(Mis)alignment Between Stance in Social Media Data and Public Opinion Surveys

Sarah Shugars
Assistant Professor, Rutgers University
they/them/theirs
sarah.shugars@rutgers.edu

Based on work with Kenneth Joseph, Ryan Gallagher, Jon Green, Alexi Quintana Mathé, Zijian An, David Lazer

RUTGERS

Prepared for Digisurvor Workshop
February 14, 2025

Motivation

Does social media capture "public opinion"?

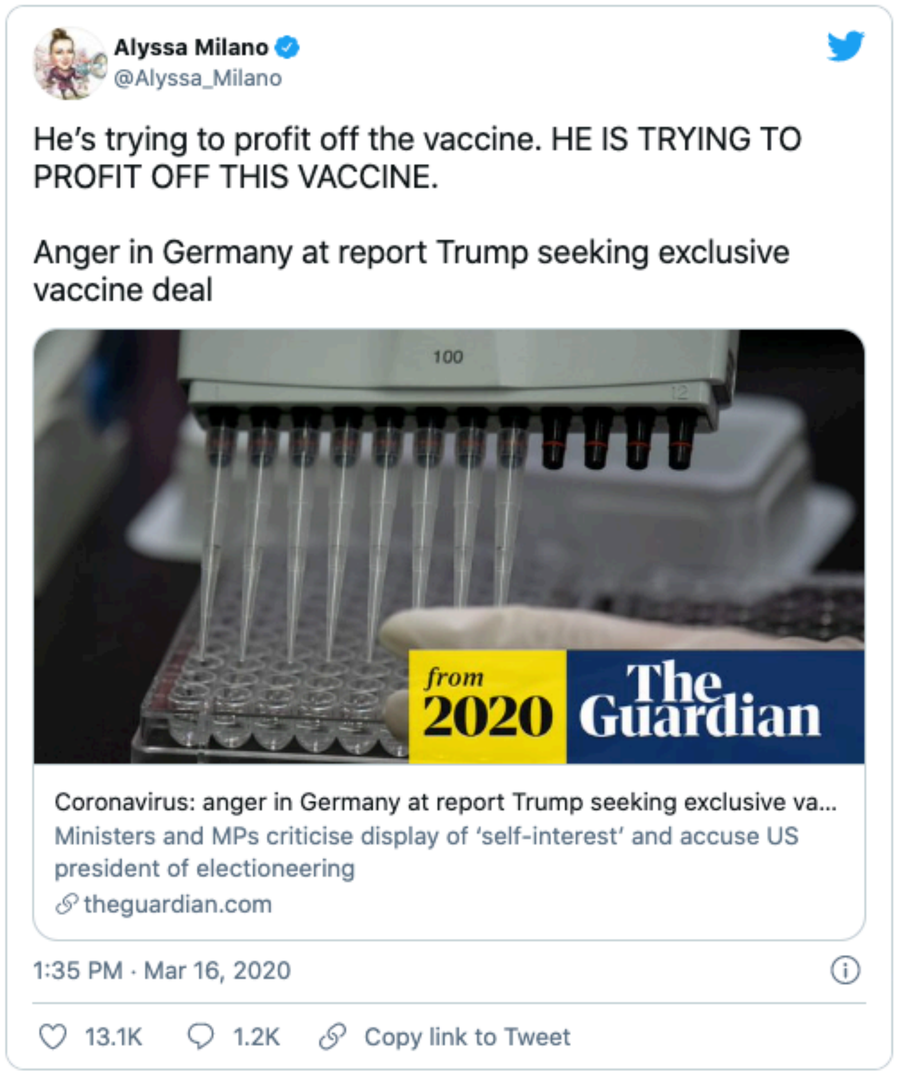

Social Media v. Surveys

Social Media v. Surveys

Strengths

Social Media v. Surveys

Strengths

- Real-time updates

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Weaknesses

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Weaknesses

- Can be slow to field

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Weaknesses

- Can be slow to field
- Cost $$

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Strengths

Weaknesses

- Can be slow to field
- Cost $$

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Strengths

- Representative sample

Weaknesses

- Can be slow to field
- Cost $$

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Strengths

- Representative sample
- Tailor specific questions

Weaknesses

- Can be slow to field
- Cost $$

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Weaknesses

Strengths

- Representative sample
- Tailor specific questions

Weaknesses

- Can be slow to field
- Cost $$

Social Media v. Surveys

Strengths

- Real-time updates
- “Publicly” available / low cost (?)

Weaknesses

- You get what you get in terms of respondents and topics

Strengths

- Representative sample
- Tailor specific questions

Weaknesses

- Can be slow to field
- Cost $$

Matching social media & survey data

Survey Data

THE COVID STATES PROJECT

https://covidstates.org

Matching social media & survey data

Survey Data

- Online survey running since April of 2020

THE COVID STATES PROJECT

https://covidstates.org

Matching social media & survey data

Survey Data

THE COVID STATES PROJECT
https://covidstates.org

- Online survey running since April of 2020
- Each wave includes ~20,000 US respondents, recruited through PureSpectrum

Matching social media & survey data

Survey Data

THE COVID STATES PROJECT
https://covidstates.org

- Online survey running since April of 2020
- Each wave includes ~20,000 US respondents, recruited through PureSpectrum
- About 100k unique respondents April 2020-Feb 2021

Matching social media & survey data

Survey Data

THE COVID STATES PROJECT
https://covidstates.org

- Online survey running since April of 2020
- Each wave includes ~20,000 US respondents, recruited through PureSpectrum
- About 100k unique respondents April 2020-Feb 2021

Social Media Data

- Respondents invited to share their Twitter handle

Validating Handles

We removed handles that were:

Validating Handles

We removed handles that were:

Just common names

- @John
- @Sarah

Validating Handles

We removed handles that were:

Just common names

- @John
- @Sarah

Not individuals

- @Google
- @McDonalds

Proxy: had more than 100,000 followers

Validating Handles

We removed handles that were:

Just common names

- @John
- @Sarah

Not individuals

- @Google
- @McDonalds

Proxy: had more than 100,000 followers

Given by multiple respondents

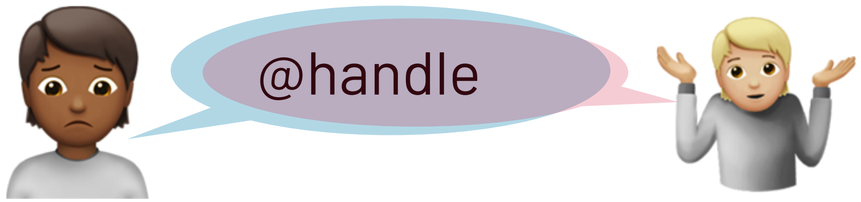

Validating Handles

We removed handles that were:

Just common names

- @John
- @Sarah

Not individuals

- @Google
- @McDonalds

Proxy: had more than 100,000 followers

Given by multiple respondents

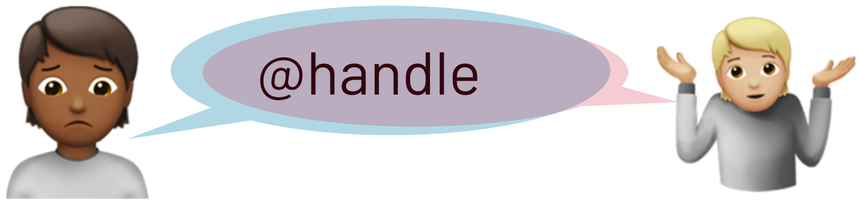

Different across same respondent

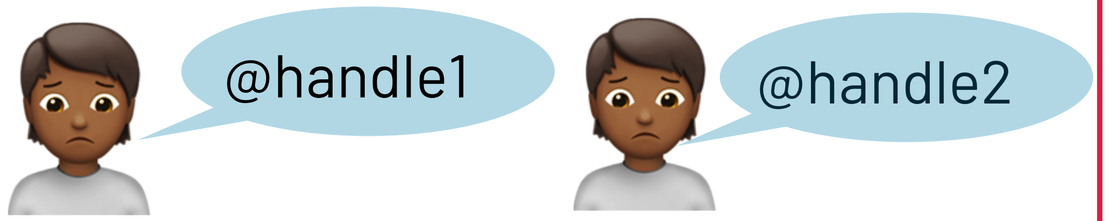

Survey & Social Media Data

Survey & Social Media Data

Across all respondents, **~ 14%** of survey respondents provided a valid Twitter handle

Survey & Social Media Data

Across all respondents, **~ 14%** of survey respondents provided a valid Twitter handle

15,160 handles matched to survey data

Survey & Social Media Data

Across all respondents, **~ 14%** of survey respondents provided a valid Twitter handle

15,160 handles matched to survey data

7,943 users for whom we have:

At least 1 wave of survey response

Validated handle with at least 1 Tweet April 2020 - Feb 2021

Survey & Social Media Data

Across all respondents, **~ 14%** of survey respondents provided a valid Twitter handle

15,160 handles matched to survey data

7,943 users for whom we have:

At least 1 wave of survey response

Validated handle with at least 1 Tweet April 2020 - Feb 2021

Primary analysis on **1,129 people** who:

Responded to the survey with a valid handle

Tweeted about specific topics

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

1. **Lockdowns**: 26,942 tweets from 2,668 respondents

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

1. **Lockdowns**: 26,942 tweets from 2,668 respondents
2. **Masks**: 28,609 tweets from 1,816 respondents

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

1. **Lockdowns**: 26,942 tweets from 2,668 respondents
2. **Masks**: 28,609 tweets from 1,816 respondents
3. **Trump**: 485,906 tweets from 3,323 respondents

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

1. **Lockdowns**: 26,942 tweets from 2,668 respondents
2. **Masks**: 28,609 tweets from 1,816 respondents
3. **Trump**: 485,906 tweets from 3,323 respondents
4. **Vaccines**: 29,935 tweets from 1,338 respondents

Topics

Took Tweets from 7,943 validated handles and used keywords to identify posts about 4 target topics:

1. **Lockdowns**: 26,942 tweets from 2,668 respondents
2. **Masks**: 28,609 tweets from 1,816 respondents
3. **Trump**: 485,906 tweets from 3,323 respondents
4. **Vaccines**: 29,935 tweets from 1,338 respondents

These topics all had associated questions in our surveys

Assessing Survey Stance

Constructed 3-point (Pro, Anti, Neutral) stance variables for each examined target:

Lockdowns

Pro: Strongly Agree on average — 60%

Neutral: Other responses — 37%

Anti: Strongly Disagree on average — 3%

Masks

Pro: Closely following mask-guidelines — 66%

Neutral: Other responses — 18%

Anti: Not following mask guidelines — 15%

Trump

Pro: Would/did vote for Trump — 32%

Neutral: Unsure — 8%

Anti: Would/did vote for Biden — 60%

Vaccines

Pro: vaccinated/Extremely Likely to be vaccinated — 43%

Neutral: Other responses — 44%

Anti: Extremely unlikely to be vaccinated — 14%

Sampling for qualitative analysis

Sampling for qualitative analysis

Step 1	Match users to survey data to assess stance

Sampling for qualitative analysis

Step 1

Match users to survey data to assess stance

Step 2

Stratify by survey stance

Sampling for qualitative analysis

Step 1

Match users to survey data to assess stance

Step 2

Stratify by survey stance

Step 3

Stratify by posting frequency

Sampling for qualitative analysis

Posting activity

Survey Stance

	Low	Moderate	High
Pro			
Neutral			
Anti			

For each target, sampled up to 40 users from each bucket

Resulted in **1,129 unique users** across 4 topics

Assessing Tweet Stance

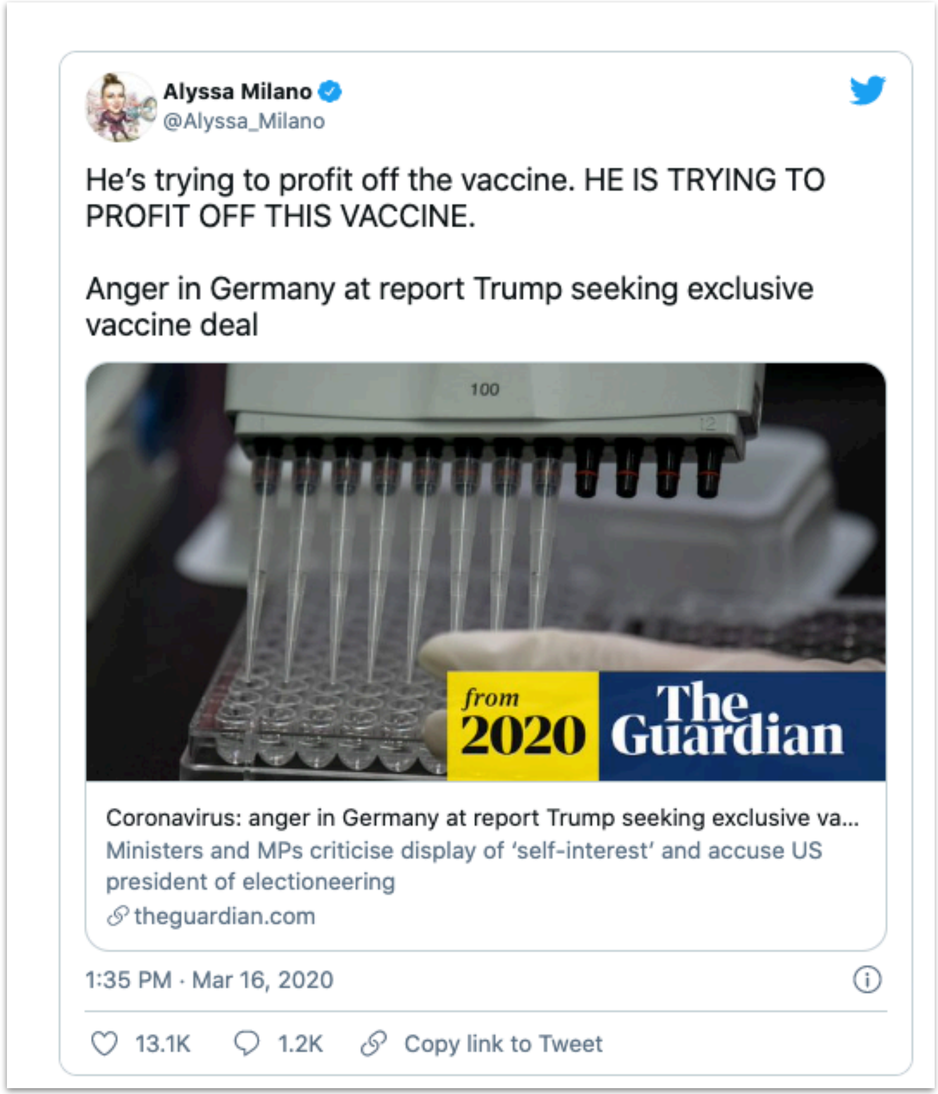

Assessing Tweet Stance

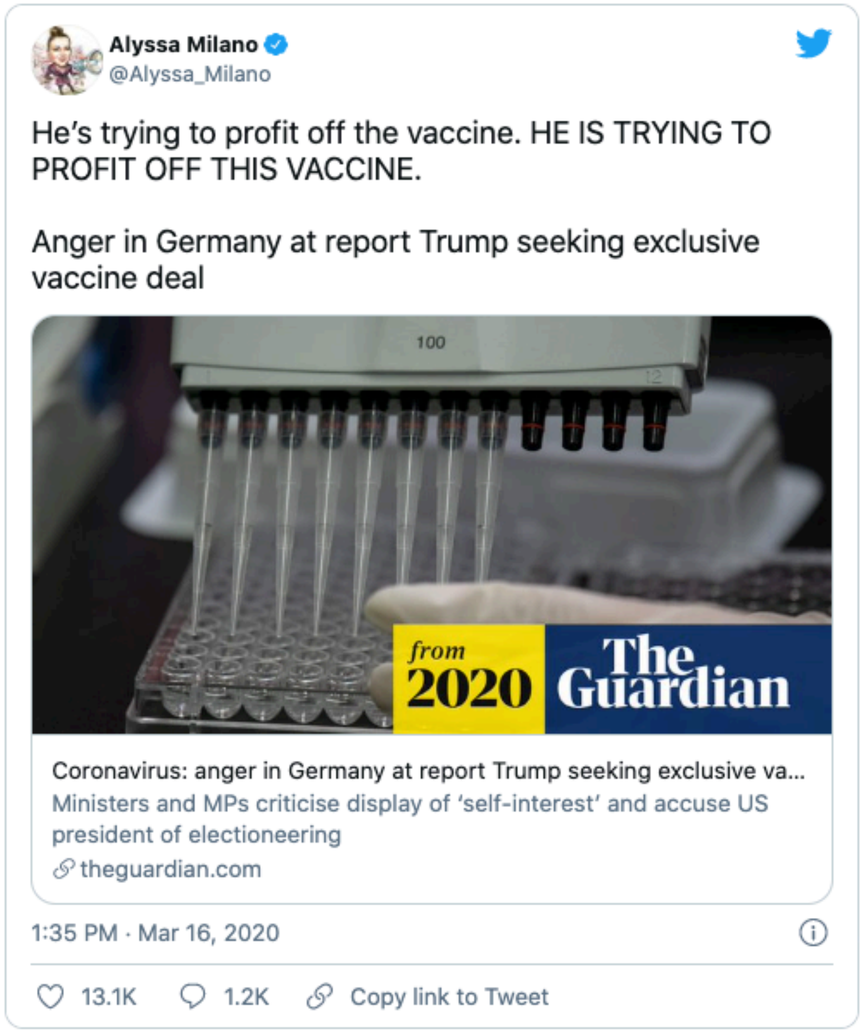

Given the tweets above, do you believe this user is ...

