T»VISION

Understanding co-viewing measurement

Why co-viewing accuracy will define 2025

Understanding co-viewing measurement

Table of contents

Why co-viewing accuracy will define 2025

BY KELLY ABCARIAN
Chief Strategy Officer | Matter More Media

In a media economy built on precision, we've come to expect that every impression is traceable, every audience measurable, and every dollar accountable. But there's a flaw — quiet, yet consequential — that threatens to undermine all three: how co-viewing is counted.

At its core, co-viewing is supposed to tell us how many people are truly in the room — watching, engaging, and sharing the ad experience together. But here's the truth: co-viewing numbers aren't directly observed. They're modeled. And those models vary — by platform, by provider, even by program.

That means the same household, watching the same show, can be counted as 1.5 viewers on one platform and just 1.2 on another. Not because behavior changed, but because the math did.

This isn't a back-end technicality. It's a front-line issue:

- Co-viewing models can overstate how many unique people are watching — leading to inflated reach.
- They can also understate how often viewers see ads — resulting in lower frequency.
- Publishers may be undervalued for shared-viewing content — depending on the platform, or even the specific program.
- Agencies are stuck navigating inconsistent data across platforms, making it difficult to understand the real value of the audience.
- And the entire market risks eroding trust in person-based currency.

Nowhere is this more urgent than in live sports — where co-viewing is highest, stakes are steep, and advertisers are paying a premium for audience impact. If we want to know the true value of these audiences, we have to get co-viewing right.

Today, advertisers are being asked to plan and buy media using numbers that simply aren't comparable. One co-viewing model might inflate reach on a streaming platform, while another undervalues shared viewing on linear — turning what should be a single source of truth into a patchwork of Frankenmetrics. Measurement partners must do more than model audiences — they must model them consistently across platforms and publishers and programs. Because if the math is broken, so is the value.

While some vendors tout transparency, many hide behind MRC accreditation — as if following your own rules is the same as those rules being statistically sound or market-ready. It isn't. The MRC validates process, not methodological accuracy. That's why I believe co-viewing accuracy must become a front-line issue in 2025. Because measurement isn't just data — it's the foundation for how media is priced, planned, and trusted.

This paper is a call to action: it's time to stop treating co-viewing as a black box. Let's demand clarity, challenge assumptions, and create a more accountable marketplace — together.

Defining co-viewing

Let's begin by understanding what we mean when we talk about co-viewing. Simply co-viewing is the number of people in the room with the opportunity to see the content on the screen.

Co-viewing varies greatly across dayparts, genres, apps and networks and programs. Understanding true co-viewing is important for understanding true audience reach.

When counted accurately, co-viewing:

- Increases reach
  (more real viewers per screen)
- Reduces CPM
  (more efficient cost per person)
- Shapes frequency models
  (shared exposure alters repetition value)
- Impacts attribution and ROI
  (by tying viewership to outcomes)

There are hundreds of apps and networks with a co-viewing rate between 1 and 2.5, with many having co-viewing rates between 1.25 and 1.5. Likewise, there are over 750 programs with a co-viewing rate between 1.30 and 1.36. **In the charts to the right, we have highlighted a few example apps, genres and programs.**

