Consensus

# Debunking social media claims

A one week assignment. Students take one health claim off social media, find out whether it is actually true, and write a verdict they can defend.

## INSPIRED BY

This assignment is inspired by an assistant professor of sports nutrition at Ohio University, whose students find an influencer's post and use Consensus to defend or debunk the claim against real scientific evidence. Teaching them how to trace the literature behind a broad claim and evaluate evidence.

## WHO THIS ASSIGNMENT IS FOR

This assignment is broadly applicable to any field where students encounter consumer-facing claims. The assignment moves students beyond simply doubting online information, teaching them how to distinguish between false, oversold, and true claims through the following skills:

- **Formulating Testable Questions:** Transforming vague marketing or social media claims into testable research questions by identifying the implied population, dose, and outcome.
- **Evidence Evaluation:** Ranking evidence by quality and identifying when preclinical research (e.g., cell or animal studies) is being improperly presented as applicable to humans.
- **Critical Verification:** Developing the habit of checking AI-generated summaries against the full text of source papers to identify what information was lost or flattened in the synthesis.
- **Nuanced Verdicts:** Learning to assign a verdict (Supported, Supported under conditions, Unsupported, or Untested) with appropriate confidence levels and explaining it in plain language without losing the critical qualifications.
- **Evidence Synthesis:** Building a balanced evidence set that includes supporting information—even for claims the student expects to be false—to ensure a comprehensive and fair assessment.

## IN THIS ASSIGNMENT

### Assignment Overview


01 Why this assignment exists

02 The five parts

03 How the week runs

04 Learning outcomes


### Example Submission


Part 0 The claim capture

Part 1 The search log

Part 2 The evidence set

Part 3 The discrepancy report

Part 4 The verdict

Part 5 The Case Presentation (optional)

# Assignment Overview

Students pick one health claim they have seen online. Over one week, students turn the claim into a testable question. They then search the peer-reviewed literature using Consensus, collect and read three sources, check the AI summaries against the original papers, and justify a verdict.

| Verdict | What it requires |
|---|---|
| **Supported** | Multiple human studies point the same way, at doses and in populations that resemble what the claim implies. |
| **Supported under conditions** | A real effect exists, but only in a narrower population, dose, duration, or outcome than the claim states. Students must name the conditions. |
| **Unsupported** | Human studies exist, they looked for this effect, and they did not find it. This needs evidence of absence, not just absence of evidence. |
| **Untested** | No human studies exist at the relevant dose or in the relevant population. Different from unsupported, and that difference is why there are four categories. |

Two constraints do most of the teaching. Students must find at least one source that generally supports the claim (i.e. it might be a cell or animal study that the initial claim was based on). And they must choose and defend one of the above verdicts.

## 01 Why this assignment exists

Most students already doubt what they see online. That is not the same as being able to check it and trace the evidence. Doubt without a method means they cannot tell the difference between a claim that is false, a claim that is oversold, and a claim that happens to be true. Faced with all three they either believe the confident person or dismiss everything.

Many health claims start from a small study, animal work, or cell culture, and get stretched to cover everyone. Consensus filters human studies apart from animal and in vitro work in one click, which turns an abstract warning about evidence quality into something a first year student can watch happen on their own screen. Learning to run that check is the skill.

| | |
|---|---|
| TIME | One week. Two class sessions of about 40 and 25 minutes, plus roughly 2.5 hours of student work in between. |
| COURSE FIT | Nutrition, exercise science, public health, health communication, first year seminar, general education. Any field with a consumer claim space. |
| LEVEL | First year non majors through upper division health science majors. No research methods prerequisite. |
| TOOL | Consensus (consensus.app). Check institutional access before you assign it. Free accounts have limited searches. |
| PREP | About an hour the first time. Run the searches yourself and build a list of pre approved claims. Ten minutes every term after that. |

## 02 The five parts

Each part below has a description and a worked example of what a real submission looks like. The weights are a suggestion. Customize this based on your class schedule.

| # | Part | What it is | Weight |
|---|---|---|---|
| 0 | **Claim Capture** | Six lines written before any searching. The claim word for word, who benefits, the measurable outcome, the population implied, the mechanism, and the dose. | Gate, ungraded |
| 1 | **Search Log** | Five searches in a table: query as typed, filters used, papers opened versus kept, and what they changed and why. One row must hunt for the best evidence in favour. | 15 of 100 |
| 2 | **Evidence Set** | Three peer reviewed sources. For each one: citation and working link, study type and tier, and one sentence on population, dose, and duration. | 25 of 100 |
| 3 | **Discrepancy Report** | About 250 words comparing what the AI said about one paper against the full text. One thing it got right, one thing it flattened, one practice they will adopt. Students should show their work (ie. link to the analysis, and show a marked up version of the paper). | 25 of 100 |
| 4 | **The Verdict** | About 300 words naming one of four verdicts, the evidence behind it, and what would change their mind. | 20 of 100 |
| 5 | **Group presentation (optional)** | In class in groups of three, students select one of the claims and present the actual searches they ran, and provide an overview of the evidence + explain their verdict. | 15 of 100 |