Resistant to Vaccines (Anti-vax)	Hesitant towards Vaccines (in between)	Supportive of Vaccines (Pro-vax)

How confident are you in your decision?

Not at all	Somewhat	Very

Assessing Tweet Stance

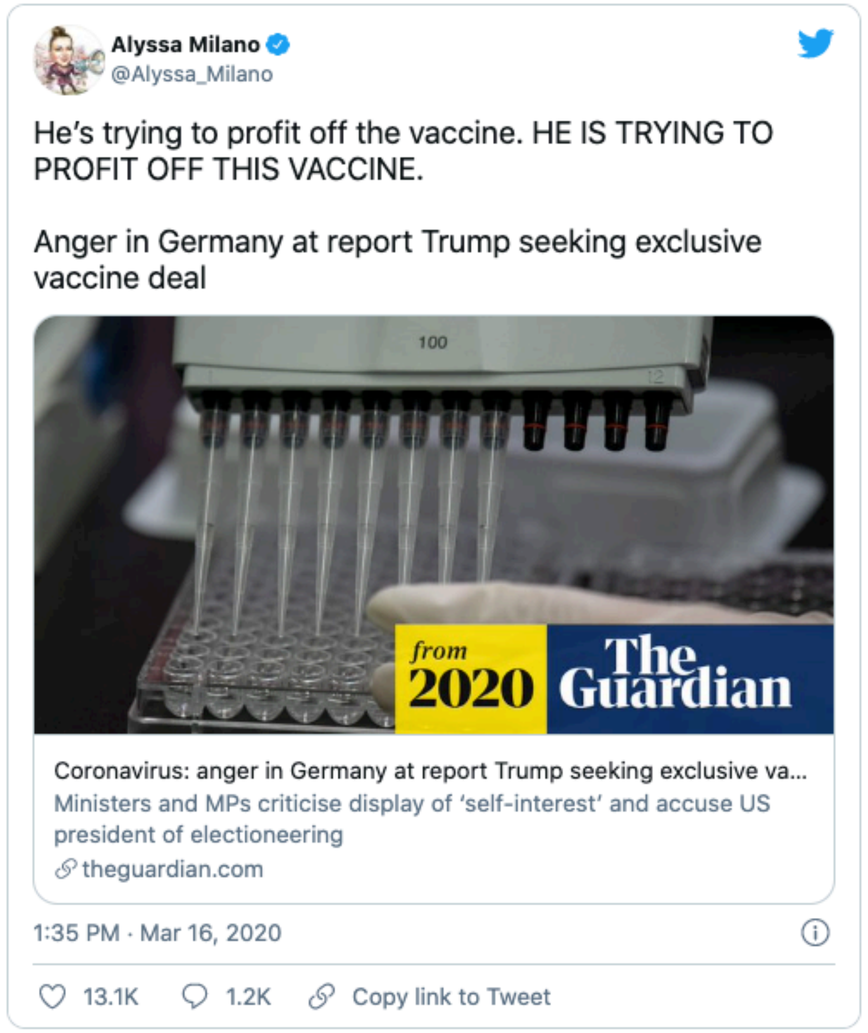

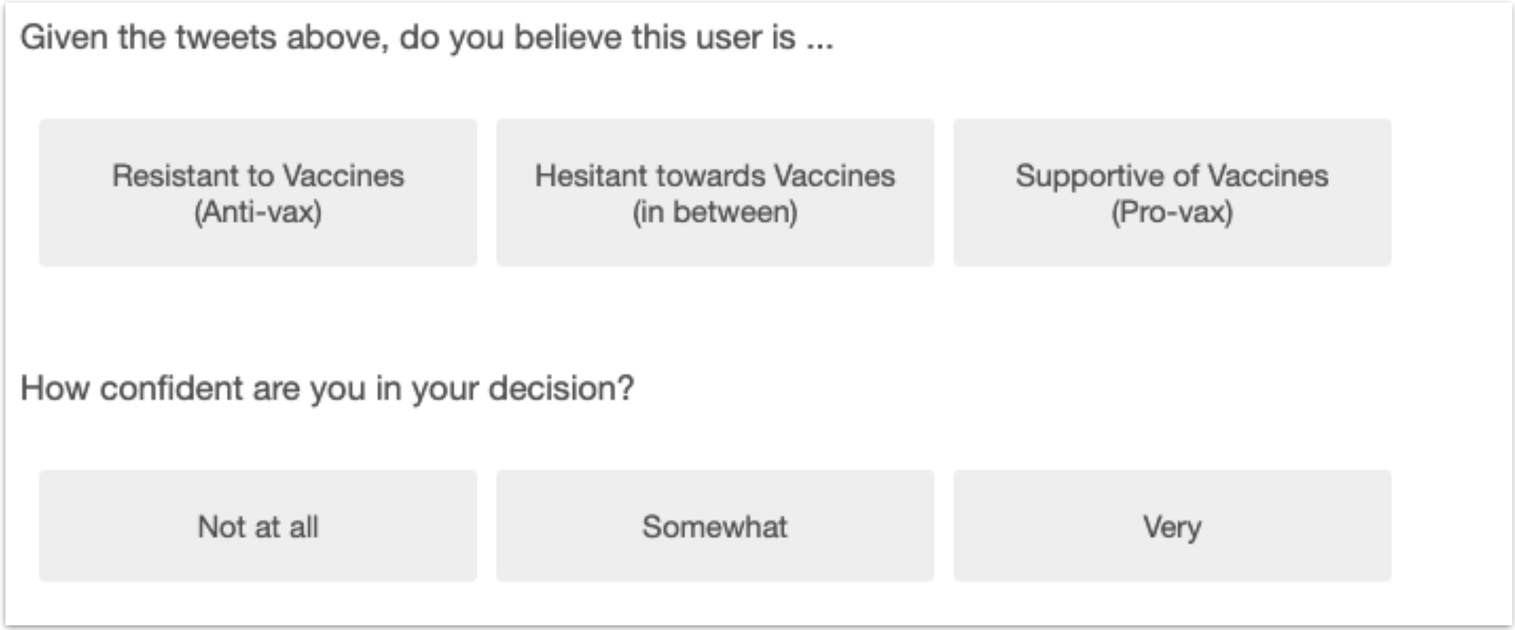

- 5 trained, human coders labeled ~570 tweets each
- Each tweet was labeled by two annotators; disagreements resolved by a third author

R1: When do stance annotators agree?

R1: When do stance annotators agree?

Krippendorf's Alpha:

- 2 annotators
- 3 confidence levels

R1: When do stance annotators agree?

Krippendorf's Alpha:

- 2 annotators
- 3 confidence levels

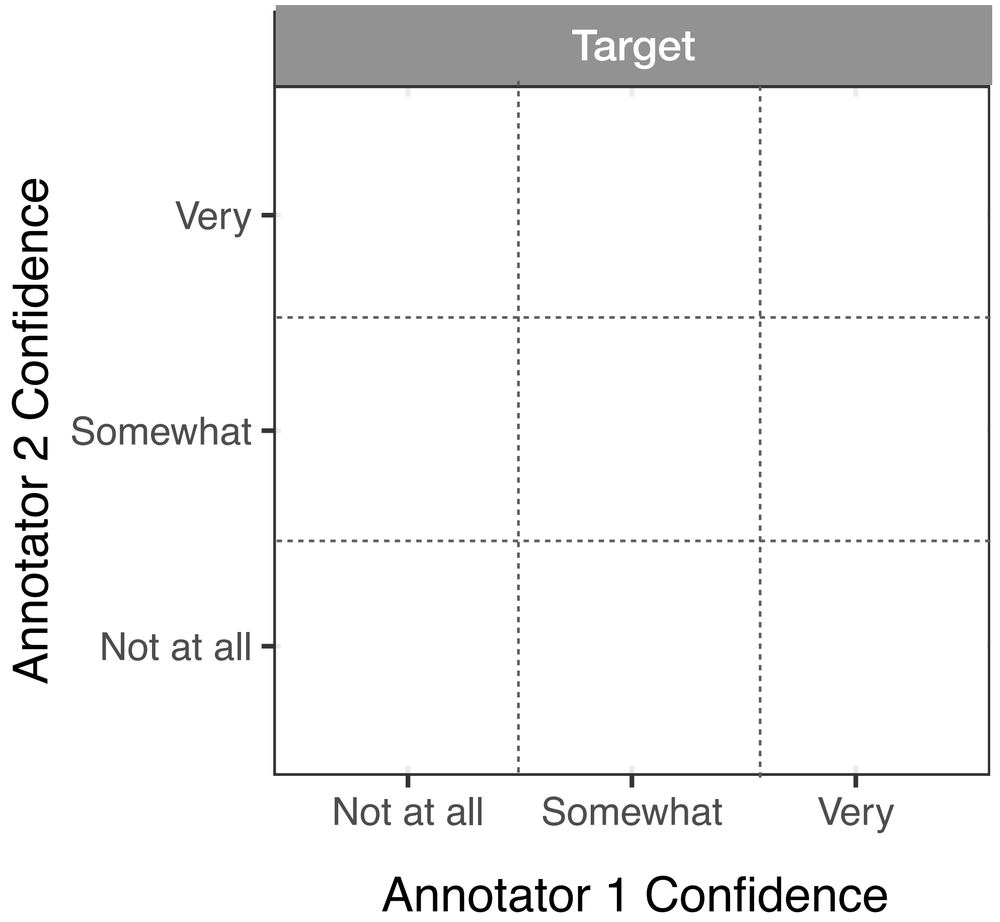

R1: When do stance annotators agree?

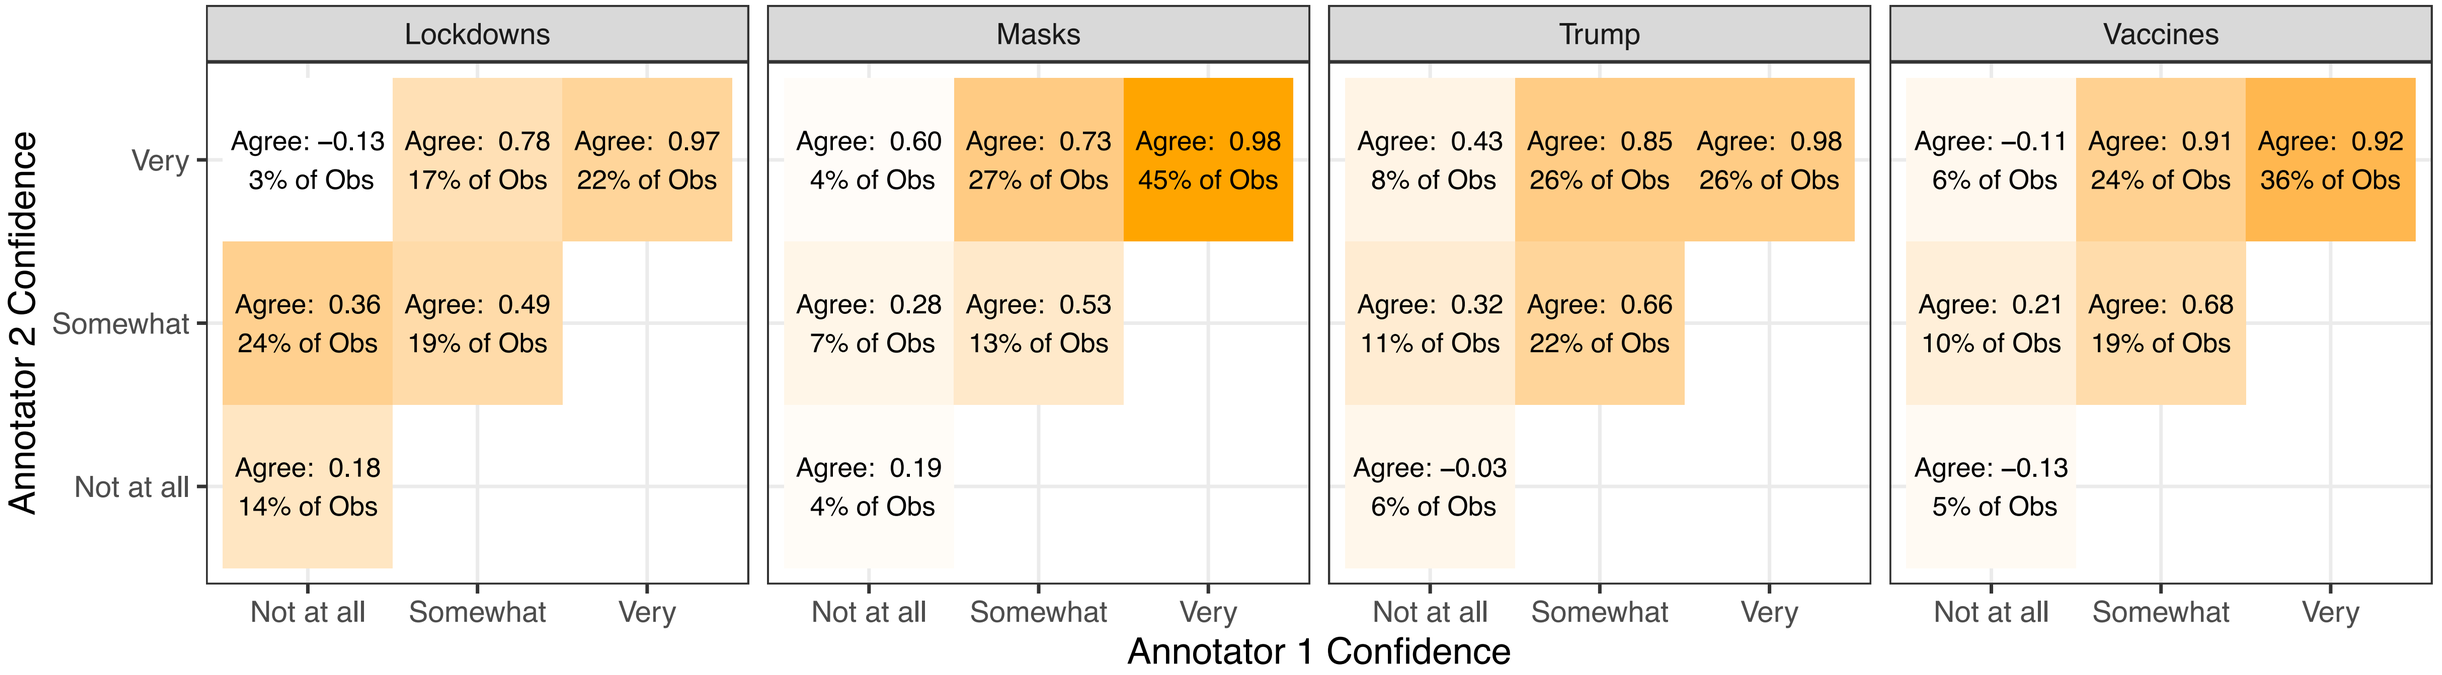

R1: When do stance annotators agree?