There is a large variance in co-viewing across apps, genres and programs

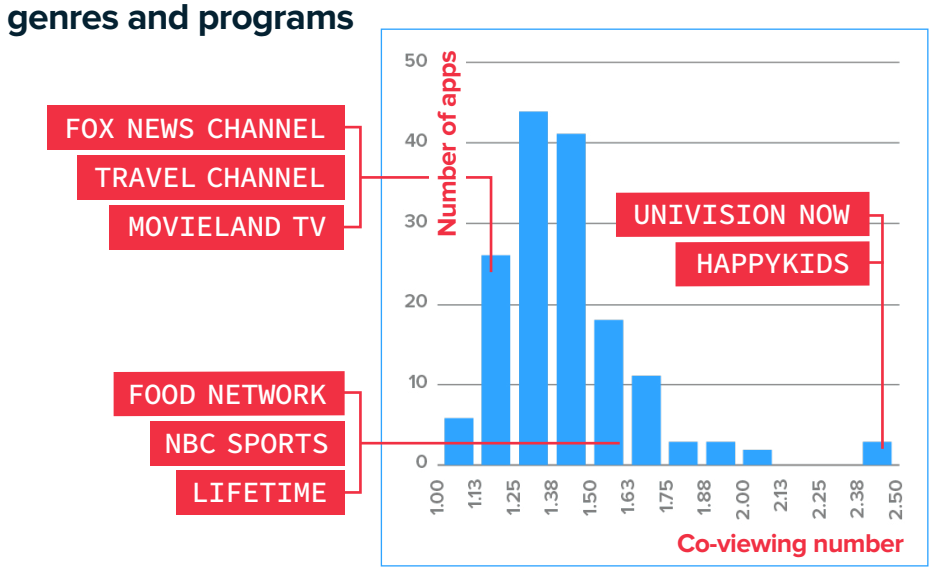

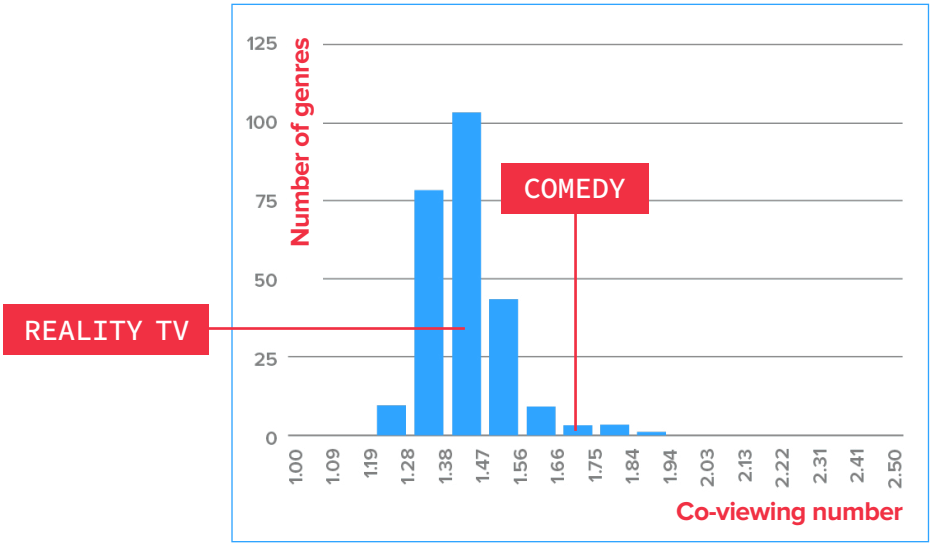

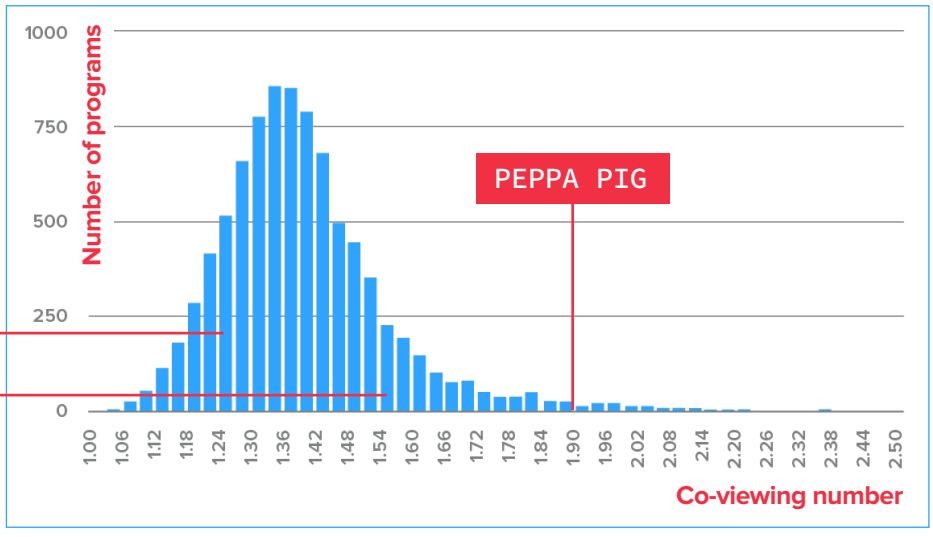

Source: TVision panel | 2024

CHAPTER ONE

Why does co-viewing matter?

Co-viewing is when two or more people watch the same piece of content together, usually on a shared screen like a living room TV. In a world of person-based media buying, it's not just about what is on — it's about who is actually watching.

It sounds simple, but in the world of ad measurement, it's a major blind spot — one that quietly affects everything from how much reach you think you're getting, to how much you pay, to whether your campaigns are truly delivering on target.