## 03 How the week runs

| When | What happens | Due |
|---|---|---|
| **Before class 1** | Students find a claim and write the Claim Capture. You approve claims. | Capture brought to class |
| **Class 1** | The Search Lab, about 40 minutes with laptops open. You run three searches live: claim to question, the evidence tier trap, and one verification pass. | Nothing. They leave with a drafted question. |
| **Days 2 to 4** | Independent work. Search, screen, read, write. Submission of written work. | The written work critiquing the AI summary and their verdict. |
| **Class 2 (optional)** | The Group presentation, about 10 minutes in groups of three. | Parts 1 to 4 at the start of class |

## 04 Learning outcomes

By the end of the week, students will be able to:

1. Turn a marketing or social media claim into a question that can be tested, naming the outcome, the population, and the dose the claim quietly assumes.
2. Rank evidence by quality, and spot when cell or animal research is being presented as if it told you something about people.
3. Build an evidence set that includes support for a claim they expected to be false, and describe what that support actually shows.
4. Check an AI generated summary against the source paper and name what was lost. Summaries keep the headline and drop the dose, the population, and the duration the headline depends on.
5. Give a verdict at the right level of confidence, then say it again in plain language without dropping the qualifications that make it true.

### HOW STUDENTS MAY AND MAY NOT USE AI

Consensus is required here. Its use is graded. The skill being assessed is using an AI research tool well, which is different from using it at all.

**Permitted:** searching, filtering, screening abstracts, and using summary and study extraction features to decide which papers are worth their time.

**Required:** they open and read every paper they cite. Part 3 exists to verify this.

**Not permitted:** submitting AI generated prose as the verdict, citing a paper they have not opened, or handing in an AI synthesis as their analysis.

# Example Submission

Every example in this document follows one student through one claim, so you can see what submissions should look like.

**The claim:** "Magnesium before bed is the single best thing you can do for your sleep. Take 200 mg and you will fall asleep faster and stay asleep."

Posted by a supplement brand that sells magnesium glycinate powder.

PART 0

## The claim capture

Students must find and write a claim down word for word, before any searching. Do not let students paraphrase. The claim should be exactly as stated by the influencer or brand. It will be refined later. Students should attach a link or screenshot.

The best claims are the ones where the student does not already know the answer. Claims that will not work: anything about motives ("the food industry is hiding the truth"), anything with no measurable outcome ("this changed my life"), anything about a specific named person, and anything so broad no single study could speak to it.

EXAMPLE SUBMISSION

| | |
|---|---|
| **The claim, word for word** | "Magnesium before bed is the single best thing you can do for your sleep. Take 200 mg and you will fall asleep faster and stay asleep." Screenshot attached. |
| **Who is making it, and what they sell** | A supplement brand's Instagram account. They sell a magnesium glycinate powder. |
| **The outcome, restated as something measurable** | Two outcomes. Time to fall asleep, measured in minutes (sleep onset latency). And total sleep time. |
| **The population implied** | Everyone. The post says nothing about age or whether you already sleep badly. |
| **The mechanism claimed** | "Calms your nervous system." Vague. Papers describe magnesium as blocking the NMDA receptor and acting on GABA receptors. |
| **Dose, duration, and form** | Dose: 200 mg. Duration: not stated. Form: not stated in the post, but the product is magnesium glycinate. Two of the three are missing. |

PART 1

## The search log

15 points. A table, five rows minimum.

Students are graded on how they searched, not on whether they got lucky. One must be a query built in good faith to find the strongest evidence for the claim. One must use the human studies filter, with the result count logged before and after.

Tell students at assignment time that dead ends earn credit as long as it is explained. “Nothing good came up” is not a reason. “Everything on the first page was rodent work, so I applied the human studies filter” is.