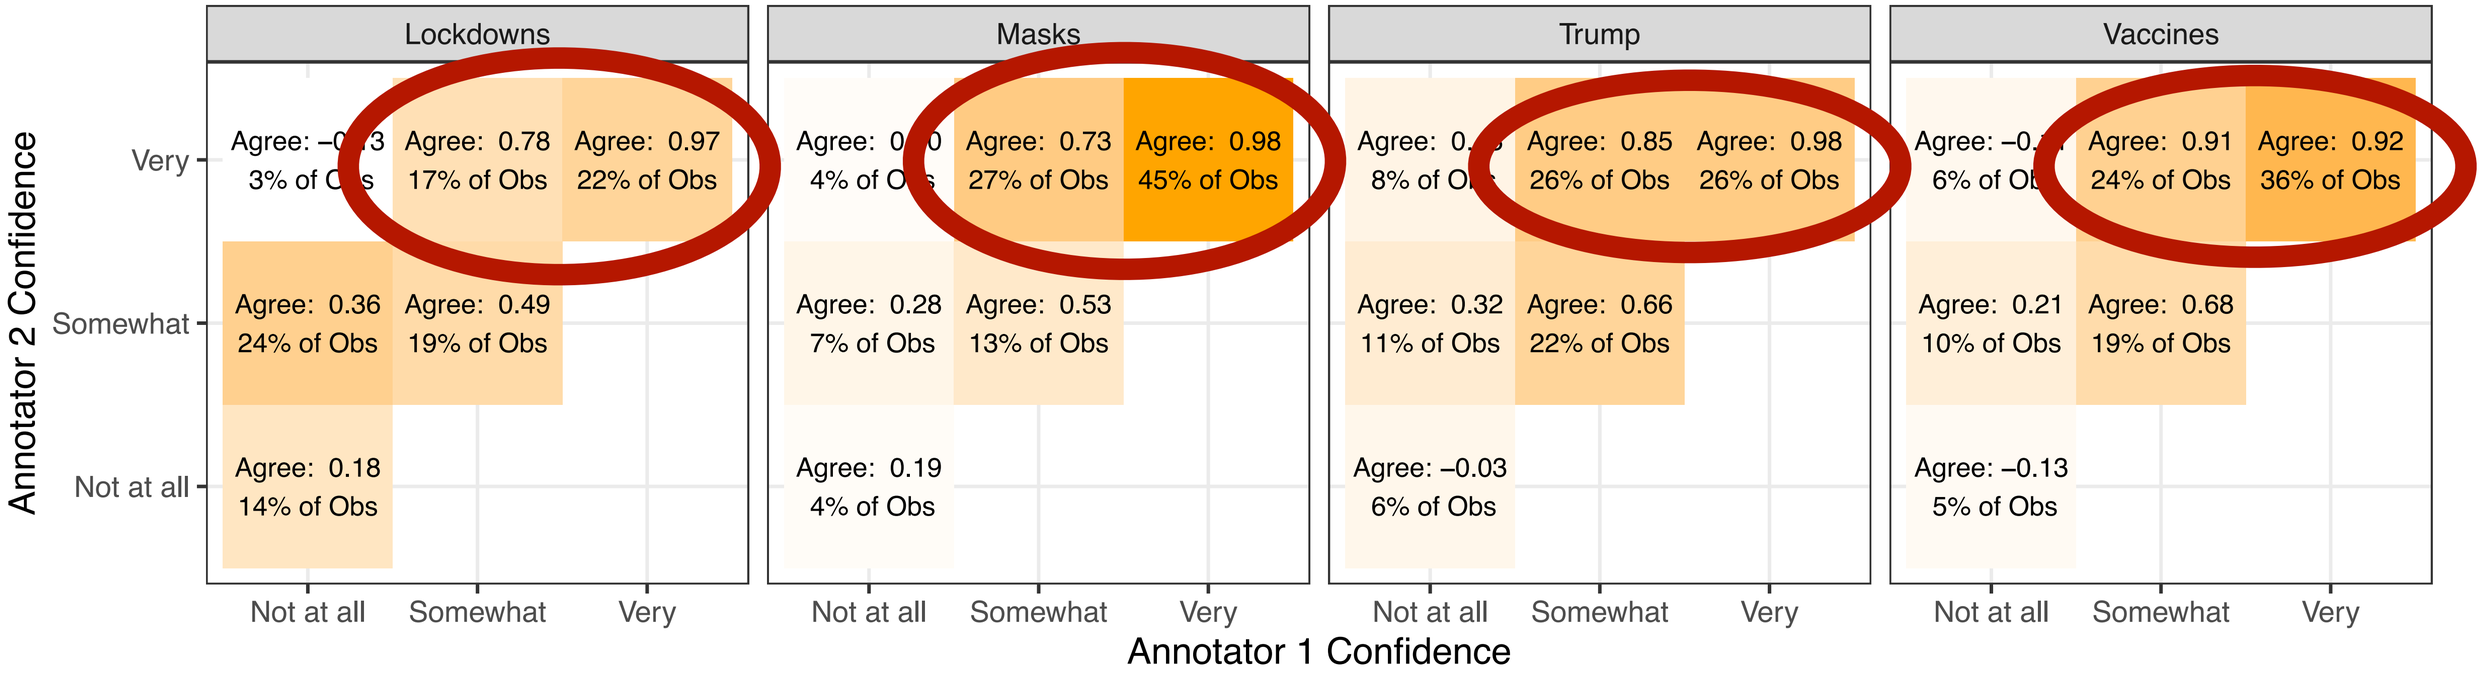

High agreement when annotators were confident

R1: When do stance annotators agree?

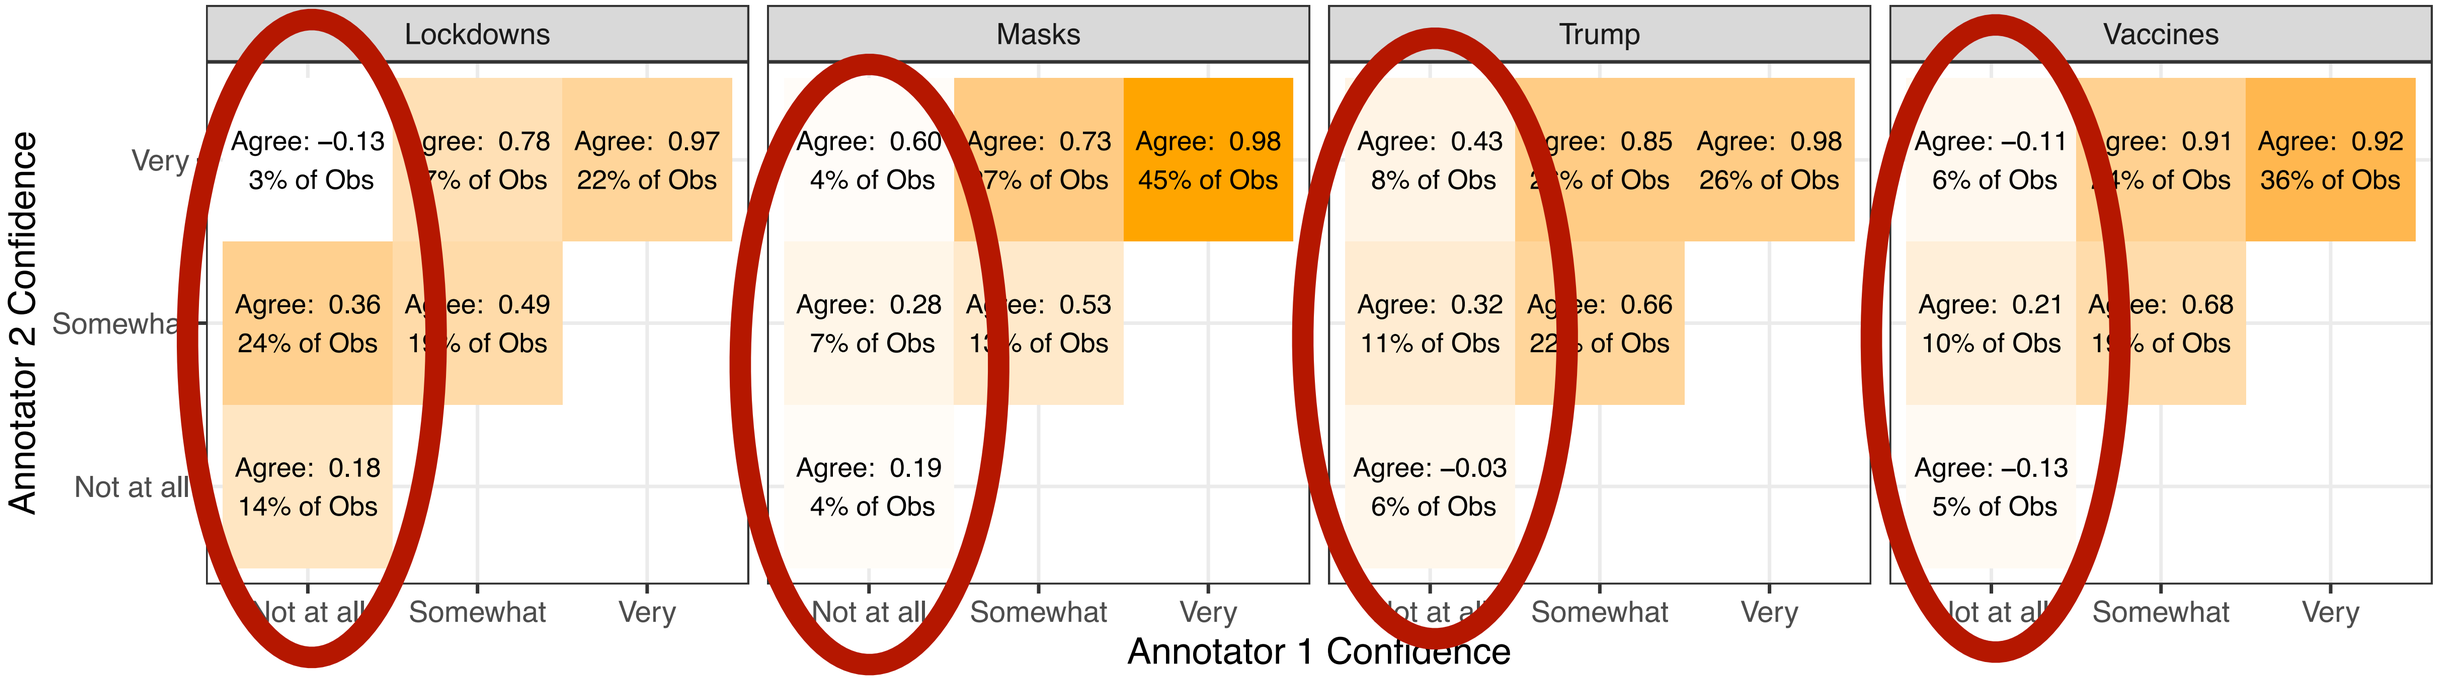

Low agreement when annotators were not confident

R1: When do stance annotators agree?

Take away: **Some posts don’t contain enough information to reliably assess stance *on a policy issue***

R1: When do stance annotators agree?

Take away: Some posts don't contain enough information to reliably assess stance *on a policy issue*

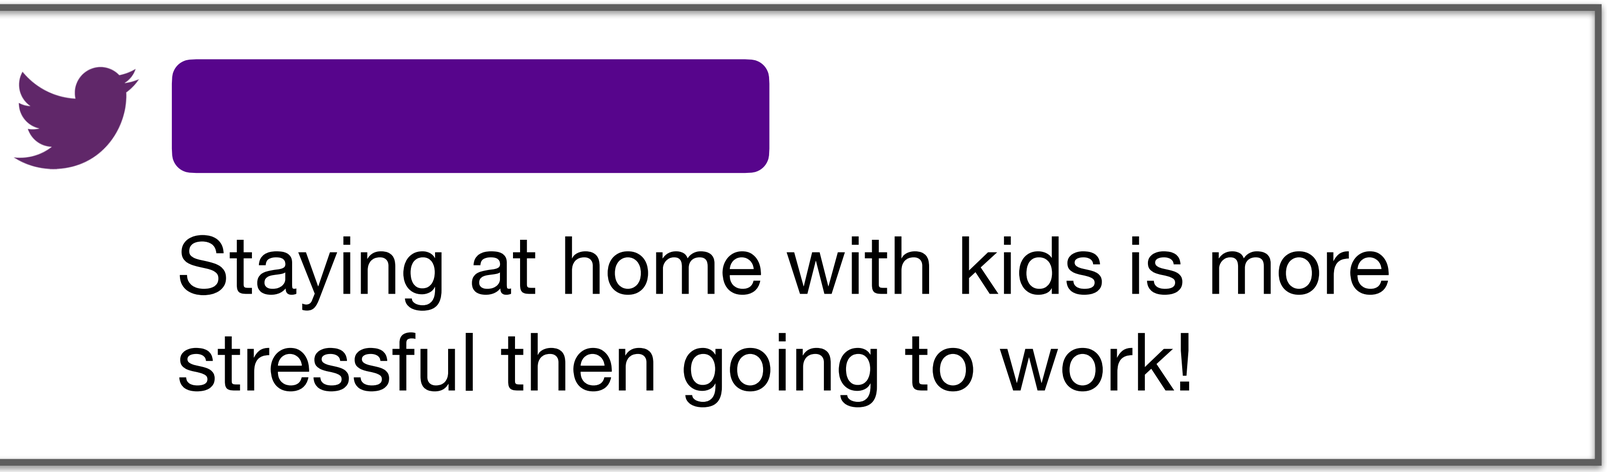

R1: When do stance annotators agree?