Here's the catch:
All co-viewing numbers are modeled — even those that are weighted and projected to a universe. And every measurement provider uses a different formula. In fact, modeling approaches can vary even within the same measurement provider, depending on the publisher, platform, or program. That means the same show, viewed by the same household, can report completely different audience sizes across platforms — not because viewer behavior changed, but because the methodology did.

And here's what we've learned:
Traditional co-viewing estimates have long relied on presence-based assumptions — namely, that if someone is in the room, they are watching or listening.

It's important to note that the definition of co-viewing used in this report aligns with current industry currency standards — where co-viewing is measured by the average number of concurrent active viewers. A viewer is considered active from the moment they enter the room until they've been absent for 10 or more minutes, with brief absences included. This methodology mirrors legacy panel behaviors, like button-pushing, and remains the basis for how co-viewing is incorporated into most currency systems today.

However, this approach is rooted in presence, not proof of attention.

Thanks to advancements in passive measurement, it's now possible to go further — to measure actual eyes-on-screen. That's exactly what companies and advertisers are doing using TVision's attention data, which captures who is truly watching. This level of precision is more than just a technical upgrade — it has been shown to correlate strongly with higher recall, stronger emotional resonance, and greater intent to act.

As advertisers demand greater accountability for every impression, more are turning to attention-based insights to supplement currency data, close the gap between exposure and effectiveness, and ensure their campaigns are built not just on presence — but on impact.

Co-viewing estimates directly impact:

- **Reach and frequency —** How many actual people saw your ad?
- **CPMs and budget efficiency —** Are you paying for impressions that never existed?
- **Campaign guarantees —** Did your media buy really deliver on its audience promise?
- **Cross-platform comparisons —** Are you comparing apples to apples... or apples to algorithms?

EXAMPLE

Co-viewing inaccuracy creates waste

A national CPG brand launched a cross-platform campaign targeting young families, running across live sports (linear TV), bingeable CTV streaming and YouTube.

The media plan assumed:

- 1.4x co-viewing on linear (panel data)
- 1.6x on streaming (first-party logs)
- 1.7x on YouTube (survey data)

But in reality:

- Streaming was mostly solo viewing — true co-viewing was closer to 1.1x.
- Linear TV (live sports) was under-reported due to poor metering compliance among younger viewers.
- Streaming co-viewing was adjusted using logs, but linear programming wasn't modeled the same way — skewing cross-platform comparisons.
- YouTube survey data overstated shared viewing without behavioral validation.

THE RESULT:

It's possible that the plan overestimated reach by 30%. Frequency caps misfired.

Takeaway

Inconsistent co-viewing measurement kills precision — and wastes media dollars.

CHAPTER TWO

Why delivery format matters

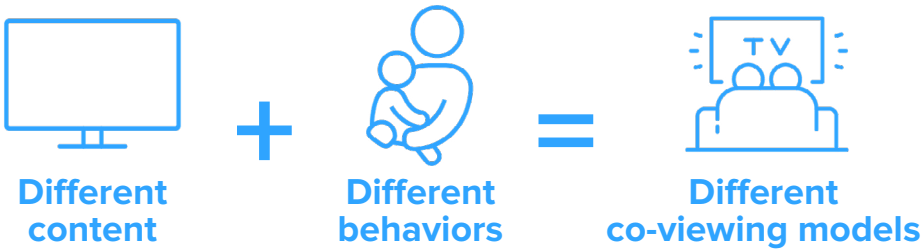

Linear-only (live, DVR, on demand):

Large-screen, living room, communal

RISKS:
Panel fatigue, under-reporting younger viewers, and modeling based on limited, highly compliant sample groups

On-demand streaming (UGC and professionally produced):

Time-shifted, bingeable, often solo

RISKS:
Repurposed linear co-viewing rates and reliance on self-reported surveys

Live sports streaming:

High co-viewing potential but on fragmented devices and platforms

RISKS:
Scaled-up modeling using first-party logs and uniform demo uplifts — inflating audiences

Hybrid simulcast (linear + streaming):

One event, two platforms, two measurement methodologies

RISKS:
Different methodologies produce two different truths for the same program

Key insight

If the same show is modeled differently across delivery platforms, the market can't rely on those comparisons to make meaningful decisions.

CHAPTER THREE

How a co-viewing number is made

Not all co-viewing numbers are created equal — even from the same measurement provider

Myth: “one source = one standard”

What people think:

“If I use one measurement provider, all my co-viewing numbers must be built the same way, so they're apples to apples.”

