

Talking to Digital Twins

Selective Disclosure and Belief Measurement in Financial Social Media

Boone Bowles

Raymond Duch

Sorin Sorescu



TEXAS A&M UNIVERSITY

Mays Business School



Nuffield
College
UNIVERSITY OF OXFORD



TEXAS A&M UNIVERSITY

Mays Business School

AI in Business Conference

The Ohio State University

October 15, 2026

What does silence tell us?

Stock A



Stock B



Bullish? Bearish? Or simply nothing to say?

Markets run on information and beliefs.
Even influential investors leave some views unspoken.

Can we make the silent region talk?

The silent region is large

81 influencers 429 large-cap stocks 50 trading dates

December 15, 2025–March 16, 2026




Only 3,251 of 21,410 stock-dates have a public recommendation from **any** monitored account that day.

A structured daily LLM protocol makes the silent region measurable.

How do we make the silent region measurable?


1. Public inputs



Data for each account
+ relevant web context

How do we make the silent region measurable?

1. Public inputs



Data for each account
+ relevant web context

2. LLM interview




“Digital Boone”

Answers as this account
using the day's context

How do we make the silent region measurable?

1. Public inputs



Data for each account
+ relevant web context

2. LLM interview



"Digital Boone"

Answers as this account
using the day's context

3. Repeat daily



Schematic schedule

Every account and stock,
every scheduled day

Public-persona belief proxies make the silent region measurable.

The question and the recorded answer

Stock question

Based on your general knowledge of market conditions, indicate whether you would BUY, HOLD, or SELL (SHORT SELL) the stock ticker:
[COMPANY NAME (TICKER)].



Example response

question:
[COMPANY NAME (TICKER)]

recommendation: 75

speculation: 30

Validation: Do twin responses match public behavior?

1. When they post

91.5%

Buy/sell agreement

Same account, stock, and date

Meaningful, low-speculation
interview responses

Validation: Do twin responses match public behavior?

1. When they post


91.5%

Buy/sell agreement

Same account, stock, and date

Meaningful, low-speculation
interview responses

2. Before they post



Validation: Do twin responses match public behavior?

1. When they post


91.5%

Buy/sell agreement

Same account, stock, and date

Meaningful, low-speculation
interview responses

2. Before they post



3. Identify the account

Held-out responses;
top-five match among 81 accounts

35.2%

Observed

vs.


6.2%

Chance

How do answers become a stock-level signal?

Stock A · Tuesday

Ten digital-twin responses



$$\text{Net Buy Share} = \frac{\# \text{ buys} - \# \text{ sells}}{\text{all parsed responses}} = \frac{2 - 1}{10} = 10\%$$


Aggregate belief proxies, then test their ability to predict returns.

Responses are time-stamped; returns begin after the assigned trading-day close.

Twin interviews contain a return signal

Each day, sort stocks into ten groups by Net Buy Share.

Compare their returns over the next ten trading days, in excess of the S&P 500.




Most bullish minus most bearish: 153.3 bp over just ten trading days.

Positive estimates across four return tests

1. Fama–MacBeth


Daily cross-stock comparisons; ten trading days.



Positive estimates across four return tests

2. Add past-return and volatility controls


Daily cross-stock comparisons; ten trading days.



Positive estimates across four return tests

3. Non-overlapping five-day blocks

Five-day Wed–Tue blocks.




All three positive and significant at 10% (*).

Positive estimates across four return tests

4. Staggered vintages

Ten alignments; non-overlapping ten-day windows within each.




Unrestricted whole-history placebo: *** $p < .01$.

Shared alignments require resampling inference.

Industry does not explain the whole signal


Reassign complete signal histories among stocks in the same industry.




10,000 whole-history reassignments within FF12 industries.

Where is the return signal?


Full stock-day sample
21,410 stock-days



Where is the return signal?




Where is the return signal?



Two stock-day subsets; three signal comparisons.

Where is the return signal?


Net Buy Share: +10 percentage points



Return estimates are largest in the silent region.

Where is the return signal?


Conviction-weighted meaningful Tilt: +1 unit



Return estimates are largest in the silent region.

Disagreement and uncertainty carry different information

Ten-day outcomes per +0.10 increase in each 0–1 measure



Disagreement: lower returns, higher volatility.

Uncertainty: lower volatility; its negative return estimate is imprecise.

Macro disagreement is followed by lower market returns

Ask each twin the same macro questions

What is the chance of
a U.S. recession
in the next year?

Will business conditions
improve over the
next six months?

Will the stock market
rise over the
next six months?



Net Sentiment


How bullish is the panel
on balance?

Disagreement

How far apart are
the twins' views?

Macro disagreement is followed by lower market returns


Ten-day S&P 500 return association per +1 SD in each measure



More disagreement is followed by lower market returns.

Net sentiment's return estimate is positive but imprecise.

Coloring in the silent region



The added coverage contains return information.

Public disclosure gives us
only part of the picture.