Take away: Some posts don't contain enough information to reliably assess stance *on a policy issue*

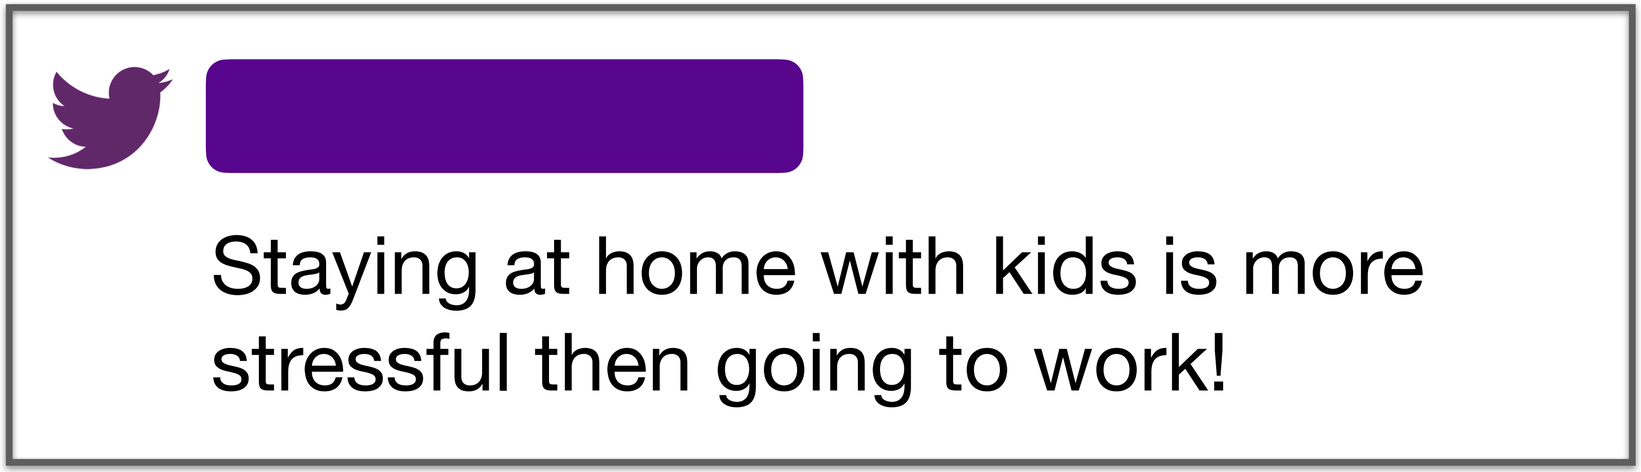

Expressing lockdown fatigue is not necessarily the same as being "anti-Lockdown"

R2: Does “stance” map to survey data?

R2: Does “stance” map to survey data?

Using survey responses as ground truth, how accurate were human annotators?

(Subsetted to tweets where annotators were somewhat or confident)

R2: Does “stance” map to survey data?

Recall

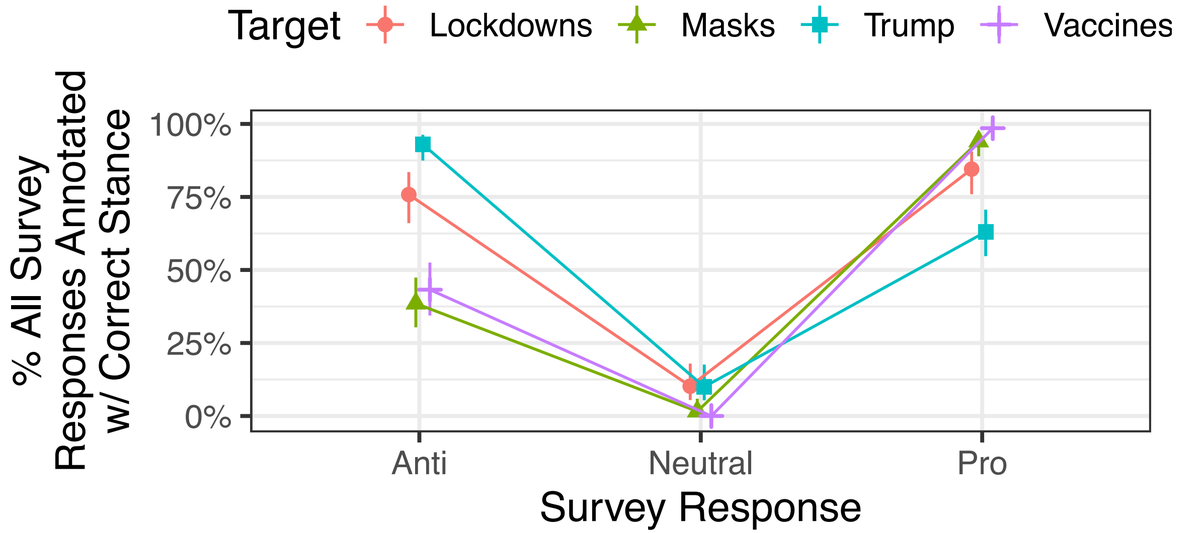

Precision

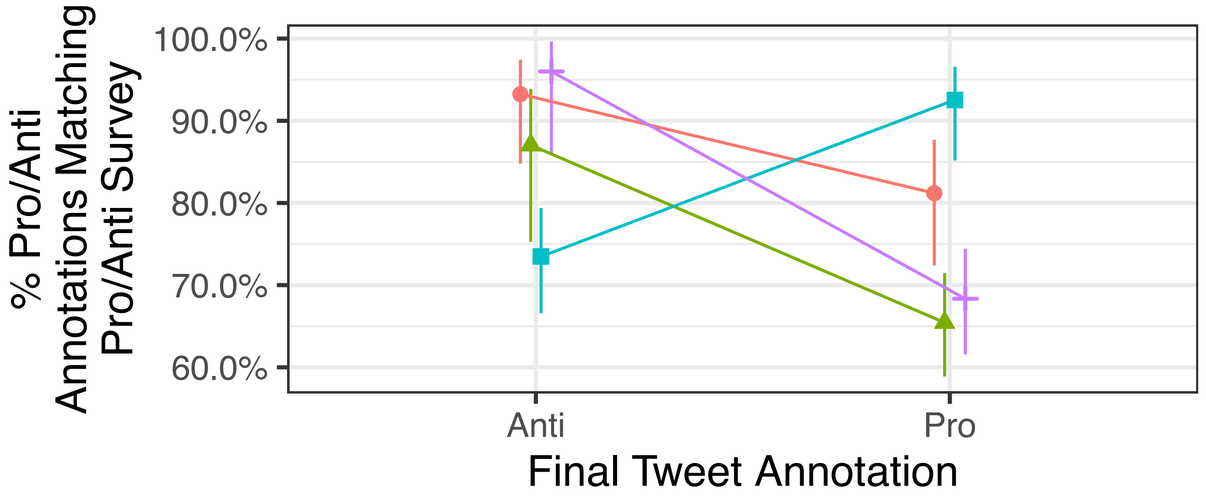

Using survey responses as ground truth, how accurate were human annotators?

(Subsetted to tweets where annotators were somewhat or confident)

R2: Does “stance” map to survey data?

Recall

Precision

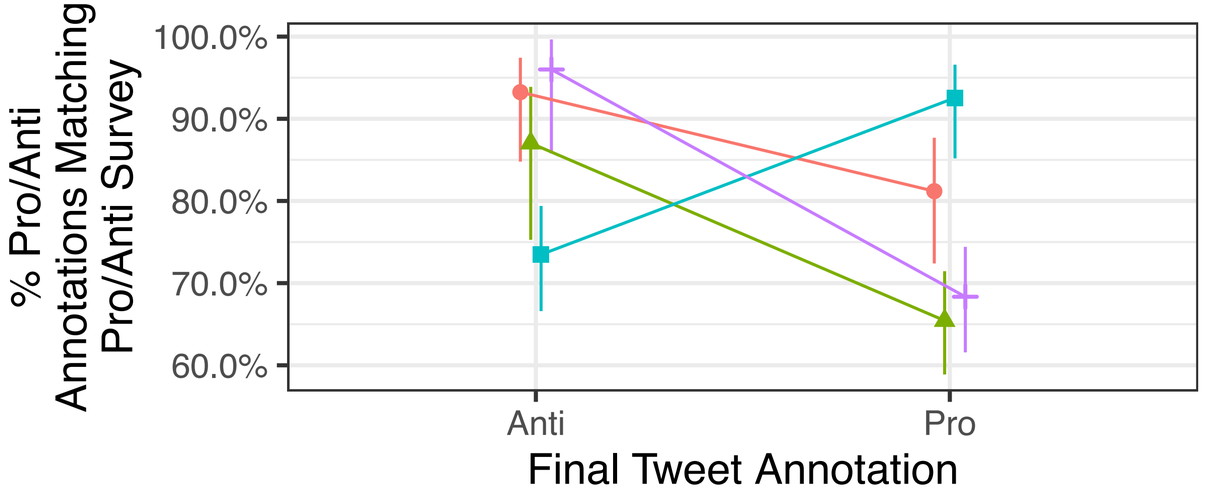

Anti-Trump,
Anti-lockdown,
Pro-lockdown,
Pro-mask,
Pro- vaccine
almost always annotated
with the same stance

R2: Does “stance” map to survey data?

Recall

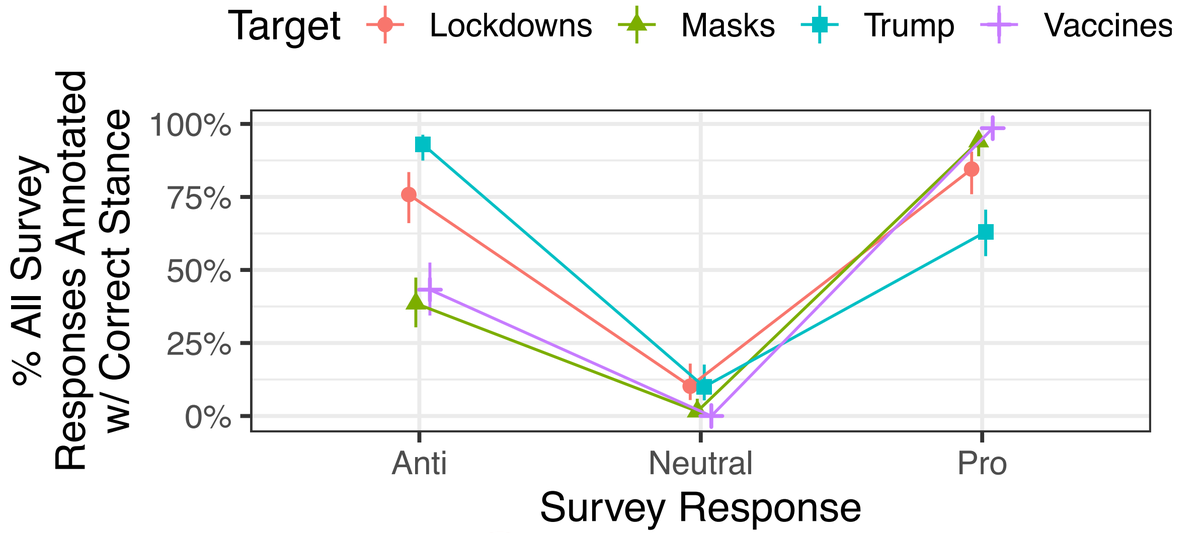

Precision

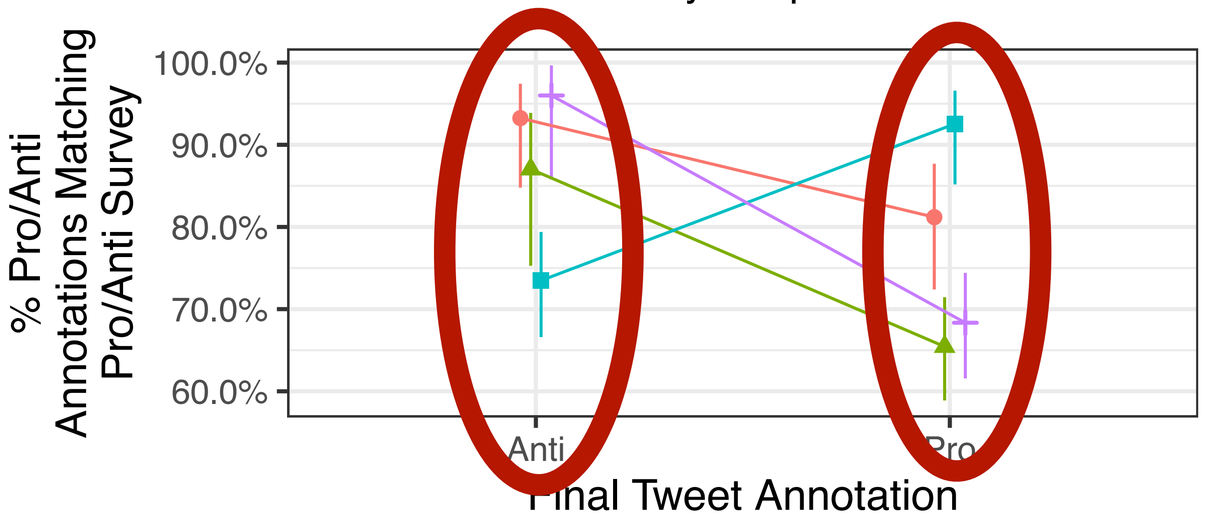

Anti-masks,
Anti-vaccines,
Anti- lockdowns,
Pro-Trump,
matched survey stances
over 90% of the time

R2: Does "stance" map to survey data?

Recall

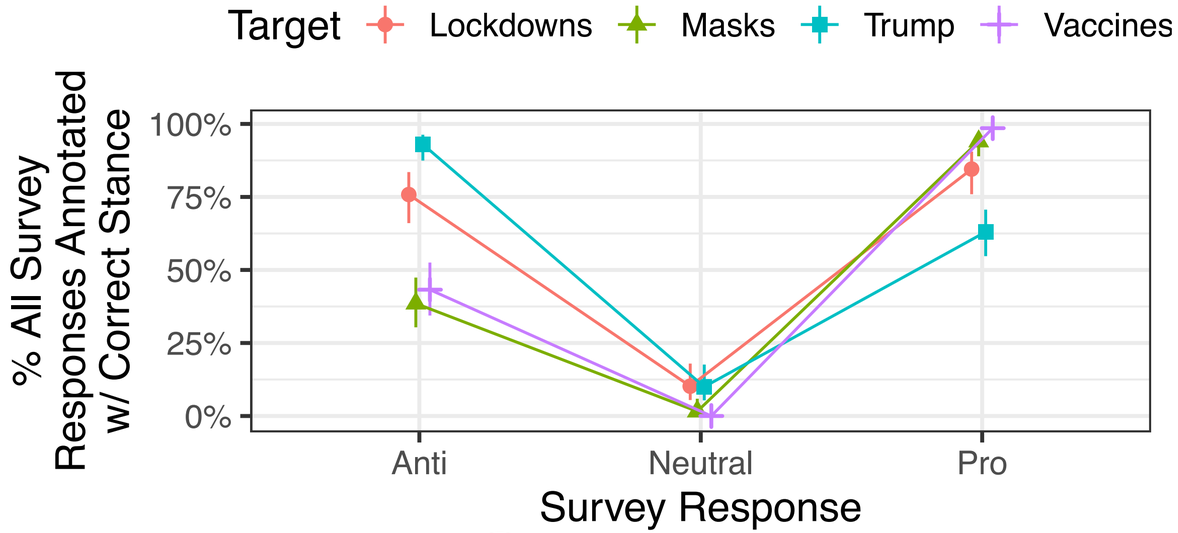

Precision

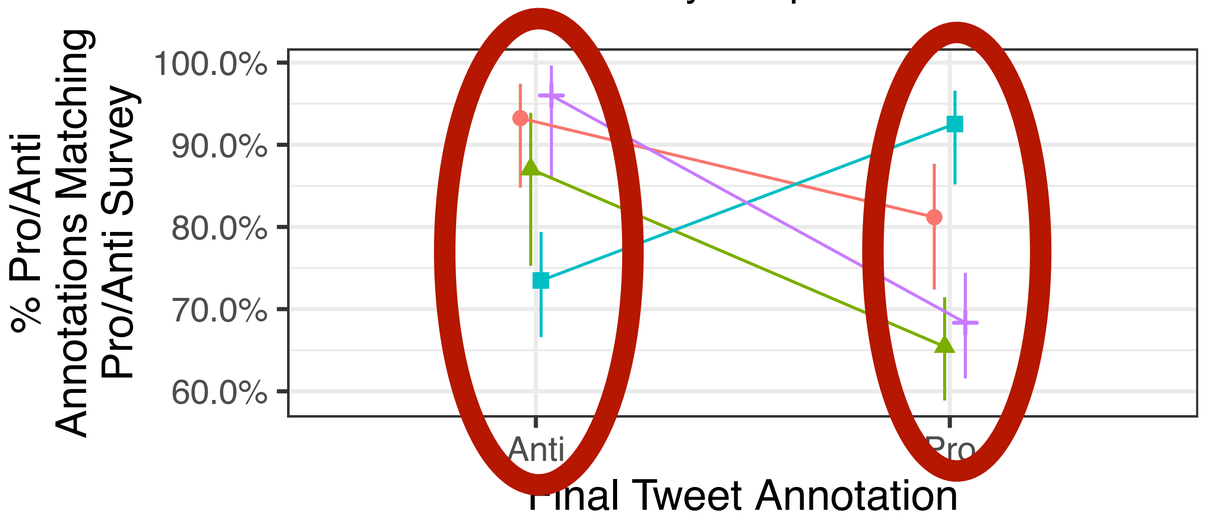

Anti-masks,
Anti-vaccines,
Anti- lockdowns,
Pro-Trump,
matched survey stances over 90% of the time

Is that "high"? Given that tweet authors and survey respondents are matched??