Reality:

Even within the same measurement provider, co-viewing numbers can vary dramatically because the underlying inputs, models, and assumptions change depending on:

- The platform (linear, streaming, UGC)
- The content type (sports, sitcoms, news)
- The delivery format (live, time-shifted, simulcast, on-demand)
- The publisher (based on data availability and business rules)

Same sport, different numbers

Provider	Platform	Program	Co-viewer multiplier	Data source	Modeling method
Provider X	Linear TV	Live sports broadcast	1.4x	Button-push or passive panel	Donor model
Provider X	Streaming (CTV)	Live sports broadcast	1.9x	Button-push or passive panel + first party server logs	Behavioral weighting, uniform demo scaling
Provider X	YouTube	Highlight clips	1.7x	Self-reported survey	Scaled UGC uplift

Source: Indicative co-viewing averages based on various methodological approaches used in linear and streaming audience measurement.

Co-viewing numbers aren't directly counted — they are built through a series of steps using different types of data, assumptions, and math. Here is a simplified look at how it works:

STEP 1 Where the data comes from

To estimate how many people were watching, measurement providers start with different types of data, like:

- **People panels** — real households that log who's watching (either by pressing a button or being passively detected by sensors or cameras)
- **Device data** — first-party server logs from streaming apps that show when and where content played and for how long, but not necessarily who watched
- **Self-reported surveys** — people self-report what and how they watch, highly subjective (common on user generated platforms)
- **Demographic databases** — third party data that estimates age, gender, or household makeup, but isn't tied to behavior in real time

STEP 2 How the numbers are estimated

Next, measurement providers use models to fill in the gaps and figure out who was likely watching. These models can vary widely:

1. **"Lookalike" households** — Data from one household is used to fill in the gaps for another, based on similar traits. This can lead to blind spots and behavioral bias if too few homes are used or the same home is re-used too often.
2. **Probability models** — Statistical tools estimate who's most likely watching at a given moment, based on habits.
3. **Behavior-based models** — AI tools match viewing behavior with demographics to predict who was likely in the room — but these need a lot of high-quality data to be accurate.

STEP 3 How the numbers are scaled up or calibrated

Once estimates are made, providers often adjust or "calibrate" the numbers to align with broader population patterns. This might include:

- Blending in third party identity data (like from TransUnion or Experian) and comparing it against smaller panel data to identify and correct for biases and gaps
- Using calibration factors to adjust big data to better align with smaller panel findings
- Finding a single home within the panel that shares similar characteristics, then using that home to "donate" its viewing probabilities to other homes where demographic data is missing
- Bringing in first party survey data to calibrate and/or report directly co-viewing habits for UGC content

This lack of methodological alignment leads to:

- Artificial inflation or deflation of audience counts
- Platform-based inconsistencies (streaming vs. linear vs. UGC)
- Inaccurate reach and frequency capping

Bottom line:

Co-viewing is one of the few areas in modern marketing where assumption still outweighs observation — and the risks are accelerating.

CHAPTER FOUR

The limitations of active metering —

And its impact on co-viewing accuracy

In an era where media investment is increasingly precision-driven, the systems used to measure those audiences must match that precision. And yet, many co-viewing estimates in use today still rely on active metering — a legacy approach that requires individuals to physically “check in” during TV viewing by pushing buttons, using secondary remotes, or other forms of manual compliance.

This manual approach may have sufficed in an earlier era of TV consumption, but today, it creates a fragile, error-prone measurement system that cannot keep pace with the complexities of modern viewing behavior — especially in shared-viewing environments.

Why active metering falls short

Active metering introduces three core vulnerabilities:

1. **Button fatigue:**

   This creates net false absences where co-viewing did occur — but goes unrecorded.

2. **Under-representation of younger viewers:**

   Millennials and Gen Z are significantly less likely to follow compliance protocols. As a result, co-viewing among younger demos is frequently under-reported — skewing media plans targeting these valuable audiences.

3. **No passive correction layer:**

   When a panelist forgets to log in, the system assumes no one was present. Unlike passive systems, there is no behavioral backup to verify or correct that assumption.

EVIDENCE

Understanding the issue of age-based bias in active audience measurement panels

To quantify the inaccuracy of active metering, **TVision and the Coalition for Innovative Media Measurement (CIMM)** conducted a head-to-head simulation, comparing traditional active metering to passive, second-by-second person-level observation in real homes.

