

QUANT RUSH · FINAL SPRINT

THE LAST DIG

Ensembling, final picks, and Q&A.

TUE NOV 17

5 DAYS TO THE ASSAY

Engraving: “San Francisco Harbor during the gold craze,” from J. P. Young, San Francisco (1912) · Public domain

TODAY

Five days to the Assay

1

BLEND

Combine what you already have, the right way.

2

PICK

Choose your two final submissions on purpose.

3

SHIP

Everything to check before Sunday night.

Then Q&A: bring your notebook and your questions.

ENSEMBLING

Two ideas beat one

Scored with the competition metric on the last 1,080 training eras:

Prediction	Score
`-ret_00` alone	0.0576
`-ret_01` alone	0.0299
both, rank-averaged	0.0633

The blend beats even the stronger rule on its own.

WHY IT WORKS

The two rules are almost unrelated: their correlation is −0.06. When one is wrong, the other often isn't, so part of the noise cancels.

HOW TO BLEND

Rank within the era, then average

```
def era_rank(df, col):
    return df.groupby("era")[col].rank(pct=True)

test["prediction"] = (era_rank(test, "pred_a") + era_rank(test, "pred_b")) / 2
```

WHY RANK

Models output different scales. Ranks put them on one.

WHY PER ERA

The metric only compares assets within the same era.

WHY EQUAL

Equal weights are a strong, safe default.

NOT EVERY BLEND HELPS

Blend different, not more

Starting from -ret_00 (0.0576), add…	Score	Lesson
a near-copy of itself (correlation 0.96)	0.0573	Copies add nothing.
-ret_01 (correlation −0.06)	0.0633	Different ideas add.
-ret_01, -taker_00 and -range_00	0.0562	Weak, overlapping signals dilute.
-ret_01, with -ret_00 counted twice	0.0655	+0.002 is inside the ±0.003 noise.

Check every blend on recent eras, and don’t fine-tune the weights.

THE PUBLIC BOARD

Trust your validation, not the leaderboard.

The public board scores about 30% of the test set. Gaps of a few thousandths there are mostly luck.

Tuning to it is how strong teams fall on the private board. That is what Fool's Gold is for.

YOUR TWO FINALS

Pick two, on purpose

Final 1

YOUR BEST VALIDATED

The submission that scores best on the latest training eras, not the one with the best public score.

Final 2

A DIFFERENT, STURDY BET

An equal-weight blend, or your best idea from a different family. If one stumbles, the other may hold.

Pick none, and Kaggle uses your best public scores: possibly your luckiest, not your best.

BEFORE SUNDAY

The last-dig checklist

1 SELECT YOUR FINALS

Mark 2 submissions on Kaggle before Sun Nov 22, 11:59 p.m. PT.

2 KEEP THE CODE

Fixed seeds, and it reruns with only the competition files. Winners send it within 5 days.

3 CHECK YOUR REGISTRATION

Your Kaggle username must be on it. Missing? Email qfec.smc@gmail.com.

4 SPREAD YOUR SUBMISSIONS

5 a day per team, reset at 4 p.m. PT. Don’t save them all for Sunday.

WHAT'S LEFT

The last three dates

Sun Nov 22

THE ASSAY

Submissions close at 11:59 p.m. PT. The private board appears after.

Wed Nov 25

MEMOS DUE

The Prospector's Report: email it to qfec.smc@gmail.com by 11:59 p.m. PT.

Tue Dec 1

PAYDIRT

Winners announced, and top teams walk through their solutions.

Red outlines mark the hard deadlines.

Q&A

Dig in.

Bring your notebook: we'll help you blend, validate, and pick your finals.

DISCORD

Questions any time until Sunday:
discord.gg/MUPdvSAV4M

EMAIL

Registration and anything else:
qfec.smc@gmail.com

Good luck at the Assay.