R2: Does “stance” map to survey data?

Recall

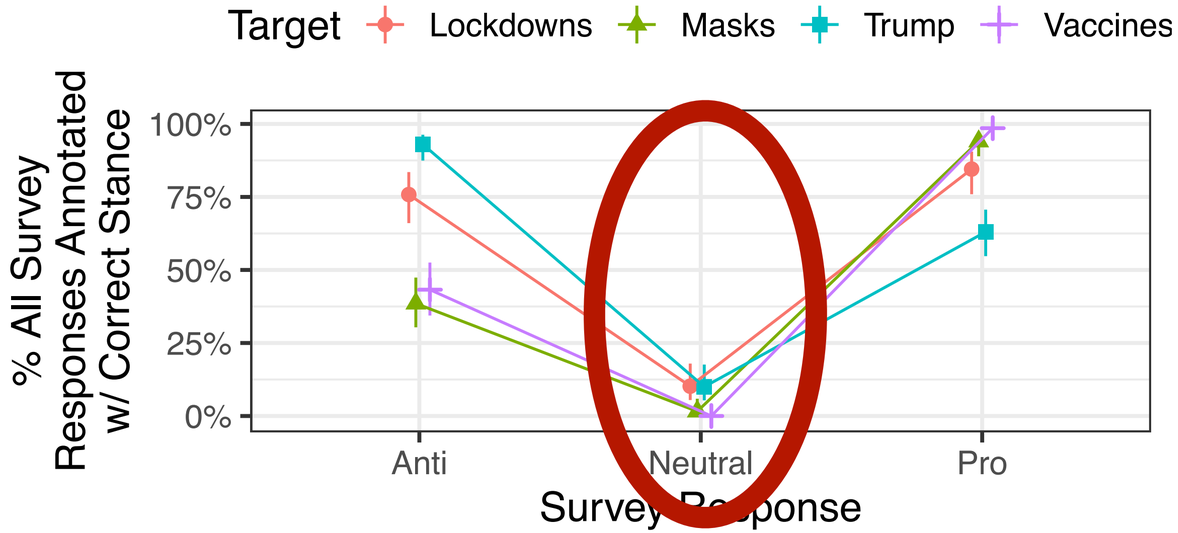

Precision

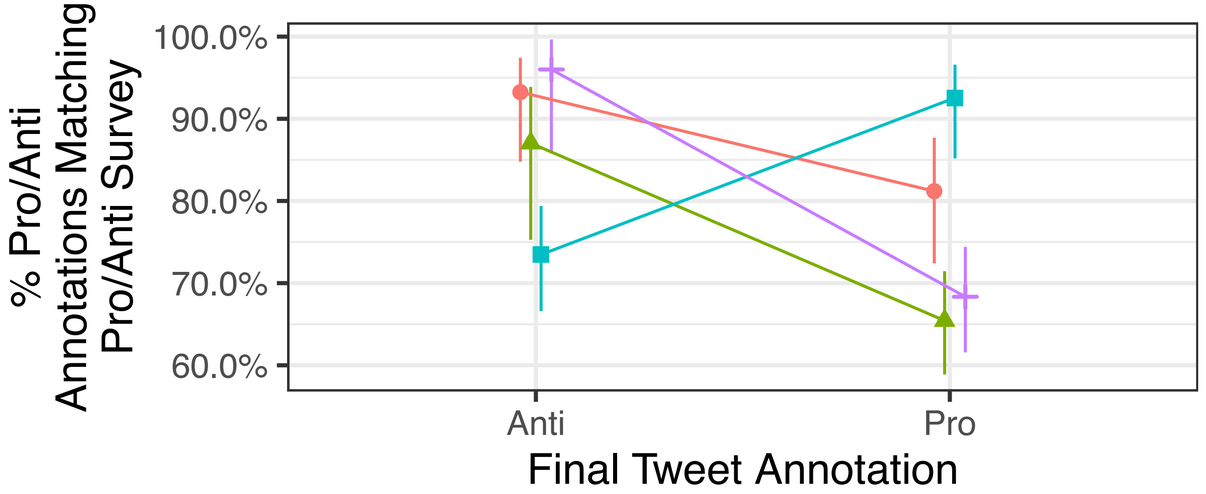

Recall of Neutral survey stances was near 0% across all four targets!