KEY FINDINGS:

1 Massive reporting inaccuracy

- Active and passive measurement aligned only 56% of the time when the TV was on.
- That means **44% of co-viewing time** was either *missed or overcounted.*
- This directly affects impression counts, CPMs, reach calculations, and campaign guarantees.

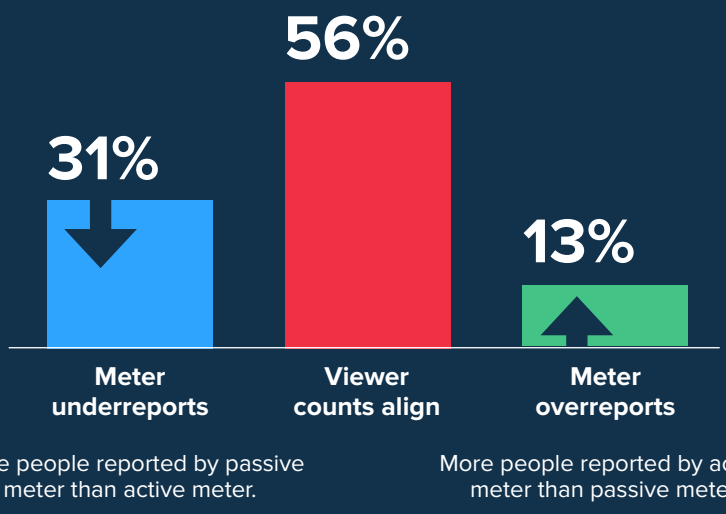

2 Systemic under-reporting among younger demos

- Viewers **aged 18–34** were **7.9x more likely to under-report** their viewing than to over-report.
- This creates a persistent undercounting of a high-value, media-hungry audience segment.
- Measurement systems relying on active compliance **misrepresent the performance** of Gen Z and millennial-targeted content.

3 Rapid fatigue in panel engagement

- Within just four weeks, non-compliance rose by **50%**, as panelists disengaged from required button-pushing.
- This leads to what researchers call “compliance drift,” where over time, the data becomes increasingly unreliable.

4 Daypart-specific distortion

- Under-reporting wasn’t evenly distributed throughout the day.
- Morning and late-night viewing were the most affected — both critical for understanding family routines, time-shifted content, and solo-device behavior.

CIMM report | “Understanding the issue of age-based bias in active audience measurement panels”

Why CMOs should care

PHANTOM REACH = ILLUSION OF SCALE

Co-viewing inflation isn't new — it's a modern version of the old "average viewers per set" problem. When co-viewing is overstated, your reach appears higher than it really is, leading to misleading ROI and media effectiveness metrics.

You're paying for impressions that may never have existed — and drawing strategy from inflated results.

MISMATCHED MEASUREMENT = UNRELIABLE PLATFORM COMPARISONS

Streaming platforms often use modeled co-viewing based on behavioral signals, while linear TV relies on meters or panel-based estimates. Comparing the two directly creates a false sense of equivalence.

This undermines your media mix decisions and masks which platforms are truly delivering value.

UNDERCOUNTING VIEWERS = MISSED AUDIENCES AND FATIGUED ONES

When co-viewing is underestimated, reports may show audience goals were met — even if critical demos were missed. Meanwhile, those who were reached may have seen your ad too many times.

This is a dual threat: underperformance goes unnoticed, and frequency balloons where you least expect it.

BROKEN FREQUENCY CONTROL = BRAND RISK

Frequency caps set at the stream level break down when the true number of viewers per stream is miscounted. You risk overexposing real people to the same ad over and over — without realizing it.

This leads to wasted impressions, rising costs, and consumer annoyance — all while your reports suggest control.

Key takeaway: passive data powers smarter investment

You can't optimize what you can't see.

Accurate co-viewing is essential for cross-platform planning and ROI.

CHAPTER FIVE

Why co-viewing accuracy matters for everyone

For CMOs, brands and agencies:

- Overstated reach = wasted spend
- Cross-platform comparability collapses, affecting pricing
- Media mix modeling degrades

For publishers:

- Under-credited impressions mean lost revenue
- No parity between linear and streaming valuations
- Younger audiences are being undercounted, causing advertisers to move money out of TV—even when those viewers are actually watching

For the industry:

- Frankenmetrics (different methodologies by platform) destroy confidence
- Persons-based currency becomes compromised
- Reach, frequency, and CPMs are modeled on incomplete and non-comparable truths across platforms

Bottom line

Co-viewing is not a back-end issue. It's a pricing issue, a planning issue, and a fairness issue.