EXAMPLE SUBMISSION

| # | Query, as typed | Filters | Opened / kept | What I changed and why |
|---|---|---|---|---|
| 1 | `magnesium helps you sleep` | none | 6 / 0 | The claim as posted. Top results were mechanism reviews and mouse studies. Nothing here is about a person taking a supplement, so I named a measurable outcome instead. |
| 2 | `magnesium supplementation and sleep onset latency` | none | 5 / 2 | Now I am getting trials. Still mixed in with rodent work. |
| 3 | `Does magnesium supplementation reduce sleep onset latency in adults?` | Human studies | 8 / 3 | Yes or no phrasing gave me a Meter. The human filter cut my results by roughly three quarters, which told me most of the first list was not about people. |
| 4 | Same question as row 3 | Meta-analysis, Systematic review | 4 / 2 | Went to the top of the hierarchy. Found the Mah 2021 review, which pools only three trials and only in older adults. That restriction ended up deciding my verdict. |
| 5 | `benefits of magnesium supplementation for sleep quality` | Human studies, RCT | 7 / 2 | Deliberately looking for the best case in favour. Found the largest trial I could (n = 155) and the 2012 trial in older adults. |

| # | Query, as typed | Filters | Opened / kept | What I changed and why |
|---|---|---|---|---|
| 6 | `Does magnesium supplementation fail to improve sleep quality?` | Human studies, RCT | 4 / 1 | Reversed my own question to find the counterweight. Found a null result in women with PCOS. Kept it, because it changed my verdict. |

PART 2

## The evidence set

25 points. Three peer reviewed sources, four fields each.

Two constraints.

1. At least one source must support the claim, at least in part.
2. At least one must be a primary study with a methods section, not a review.

The supporting source is the constraint students fight hardest and learn the most from. If after a genuine search no supporting human evidence exists, they may substitute a source showing why the claim is plausible in principle, and they must show in the log exactly what they searched.

EXAMPLE SUBMISSION, SOURCE 1 OF 3

| | |
|---|---|
| **Citation and link** | Mah, J. et al. (2021). Oral magnesium supplementation for insomnia in older adults: a Systematic Review and Meta-Analysis. BMC Complementary Medicine and Therapies. [Consensus link pasted here] |
| **Study type and tier** | Systematic review and meta-analysis of randomised controlled trials. Top of the evidence hierarchy, and human research. |
| **Who, how much, how long** | Three RCTs pooled, 151 older adults across three countries, oral magnesium at under 1 g per day compared with placebo. |
| **The finding, and why this source** | Sleep onset latency was 17.36 minutes shorter than placebo (95% CI -27.27 to -7.44, p = 0.0006). Total sleep time improved by 16.06 minutes but was not statistically significant. The authors rate all three trials at moderate to high risk of bias and grade the evidence as low to very low quality. I chose this over a broader supplement meta-analysis I found, because that one did not have enough magnesium data to pool at all. |

**EXAMPLE SUBMISSION, SOURCE 2 OF 3**

| | |
|---|---|
| **Citation and link** | Abbasi, B. et al. (2012). The effect of magnesium supplementation on primary insomnia in elderly: A double-blind placebo-controlled clinical trial. Journal of Research in Medical Sciences, 17(12), 1161 to 1169. [Consensus link pasted here] |
| **Study type and tier** | Double blind randomised placebo controlled trial. Primary human study, so this is the one where I read a methods section. It is also one of the three trials inside the Mah review. |
| **Who, how much, how long** | n = 46 older adults with primary insomnia, 500 mg magnesium daily or placebo, for 8 weeks. |
| **The finding, and why this source** | Insomnia Severity Index fell in the magnesium group compared with placebo (p = 0.006), sleep onset latency fell (p = 0.02), and sleep efficiency rose (p = 0.03). I chose it over another trial in older adults because the reporting in that one was too thin for me to check the dose. |

**EXAMPLE SUBMISSION, SOURCE 3 OF 3**

| | |
|---|---|
| **Citation and link** | Gholizadeh-Moghaddam, M. et al. (2022). Effect of magnesium supplementation in improving hyperandrogenism, hirsutism, and sleep quality in women with polycystic ovary syndrome: A randomized, placebo-controlled clinical trial. Health Science Reports. [Consensus link pasted here] |
| **Study type and tier** | Randomised placebo controlled trial. Primary human study. |
| **Who, how much, how long** | n = 64 women with polycystic ovary syndrome, 250 mg magnesium oxide daily or placebo, for 10 weeks. |
| **The finding, and why this source** | No significant effect on sleep quality compared with placebo (p = 0.85). This is the source that stopped me writing “supported.” Different population, different form of magnesium, lower dose, and a flat null result. I went looking for it on purpose in row 6 of my log. |

**PART 3**

## The discrepancy report

25 points. About 250 words.

Students pick **one source** from their evidence set and compare what the tool told them about it, meaning the result summary, the extracted study details, and any AI synthesis, against the full text. They read the methods section and the limitations the authors name themselves.

**Three things must appear.** One thing the AI layer got right that saved real time. One thing it flattened, omitted, or framed in a way that would have misled them. And one concrete practice they will adopt, not a platitude about being careful.