Annotators almost never confidently identified Neutral stance

Neutral tag was used to express the lack of stance indicators

Opinion May Change Over Time

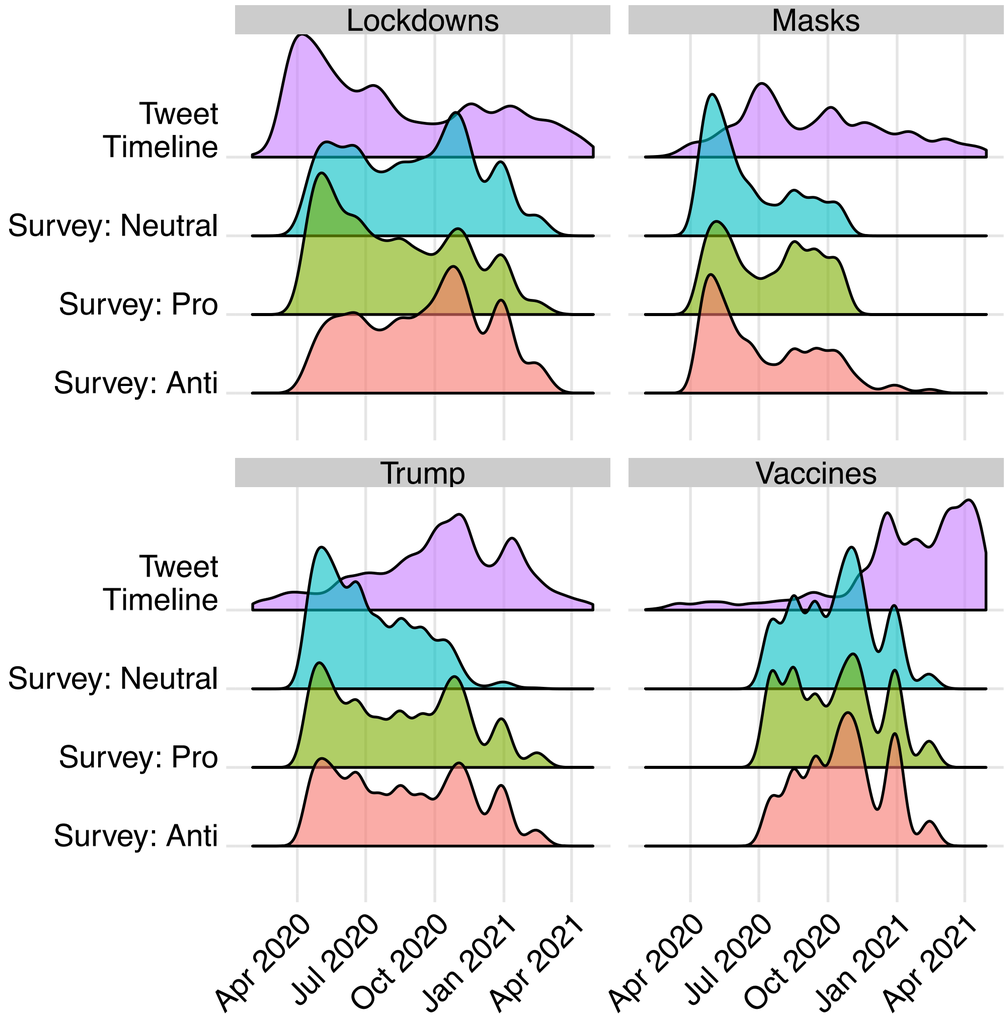

Opinion May Change Over Time

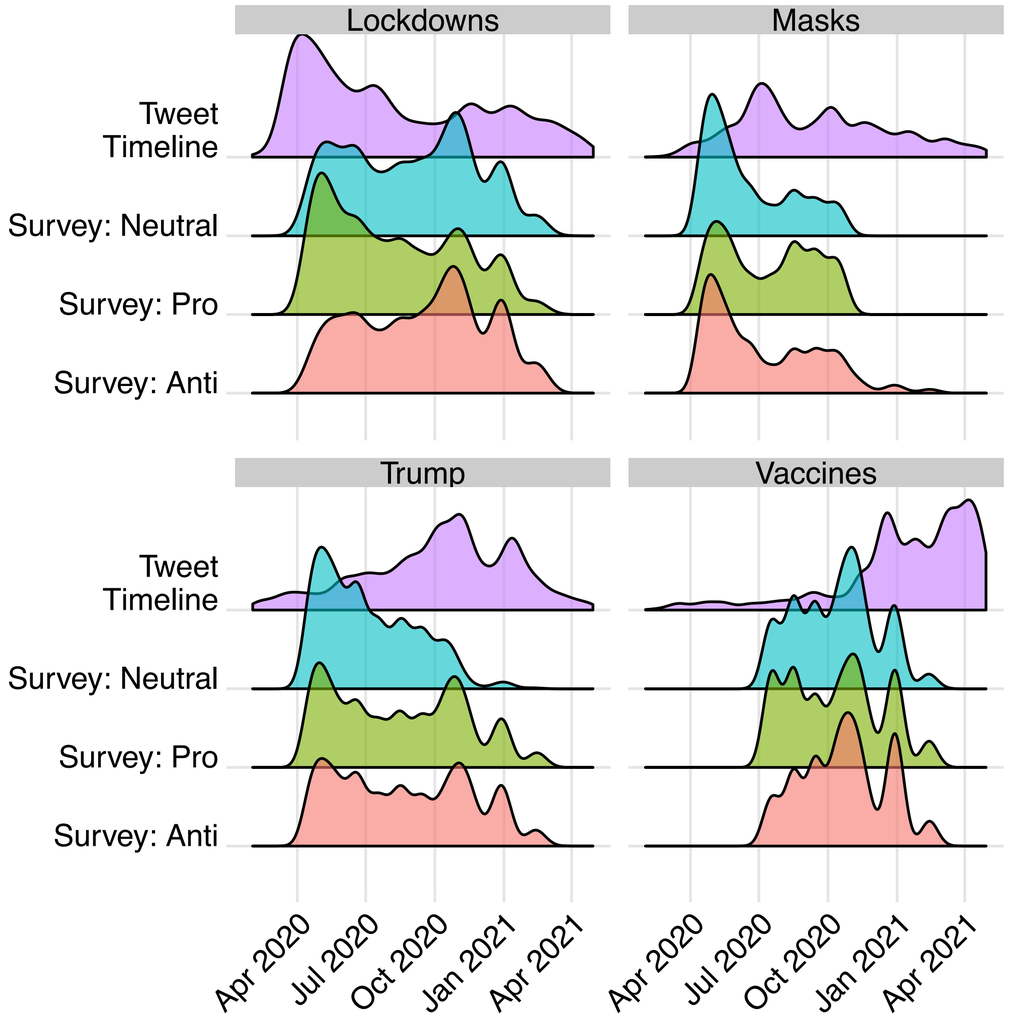

People tended to tweet about targets *after* survey responses

Opinion May Change Over Time

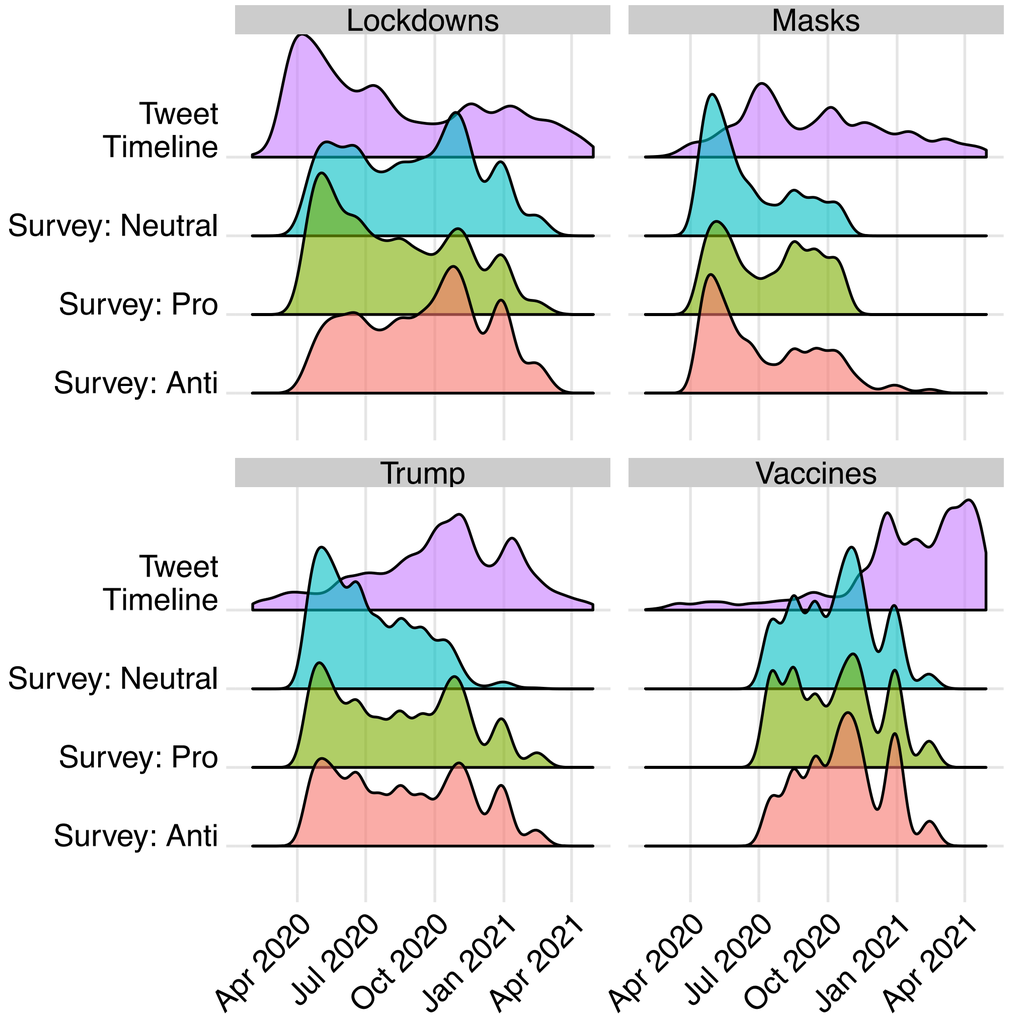

People tended to tweet about targets *after* survey responses

Opinion May Change Over Time

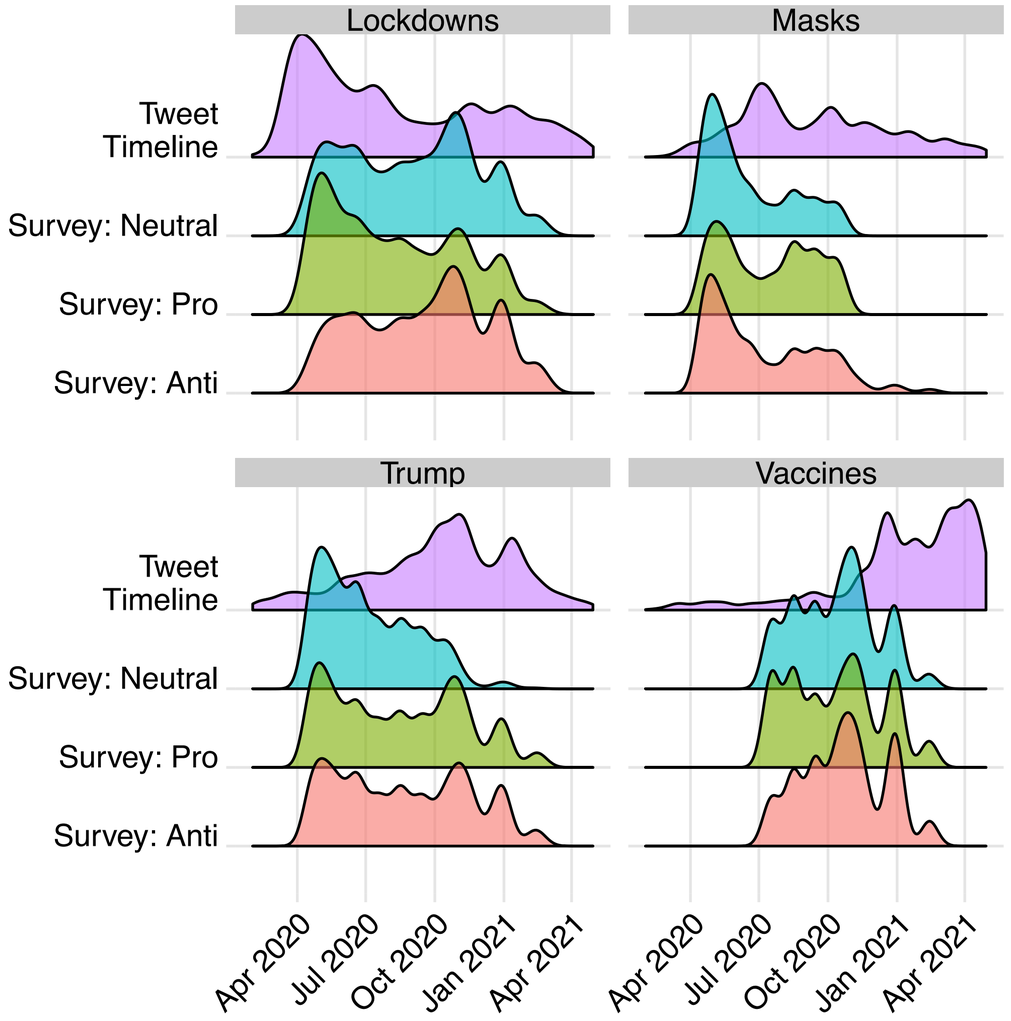

People tended to tweet about targets *after* survey responses

Survey designed to capture opinions on emerging topics

Opinion May Change Over Time

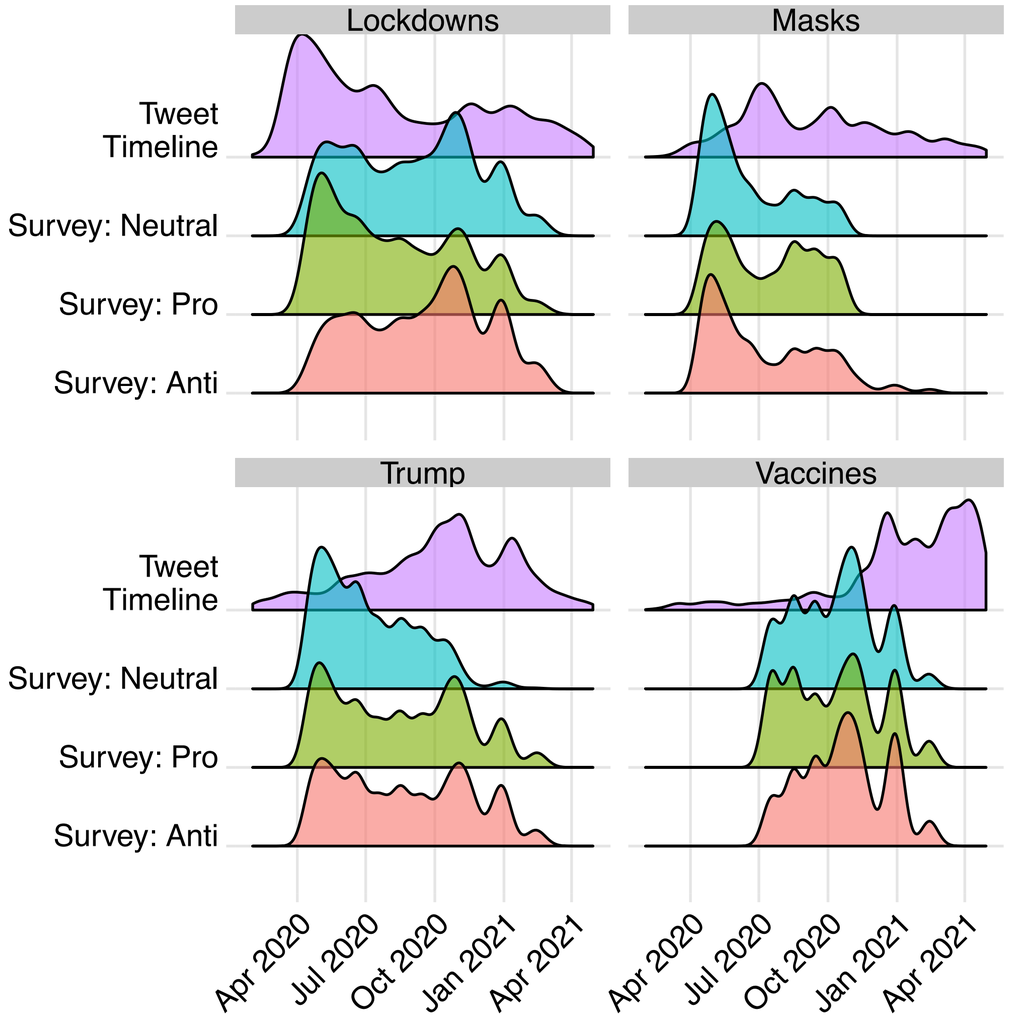

People tended to tweet about targets *after* survey responses

Survey designed to capture opinions on emerging topics

Opinion May Change Over Time

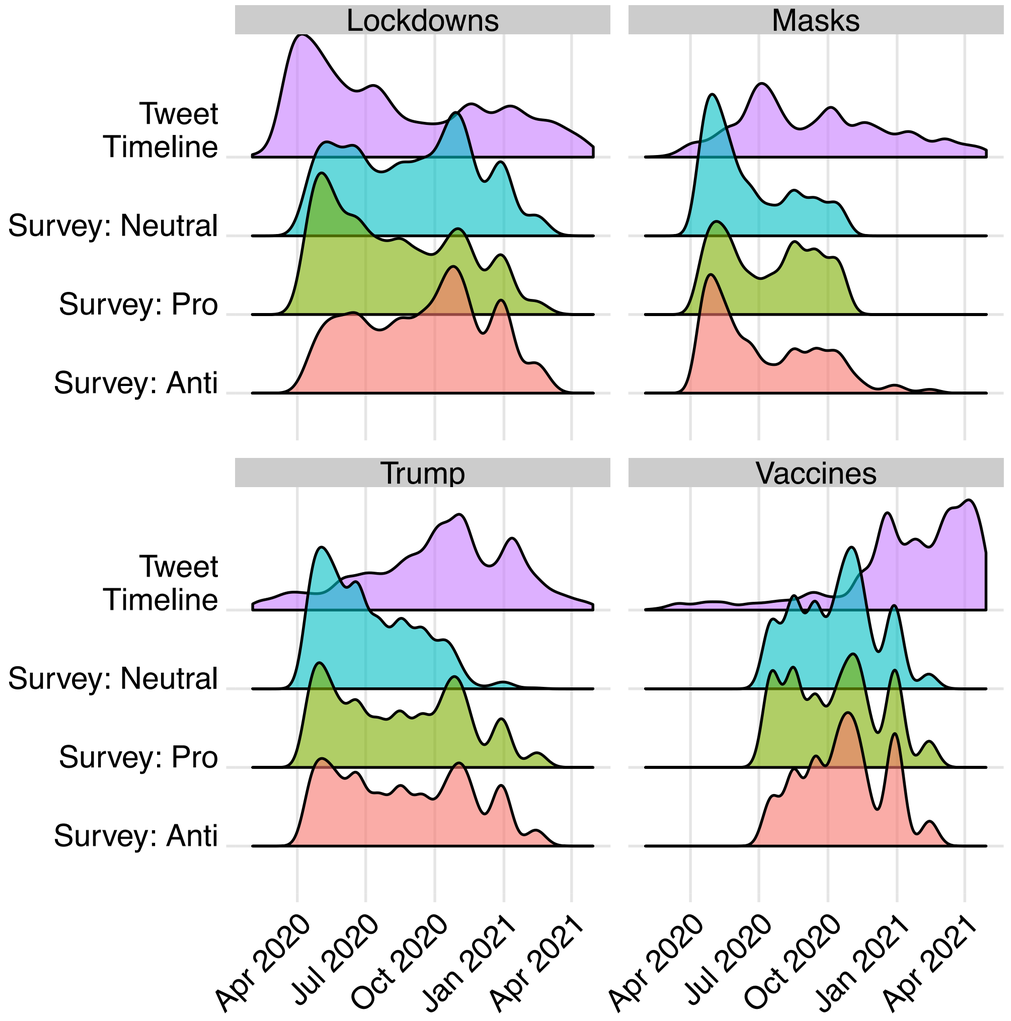

People tended to tweet about targets *after* survey responses

Survey designed to capture opinions on emerging topics

People don't tend to tweet neutral opinions

Opinion May Change Over Time

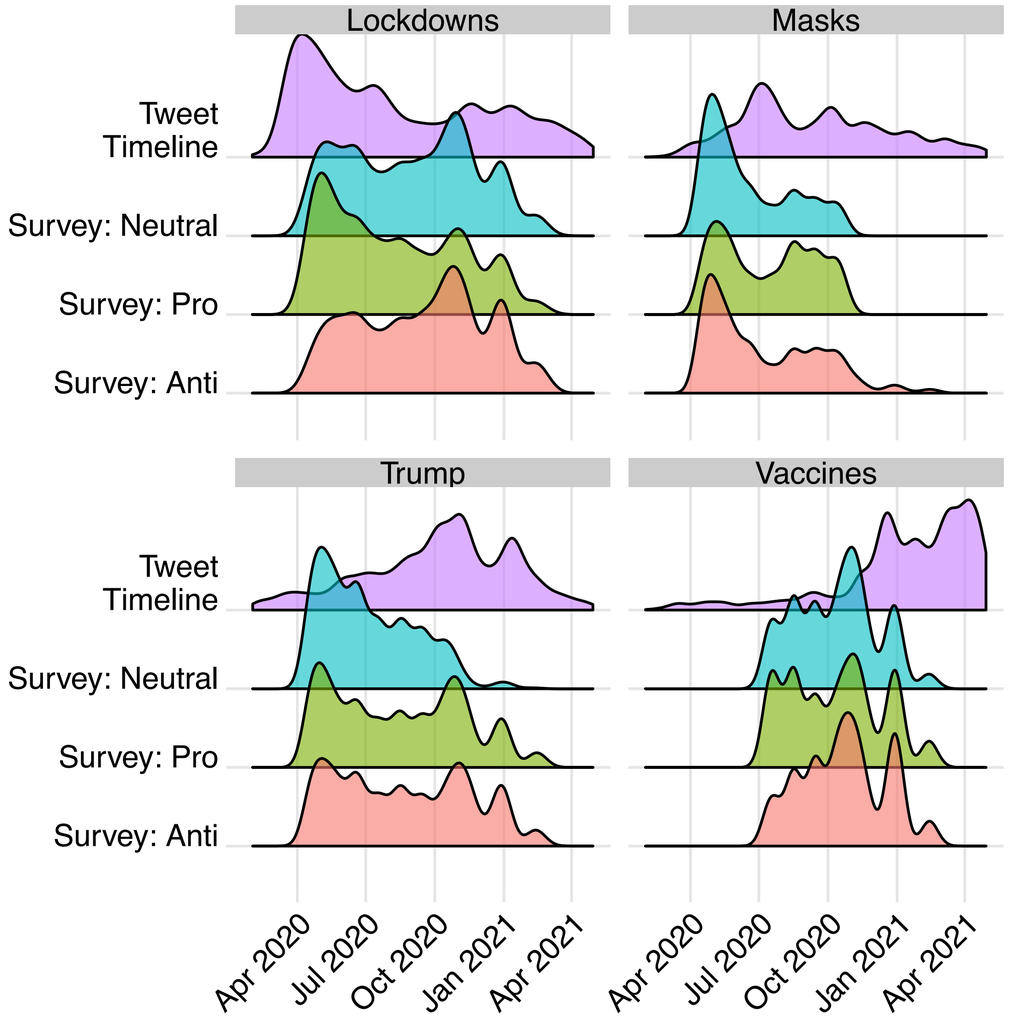

Take away:

Surveys can indicate how public opinion will develop while

Social media can reflect the most vocal opinions at a given moment

It's Not Just Opinion Change...