CHAPTER SIX

Five questions every CMO should ask about co-viewing

Co-viewing may seem like a technical metric buried deep in measurement methodology — but it directly shapes pricing, planning, frequency, and media investment outcomes. In a cross-platform world of persons based currency, every advertiser, agency and publisher must ensure their teams understand how co-viewing is calculated and when it may be inflating or undercutting audience estimates. Here are five critical questions everyone should ask their measurement provider — and what to watch out for.

1 Is your co-viewing data observed, modeled, or self-reported? What data sources were available for this publisher or platform?

Why it matters:

If your vendor uses different methods across platforms (e.g., first party server logs for one and self reported surveys for another), the result is a non-comparable persons count by platform or publisher or individual program. This makes reach and frequency calculations unstable and not comparable across your media plan.

What to listen for:

- The share of co-viewing from passive observation, modeled inference, and self-reporting
- How inputs differ between linear and streaming and UGC environments
- Whether methods shift based on the platform or data availability
- Validation or calibration steps across datasets

RED FLAG

"We use a mix depending on the platform."

— with no clear explanation.

This often signals **methodological inconsistency** that distorts outcomes across your plan.

Five questions every CMO should ask about co-viewing

2 Do your co-viewing assumptions change based on content type or platform (e.g., live sports vs. sitcoms, linear vs. CTV)?

Why it matters:

If the same or similar co-viewing factors exist regardless of what type of content was consumed on a given day for a given platform and across a campaign, your audience estimates will be wrong.

What to listen for:

- Differences in co-viewing treatment by genre or daypart or platform
- Whether live sports, kids programming, and reality TV are modeled separately
- Adjustments by screen size, device type, or location (family room vs. bedroom)

RED FLAG

"That's just how the model works."

— Similar co-viewing rates across all content often mean the model can't detect real behavioral differences — and is likely built on flawed assumptions.

Five questions every CMO should ask about co-viewing

3 How are person-level impressions scaled from household-level data — and do the inputs or modeling approach vary by content, platform, or publisher?

Why it matters:

Some platforms benefit from program-specific uplifts based on first-party data, while others don't. If your household-to-person scaling changes based on data access, it creates a biased playing field.

What to listen for:

- Whether person-level scaling varies by program, partner, or data availability
- How demographic panel data is reweighted with behavioral models
- If scaling is uniform across demographics or adjusted for bias

RED FLAG

"It depends on the program" or "We optimize using available data."

— signals **selective inflation** is being applied, distorting media value and comparability.

Five questions every CMO should ask about co-viewing

4 How often is the same household (donor home) used to model person-level viewing across other homes — and what prevents overuse?

Why it matters:

Overusing the same donor household leads to systemic bias. If the same "proxy" is used repeatedly, it assumes millions of homes watch like one — they don't.

What to listen for:

- Frequency of donor reuse or rotation
- Criteria for donor matching (e.g., just age/gender or also behavior, such as streaming or linear behavior)
- Whether donor usage is audited or validated regularly

RED FLAG

"We don't track or report donor reuse."

— means they're **flying blind on systemic model bias** and you may be buying misattributed viewers.

Five questions every CMO should ask about co-viewing

5 Do you publish your co-viewing methodology, data sources, and error margins — and are these independently validated?

Why it matters:

If your provider can't show their math, you shouldn't trust the outcome. MRC accreditation validates that vendors follow a process — but does not dictate methodology or ensure consistency across platforms.

What to listen for:

- Public methodology guides by platform
- Published confidence intervals or variance ranges
- Independent validation (e.g., by TVision, CIMM, ARF etc.)

RED FLAG

"We're MRC-accredited."

— if said without any published methodology, this means they may be **using process validation to mask outcome inconsistency.**

Bottom line for CMOs

If your measurement partner can't answer these five questions with clarity and documentation, it's time to dig deeper. Don't accept vague or generalized answers — push for transparency and platform-agnostic consistency. Your media dollars — and your trust — depend on it.