"I found no discrepancy" is acceptable only with a field by field comparison showing how they got there. It is rarely the right answer. An unsupported claim of no discrepancy scores as if the comparison were never done, because functionally it was not.

> **EXAMPLE SUBMISSION, ABOUT 250 WORDS**
>
> **Source checked: Abbasi et al. (2012)**
>
> **What it got right**. The extracted study details gave me the design, the sample size, the dose, and the length of the trial in about ten seconds: double blind, placebo controlled, 46 older adults, 500 mg daily, 8 weeks. Every one of those matched the paper. Finding that by hand would have taken me twenty minutes of scrolling.
>
> **What it flattened.** The summary I read led with sleep improving. Inside the paper, the picture splits. One sleep duration measure is reported as improved, and total sleep time is reported as showing no significant difference between groups at p = 0.37. I could not fully work out from the text which measure was which, and that is the point. A summary that keeps only the improvement gives you a cleaner story than the paper does. I would have walked away thinking magnesium reliably makes you sleep longer.
>
> The second thing it dropped was who this was. Everyone in the trial was an older adult with diagnosed primary insomnia, on 500 mg for 8 weeks. The post I am checking says 200 mg for anyone. That is less than half the dose, in a population nobody in this study resembles.
>
> **What I will do differently**. Before I cite any paper, I will write down its n, population, dose, and duration in one line, and then check that line against what the claim assumes. If they do not match, the paper is context, not proof.
>
> **Link to Consensus Search or Screenshot of AI Summary:** (example)

PART 4

## The verdict

20 points. About 300 words of analysis.

Students choose exactly one of four verdicts.

| Verdict | What it requires |
|---|---|
| **Supported** | Multiple human studies point the same way, at doses and in populations that resemble what the claim implies. |
| **Supported under conditions** | A real effect exists, but only in a narrower population, dose, duration, or outcome than the claim states. Students must name the conditions. This is right far more often than they expect. |

| Verdict | What it requires |
|---|---|
| **Unsupported** | Human studies exist, they looked for this effect, and they did not find it. This needs evidence of absence, not just absence of evidence. |
| **Untested** | No human studies exist at the relevant dose or in the relevant population. Different from unsupported, and that difference is why there are four categories. |

The 300 words must name the verdict, give the evidence that drove it, and say what would move it to a different category.

> **EXAMPLE SUBMISSION, PIECE ONE**
>
> **Verdict: Supported under conditions.**
>
> There is a real effect here, but it is smaller and much narrower than the post says. The strongest human evidence I found is a 2021 meta-analysis pooling three randomised trials in 151 older adults. People taking magnesium fell asleep about 17 minutes faster than people taking placebo, and that result was statistically clear. Total sleep time improved by about 16 minutes and was not statistically significant, so the "stay asleep" half of the claim is weaker than the "fall asleep faster" half.
>
> Three conditions have to be named. First, population. Every trial in that pooled analysis was in older adults, and the trial I read most closely was in older adults with diagnosed primary insomnia. The post is addressed to everyone. Second, dose. Those trials used 250 to 500 mg. The post says 200 mg, which is below the range that produced the effect. Third, evidence quality. The review's own authors rate all three trials at moderate to high risk of bias and grade the evidence as low to very low. They write that the literature is not good enough for physicians to make firm recommendations. That is the review arguing against its own headline number, and it belongs in my verdict.
>
> I also found a trial in 64 women with polycystic ovary syndrome, at 250 mg for 10 weeks, that found nothing at all on sleep quality ($p = 0.85$). One null result does not overturn the pooled analysis, but it does show the effect is not universal.
>
> What would move me. A trial in healthy adults with normal sleep, at 200 mg, showing the same 17 minute effect would move me to supported. A second and third null trial in adults who already sleep normally would move me to unsupported for that group.

PART 5

## The Case Presentation (optional)

Each student already has a completed claim autopsy. In the first five minutes, the group compares all three claims, verdicts, and single strongest sources, and picks one to present. The group then presents that one claim together, screen shared, with the work split three ways so all three students have a real job even on a claim they didn't originally research.

| Section | Time | The job |
|---|---|---|
| **Search log** | 5 min | Pull up the actual search log on screen. Show the queries as typed, the filters applied, and why the search counts as thorough, not just lucky. |
| **Evidence & verdict** | 5 min | Walk through the three sources, their tiers, and what each one found. Say which source is doing the most work for the verdict. |
| **Discrepancy and verdict** | 5 min | State the verdict and the one thing that would move it to a different category. |

Whoever didn't originally research this claim needs to understand their section well enough to present it, not just read from their groupmate's doc.