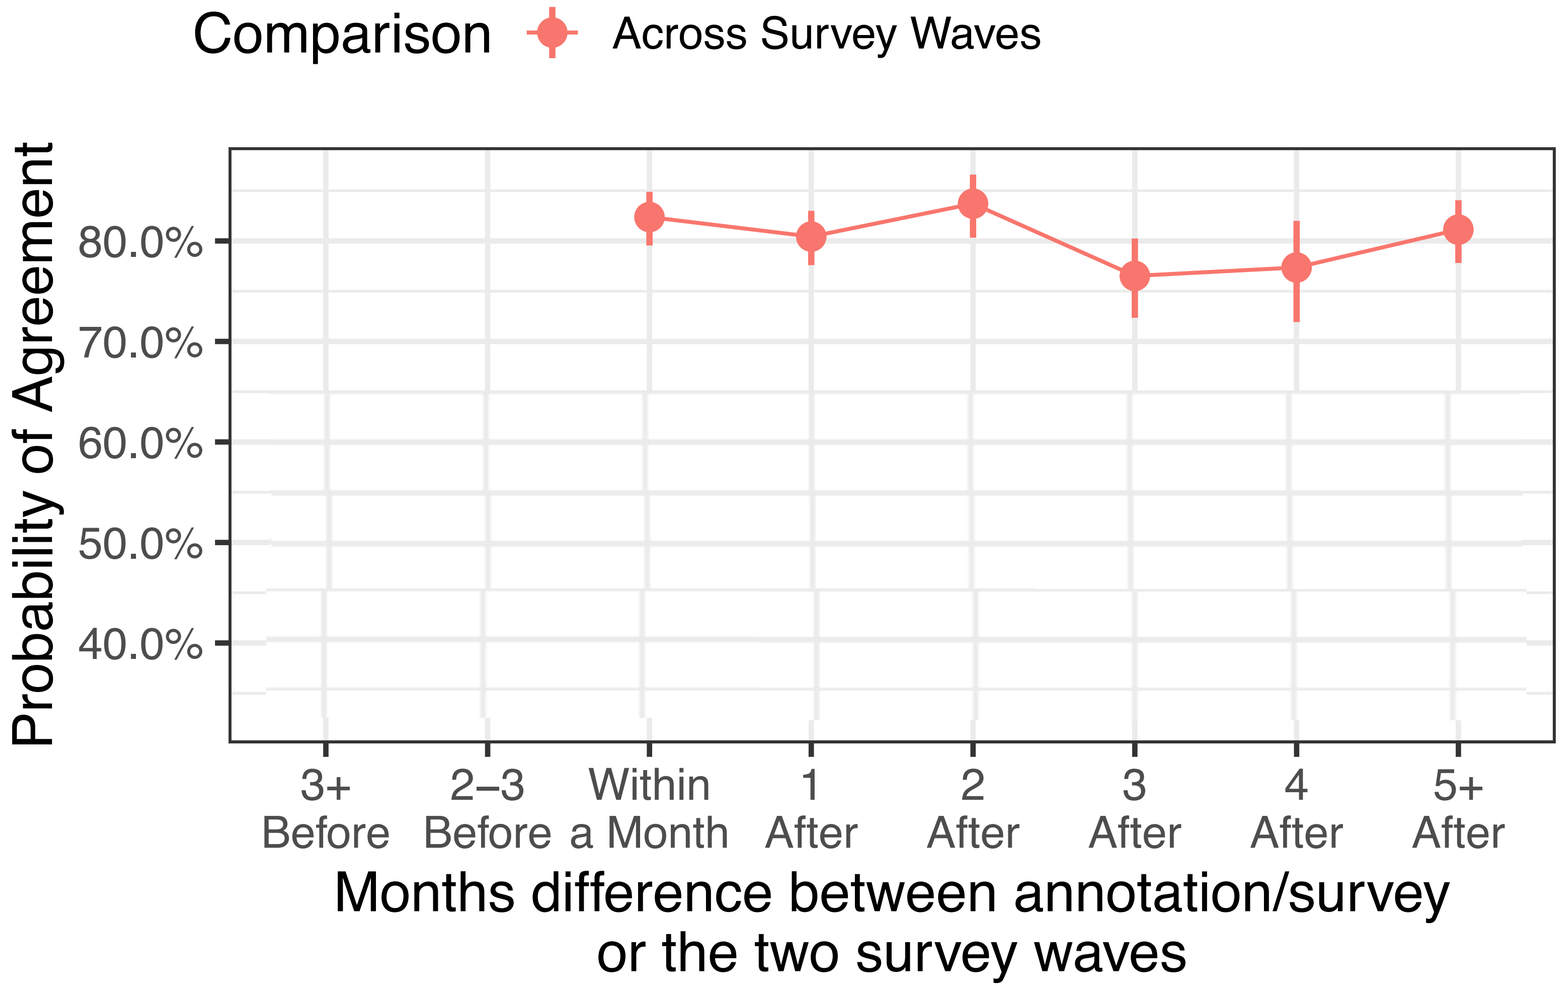

Survey respondents had fairly stable opinions across survey waves..

It's Not Just Opinion Change...

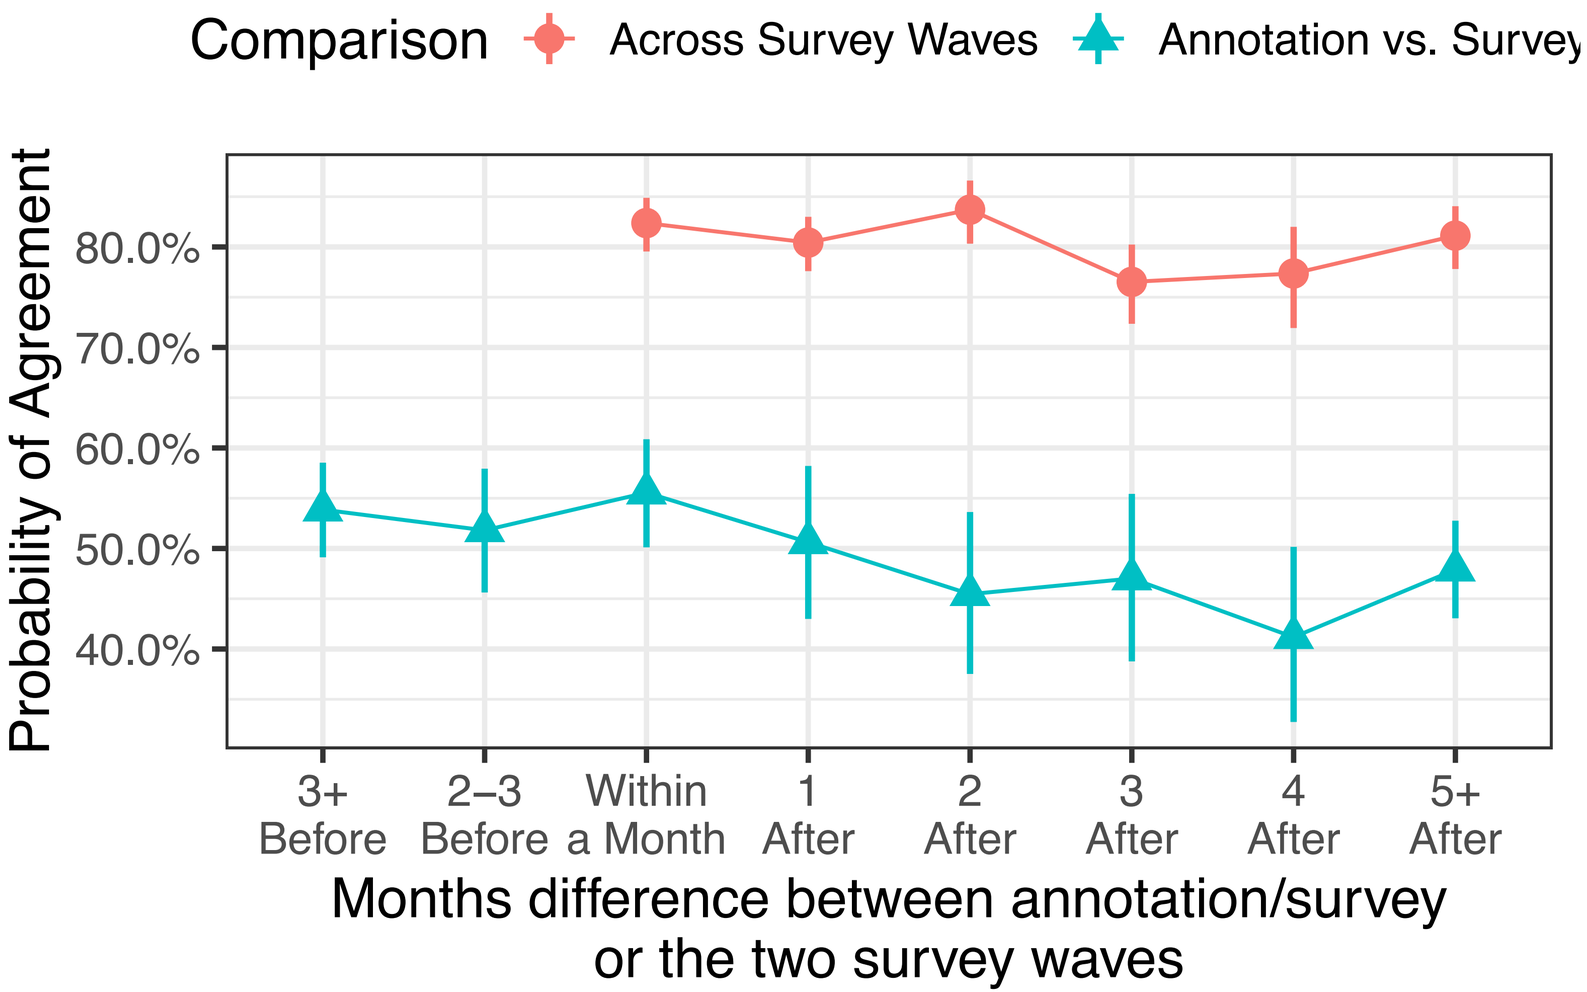

Survey respondents had fairly stable opinions across survey waves..

But tweet annotation and survey agreement was consistently lower

It's Not Just Opinion Change...

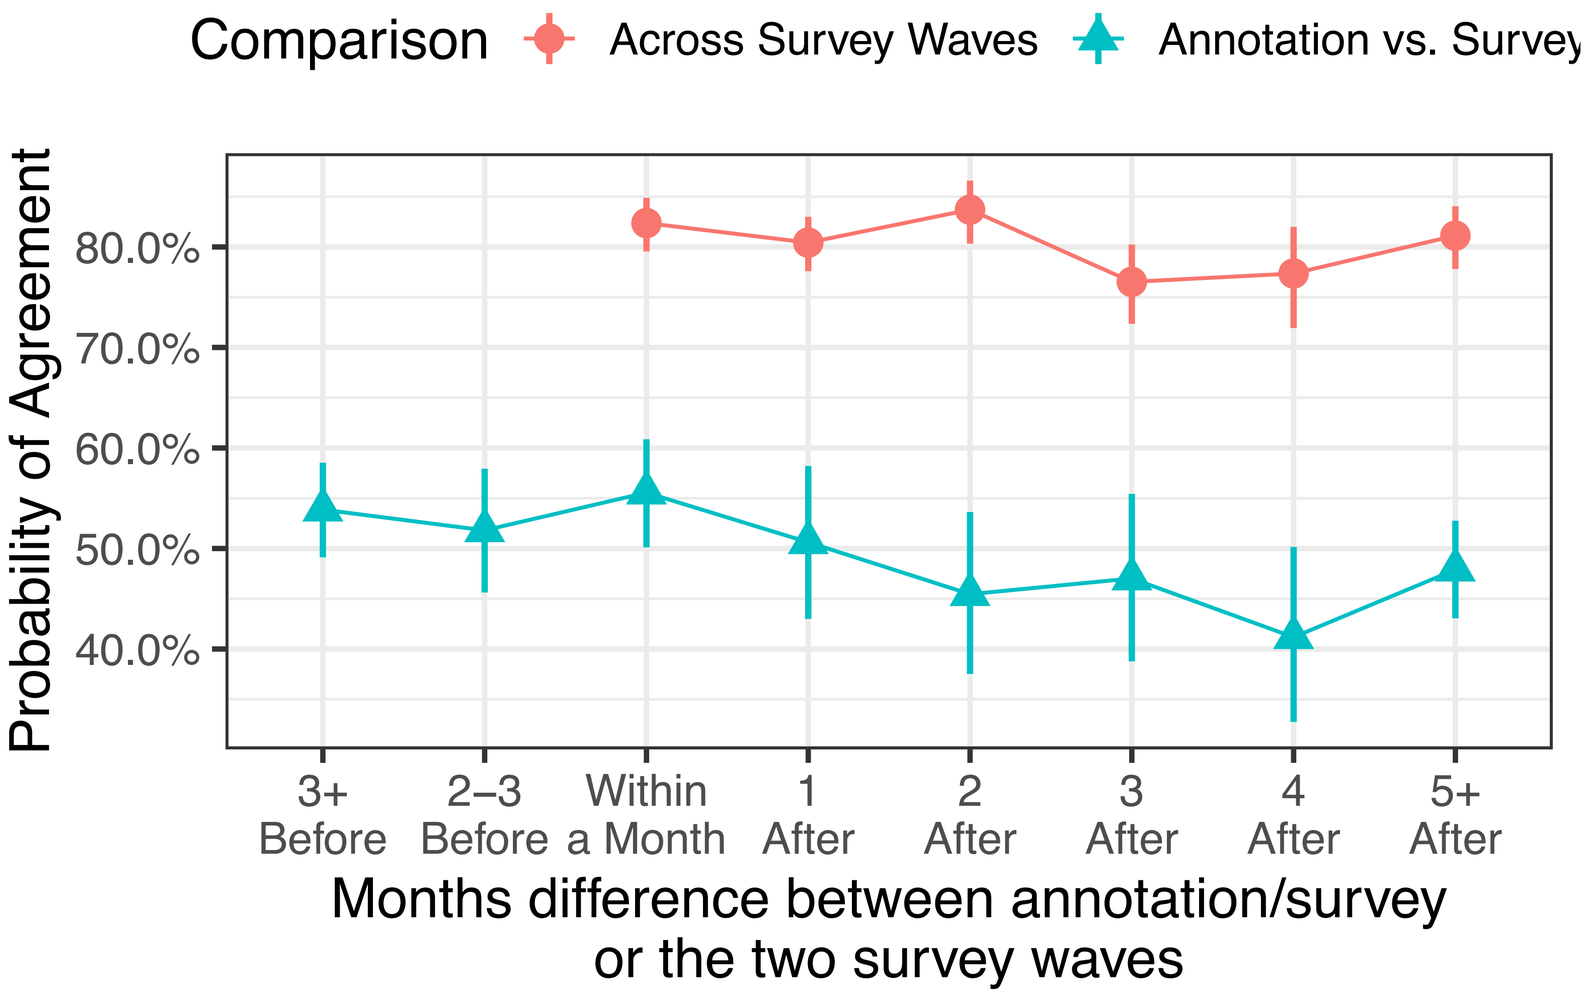

Take away:
Surveys and social media may be measuring fundamentally different things

Summary

Summary

R1: When do annotators agree on stance?

Agreement high when at least one annotator was *Very Confident*, low when both annotators were *Somewhat* or *Not at all Confident*. **"Stance" on social media is conceptually fraught.**

Summary

R1: When do annotators agree on stance?
Agreement high when at least one annotator was *Very Confident*, low when both annotators were *Somewhat* or *Not at all Confident*. **"Stance" on social media is conceptually fraught.**

R2: Does "stance" map to survey responses?
Non-Neutral survey responses often matched tweet annotations. But only when tweets contained enough information for annotators to agree. **People don't post their neutral opinions and posts don't always reflect *policy* positions.**

Summary

R1: When do annotators agree on stance?
Agreement high when at least one annotator was *Very Confident*, low when both annotators were *Somewhat* or *Not at all Confident*. **"Stance" on social media is conceptually fraught.**

R2: Does "stance" map to survey responses?
Non-Neutral survey responses often matched tweet annotations. But only when tweets contained enough information for annotators to agree. **People don't post their neutral opinions and posts don't always reflect *policy* positions.**

Social media and surveys may be measuring fundamentally different things:

Summary

R1: When do annotators agree on stance?
Agreement high when at least one annotator was *Very Confident*, low when both annotators were *Somewhat* or *Not at all Confident*. **"Stance" on social media is conceptually fraught.**

R2: Does "stance" map to survey responses?
Non-Neutral survey responses often matched tweet annotations. But only when tweets contained enough information for annotators to agree. **People don't post their neutral opinions and posts don't always reflect *policy* positions.**

Social media and surveys may be measuring fundamentally different things:

- Survey responses capture fine-grained opinion at multiple time points

Summary

R1: When do annotators agree on stance?
Agreement high when at least one annotator was *Very Confident*, low when both annotators were *Somewhat* or *Not at all Confident*. **"Stance" on social media is conceptually fraught.**

R2: Does "stance" map to survey responses?
Non-Neutral survey responses often matched tweet annotations. But only when tweets contained enough information for annotators to agree. **People don't post their neutral opinions and posts don't always reflect *policy* positions.**

Social media and surveys may be measuring fundamentally different things:

- Survey responses capture fine-grained opinion at multiple time points
- Social media captures trending events and vocal opinions

Thank you!