CHAPTER SEVEN

The path forward: building a modern, standardized co-viewing framework

To ensure accuracy, equity, and accountability, the industry must adopt a new baseline for co-viewing measurement:

1 Passive, person level observation

- Automated sensors, computer vision, or behavioral signals
- No reliance on manual entry or participant compliance

2 Cross platform methodology alignment

- Consistent logic applied across streaming, linear, UGC and hybrid formats

3 Granular calibration models based on behaviors

- Models that standardize on how to take in differentiation by genre, daypart, screen type, and viewing context

4 Transparent documentation

- Publicly available modeling logic, assumptions, and error margins

5 Independent validation

- Auditing and validation by neutral third parties like TVision, not just internal or MRC process compliance

CHAPTER EIGHT

Final call to action

For CMOs and brands:

- Demand platform-agnostic co-viewing methods
- Don’t trust black boxes — ask how numbers are made
- Reject co-viewing Frankenmetrics

For agencies:

- Run A/B tests against passive alternatives like TVision
- Push for modeling parity across platforms
- Validate reach/frequency with passive inputs

For publishers:

- Advocate for parity in linear vs. streaming vs hybrid vs UGC methodology
- Quantify and claim your co-viewing value using third party passive data like TVision
- Don't accept under crediting due to outdated methodologies

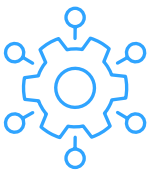

For measurement providers:

- Publish methodology, scaling logic, and error margins
- Stop selectively boosting streaming first programs with first-party data
- Evolve panels to include passive, person-level data

The future is transparent

Person-based buying depends on knowing who saw what. Let's move past process checkboxes and stop modeling shared attention — and start measuring it.

If the industry is serious about fair, platform-neutral, person-level media currency, then co-viewing must evolve: It's time to replace proxies with proof.

It's time to measure how people actually watch — together.

