDICT	\$ IMPACT / YR
All boost-eligible	+13.8% rel.	-\$0.15 / user	[-0.21, -0.09]	0.59	conversion up, revenue down	–
Light engagement (1-3)	+24.1% rel.	+\$0.07 / user	[+0.02, +0.12]	0.46	positive on both	+\$0.2M (keep \$4.99)
Mid engagement (4-7)	+12.0% rel.	-\$0.18 / user	[-0.25, -0.11]	0.52	conv. up, rev. down	+\$0.3M (back to \$7.99)
Heavy engagement (8-10)	+3.2% rel.	-\$0.42 / user	[-0.51, -0.33]	0.48	money left on the table at \$7.99	+\$0.5M (raise price)
Recently-matched (last 7d)	+6.1% rel.	-\$0.21 / user	[-0.29, -0.13]	0.41	conv. up, rev. down	+\$0.1M (back to \$7.99)

3 · THE HIDDEN COHORT EFFECT

The heavy-engagement deciles **would pay more**. At \$4.99 they cost \$0.42 / user of boost revenue per cohort member. Across the ~87,000 users in those deciles over 90 days, that's roughly **\$330K of foregone revenue** in one quarter.

4 · WHAT WE'D RECOMMEND

Bandit-personalise boost pricing on (engagement decile × days-since-last-match × current online-cohort density).

PROJECTED ANNUALISED IMPACT
+14 % boost revenue · +\$1.1M / yr

Estimates use 1,000-bootstrap doubly-robust evaluation against the customer's logged data with propensities reconstructed from the experiment configuration. ESS < n/10 cells are flagged overlap-limited and not used in the recommendation. The customer can reproduce these numbers on their own logs using `offpolicy.py` (MIT-licensed). Numbers in this report are fictive, generated as a worked example – not derived from real customer data.