- "Stance" on social media is conceptually fraught.
- People don't post their neutral opinions and posts don't always reflect *policy* positions.
- Social media and surveys may be measuring fundamentally different things:
  - Survey responses capture fine-grained opinion at multiple time points
  - Social media captures trending events and vocal opinions

Sarah Shugars
Assistant Professor,
Rutgers University

they/them/theirs
sarah.shugars@rutgers.edu

RUTGERS
THE STATE UNIVERSITY
OF NEW JERSEY

Survey Questions

Trump

We used two survey questions to assess stance towards Donald Trump, one on voting intentions, and one on vote choice asked afterwards.

We then assigned survey stance as follows:

- Anti-Trump: said they would/did vote for Biden (60% of respondents).
- Pro-Trump: said they would/did vote for Trump (32%).
- Neutral: said they were unsure of who they did or were going to vote for (8%).

Survey Questions

Masks

We used two survey questions to determine stance towards masks, one that gauged perceptions of effectiveness, and the other on mask-wearing behavior.

We then assigned survey stance as follows:

- Anti-mask: said masks were Ineffective or said they follow mask wearing Not at all Closely (15% of respondents).
- Pro-mask: said they were Closely following mask-wearing guidelines (66%).
- Neutral: answered both questions and did not match the conditions for Anti or Pro (18%).

Survey Questions

Lockdowns

We used four survey questions related to social and economic restrictions on 1) leaving the home, 2) business closings, 3) cancelling large events, and 4) closing restaurants to assess stance towards lockdowns. All questions were asked on a four-point Likert scale, ranging from Strongly Disapprove (1) to Strongly Approve (4). For each respondent, we averaged the Likert scale values of their answers and used this value to assign survey stance as follows:

- Anti-lockdown: had an average Likert score of Strongly Disagree (i.e. sum of 4-6) (3% of respondents).
- Pro-lockdown: had an average score of Strongly Agree (14-16) (60%).
- Neutral: had an average score between Strongly Disagree and Strongly Agree (7-13) (37%).

Survey Questions

Vaccines

We used two survey questions to compute stance towards COVID-19 vaccines, one on whether or not they had gotten the vaccine, and if not, what their intentions were.

We then assigned survey stance as follows:

- Anti-vaccine: Stated they were Extremely Unlikely to be vaccinated (14% of respondents).
- Pro-vaccine: Had already been vaccinated, or were Extremely Likely to be vaccinated (43%).
- Neutral: Stated they were Somewhat Likely, Neither Likely or Unlikely, or Somewhat Unlikely to be vaccinated (44%).

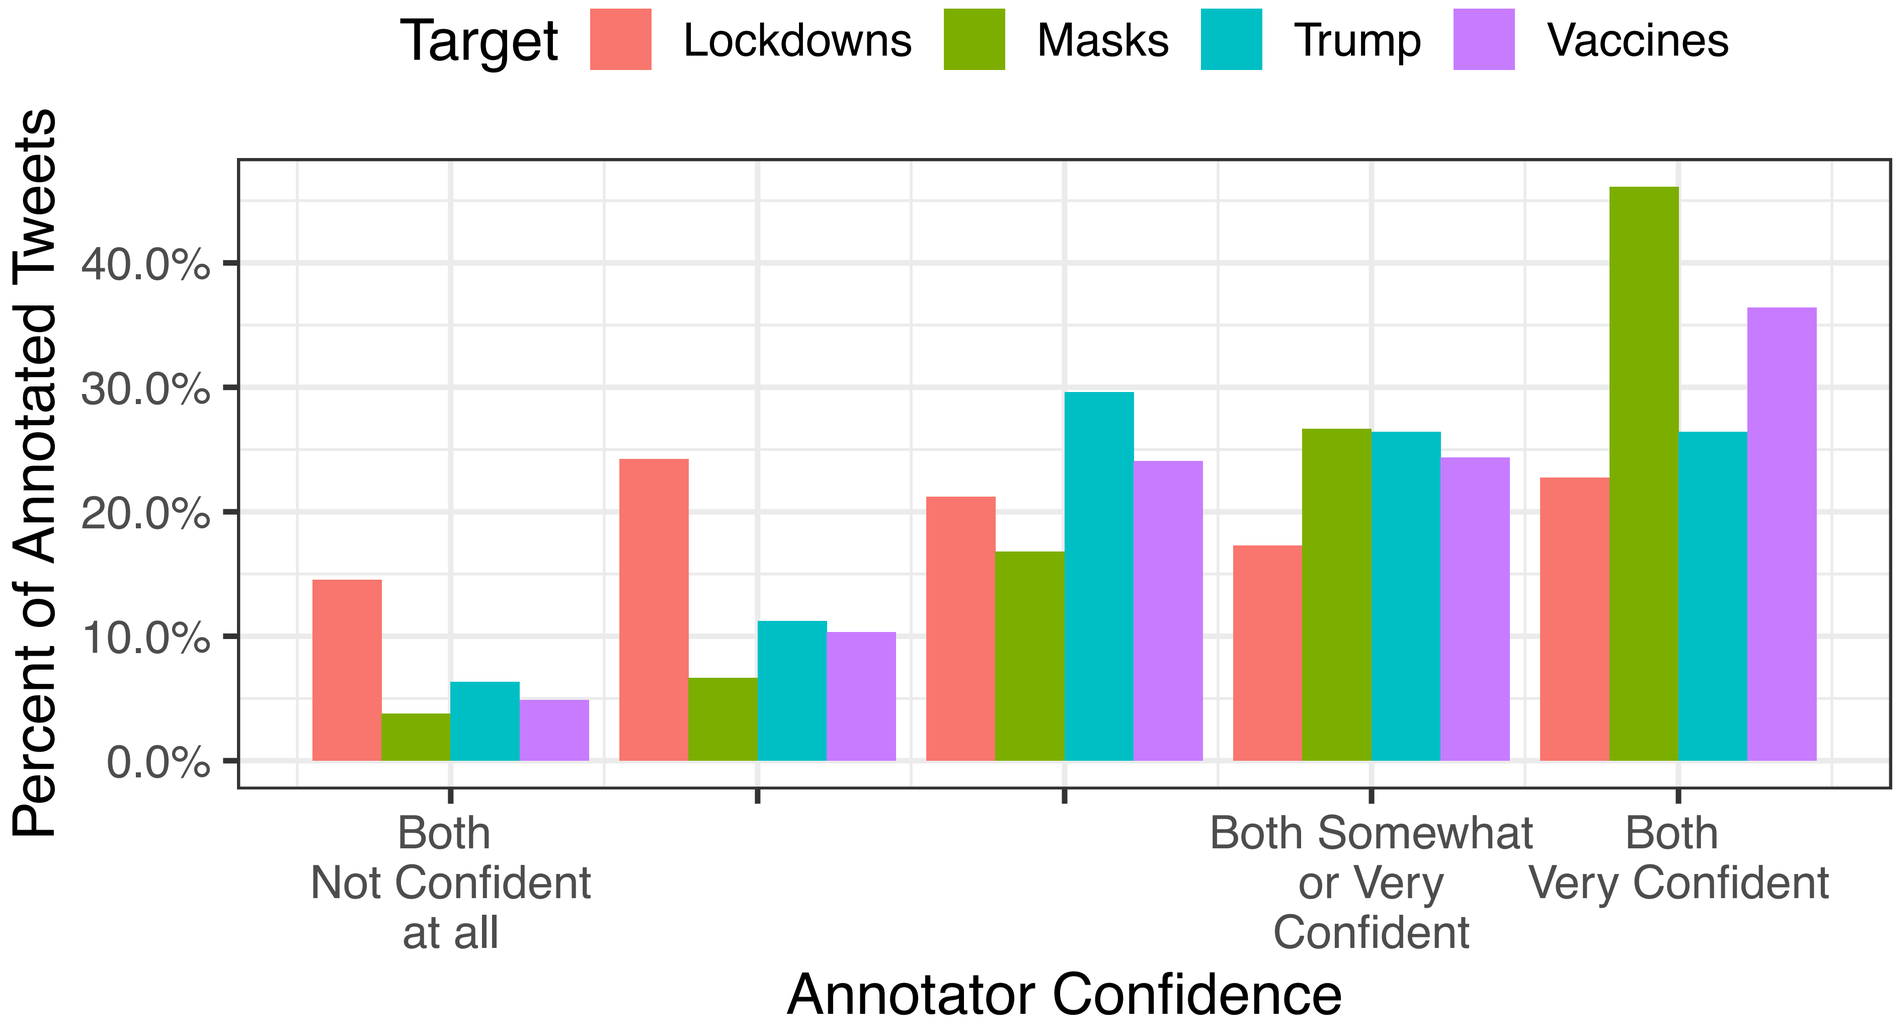
Target
Lockdowns
Masks
Trump
Vaccines
Percent of Annotated Tweets
40.0%
30.0%
20.0%
10.0%
0.0%
Both Not Confident at all
Both Somewhat or Very Confident
Both Very Confident
Annotator Confidence

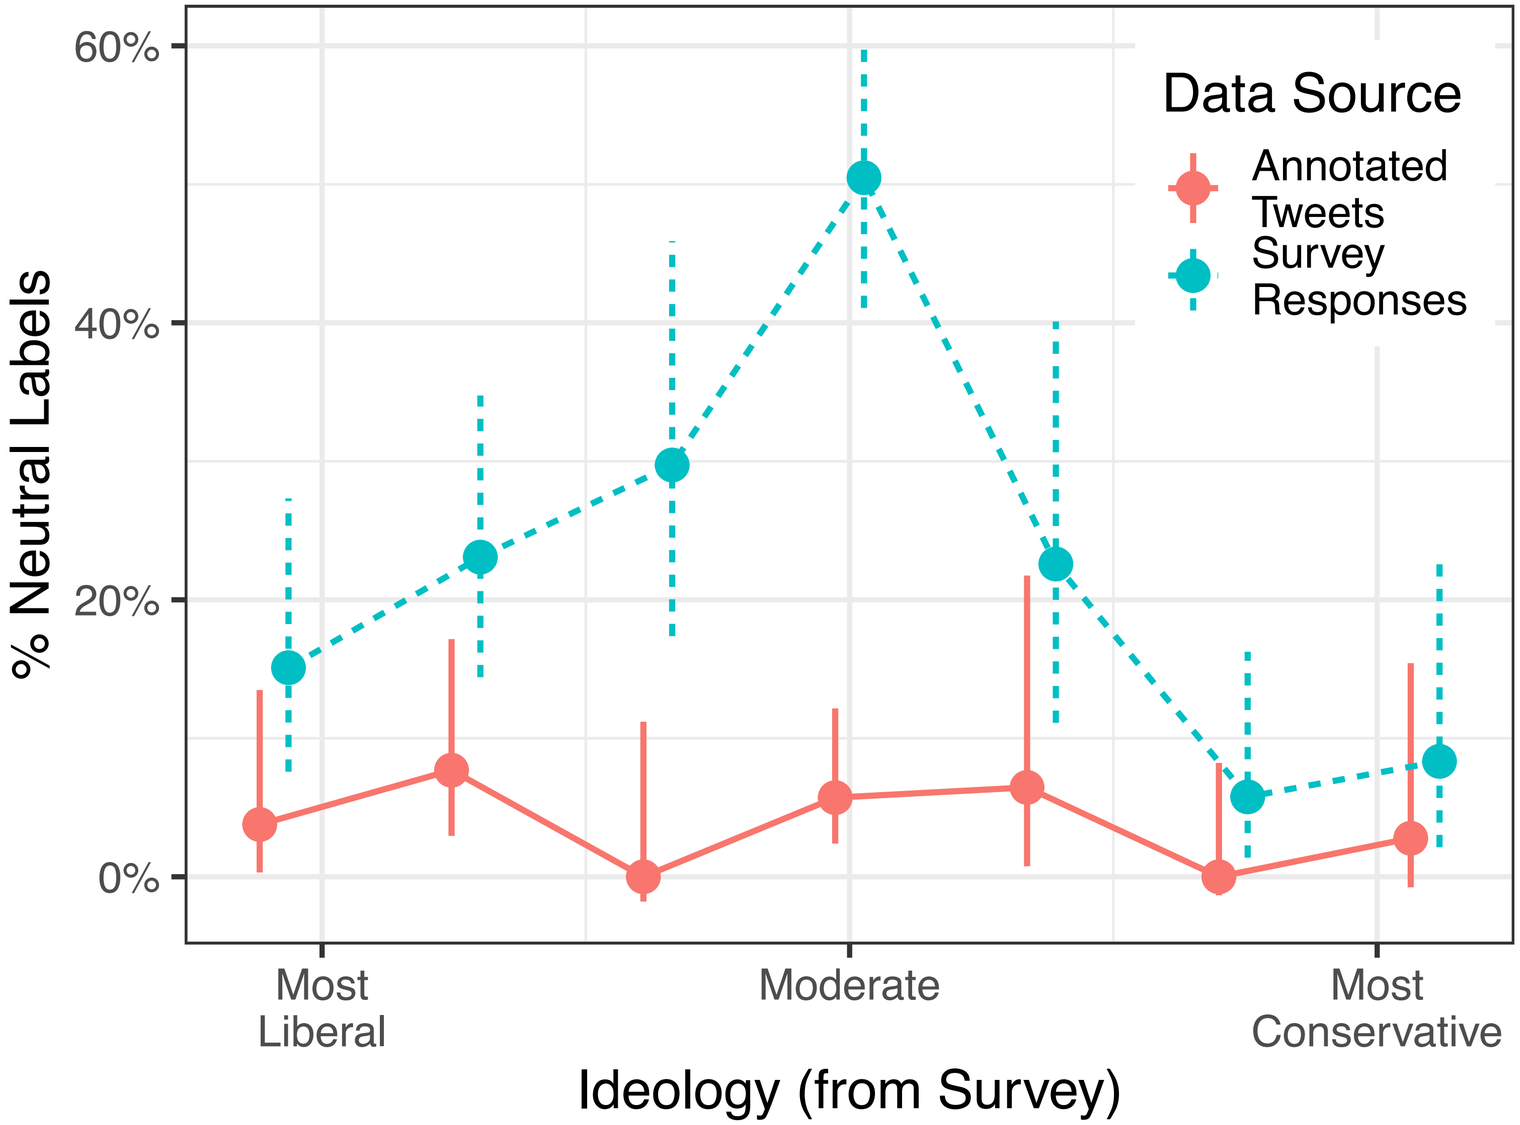
60%
40%
20%
0%
% Neutral Labels
Data Source
Annotated Tweets
Survey Responses
Most Liberal
Moderate
Most Conservative
Ideology (from Survey)