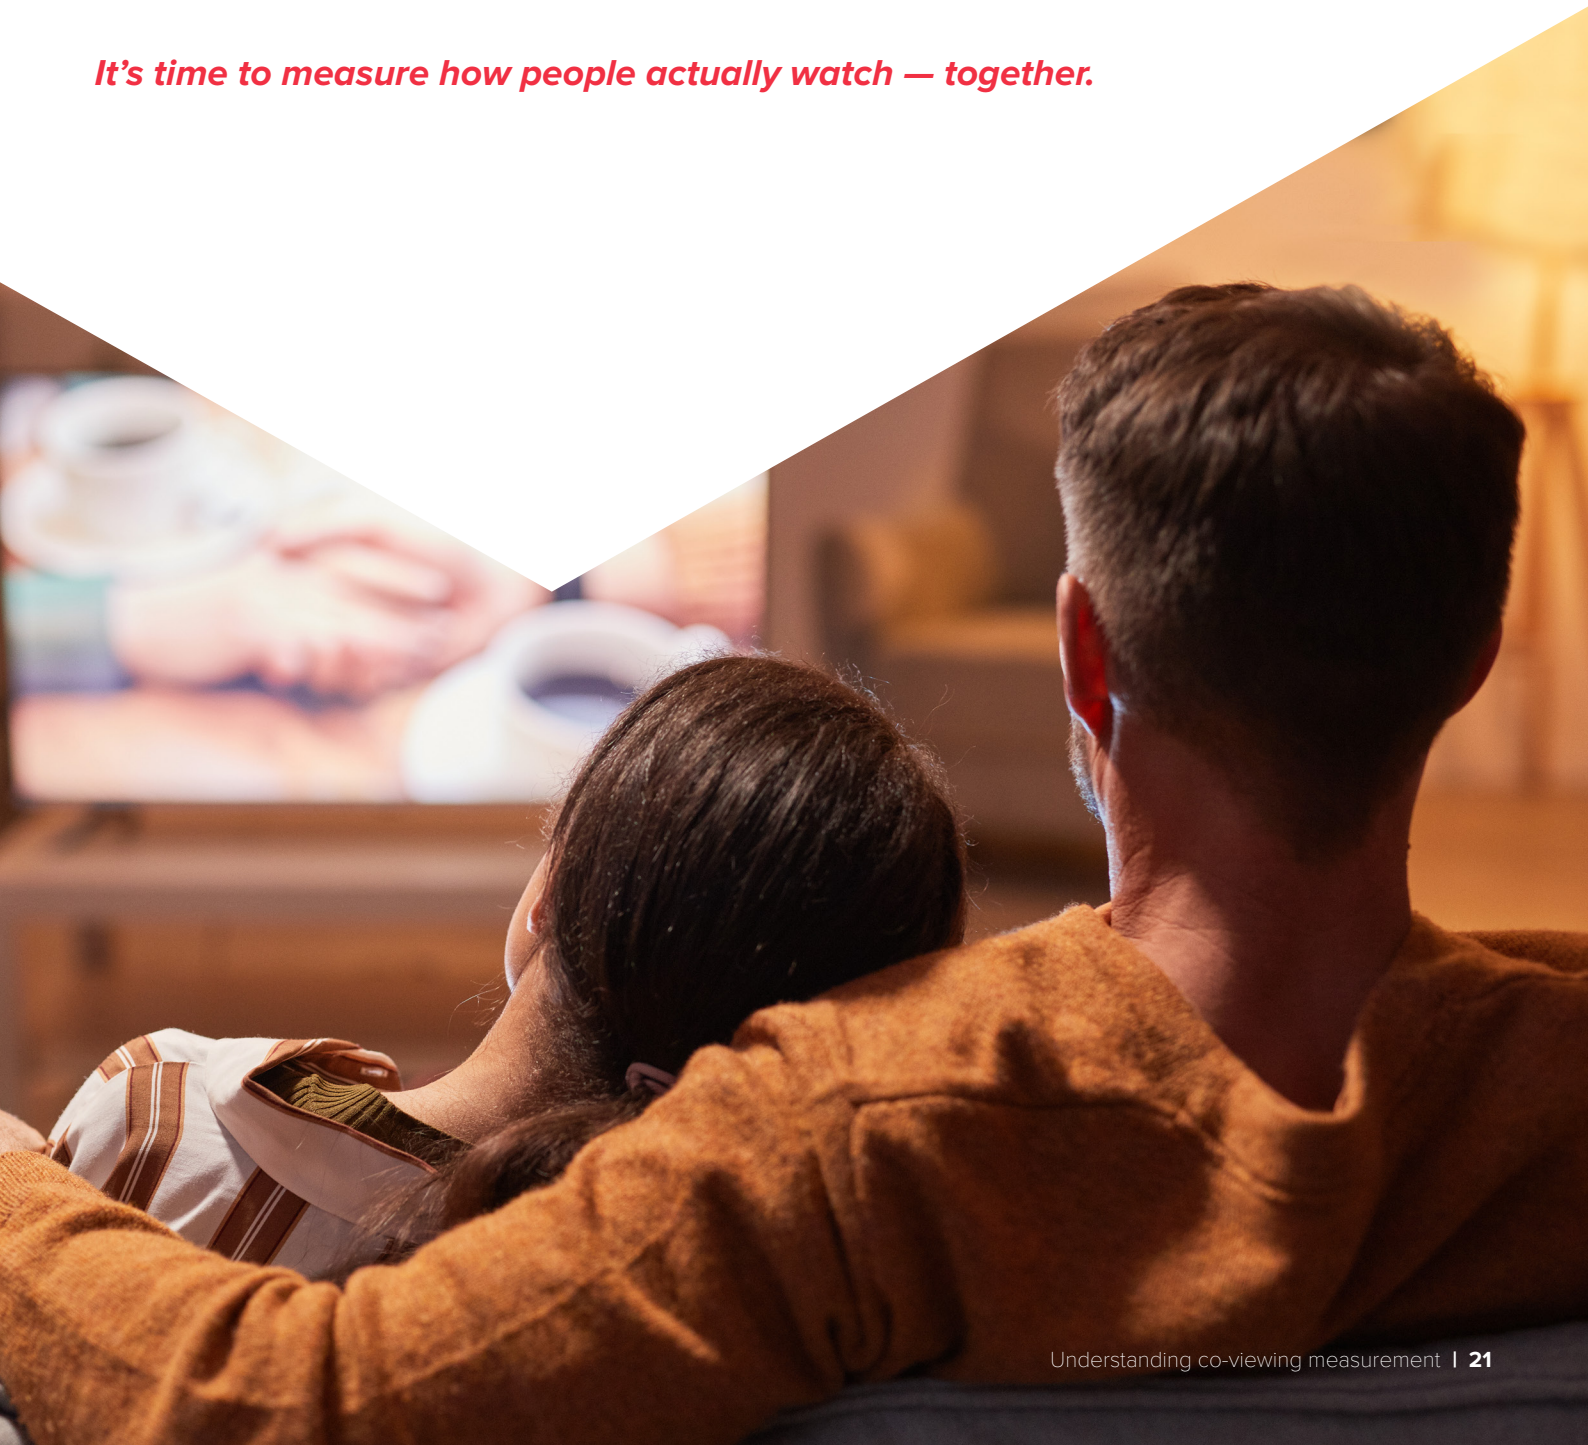

Contact us

Learn more from TVision

info@tvisioninsights.com

T»VISION