When I was about fifteen years old we had retired to our house near Belrive, w
terrible thunderstorm. It advanced from behind the mountains of Jura; and the thunder burst at once wi
loudness from various quarters of the heavens. I remained, while the storm lasted, watching its progress with curiosi
and delight. As I stood at the door, on a sudden I beheld a stream of fire issue from an old and beautiful oak, which s
about twenty yards from our house; and so soon as the dazzling light vanished, the oak had disappeared, and nothin
found the tree shattered in a singular mann

THE LOOP



HANDSHAKE CRISIS

Gen Z crumbles under pressure



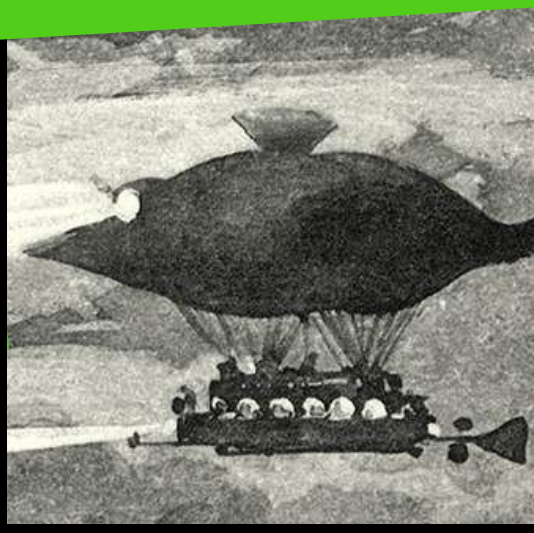

GUIDE TO WEEK-LONG EXPERIMENTS

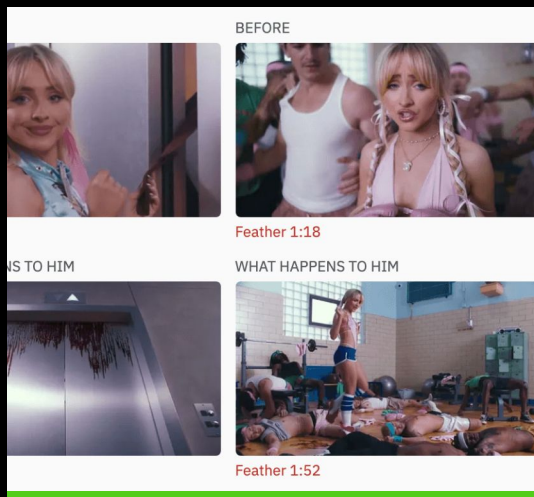

SABRINA: AVERT YOUR EYES

EXCLUSIVE! Superblessed vs. Supercursed Water

THE LOOP IS BROUGHT TO YOU BY... BRAVERY

SAY WHAT YOU REALLY MEAN



THE POWER OF COURAGE

EDITOR'S NOTE

***The LOOP** is a record of the undertakings, studies, and labours of curiosity from around the world. Our ingenious contributors offer accounts of the present Philosophical Matters from around the world. We salute our friends for their thoughtful reporting.*

–SLIME MOLD TIME MOLD ❀



Page 3: 📍 You are here 📍

Page 4: How 2️⃣ turn your life into an experiment 🔬

Page 11: Blessed 🙏 Water 💦

Page 21: 🤯 mad science mind games 🤯

Page 2️⃣ 2️⃣ : High Uncertainty Handshake Protocols

Page 30: DIY! Potion making 🧪

Page 38: *sluuuuuuuuuuuurp*

Page 44: Is it healthy to be rational ❓

Page 46: Dark Soda Detection

Page 55: Sabrina Through the Looking Glass 🪞

Page 63: Contributors 🖋️🎨

SUBSCRIBE:



WHAT 5 WEEK-LONG EXPERIMENTS TAUGHT ME ABOUT CHANGING HABITS

by BEN GROSSMAN



For one week, I wasn't allowed to be sarcastic.

I told my friends and family about it, and they called me out every single time I slipped up. Which was a lot.

So I failed. But I became far more aware of how and when I use sarcasm— every time I caught myself (or failed to), I gained some insight into my current behaviors and how I'd like them to change. I'm now much more aware of when my sarcasm is mean-spirited, which is helping me cut that out almost entirely without losing the 'good sarcasm': the kind that's just for laughs and isn't hurtful.

That was week 2 of something my friends and I have been doing off and on for the past couple months. We call them experiments.

Where the idea came from

This idea came about when I was talking to my friend Ryan (who recently spent 2 days breaking 25 world records). I was lamenting how often my behaviors don't align with my values, and yet it's so difficult for me to change my habits. Ryan recalled a post from Neel Nanda, a renowned AI researcher, about the value of running experiments on your own life.

I loved the idea so much that I decided we were going to start doing it immediately. He was down, but wanted to make sure we limited failure modes as much as possible by minimizing the number and complexity of the rules.

So there's only one rule: commit to a full week of the behavior. That's it.

How it works

Experiments are typically either daily activities (meditate every morning, write for 15 minutes a day) or explicit rules (no phone allowed in bedroom, no sarcasm). An experiment can start any day of the week, and you can do them as often or as rarely as you want, as long as they last 1 week.

My experiments so far fall into 2 categories: habits I want to get into that need a kick-start, and challenges I don't plan to commit to long-term but that I can learn from. No sarcasm and daily writing were challenges. The rest were kick-starts.

What I tried

Exp 1: No phone in my bedroom. I slipped up once or twice, walking into my room with my phone before realizing, but I didn't use my phone in my room at all. I slept much better, my alarm clock did its job, and going to bed and getting up were both much more efficient. Surprisingly, my phone usage stayed about the same, ~3.5 hrs/day. So the phone hadn't been costing me time in my room, it had been costing me sleep. 2 months later, and I still follow this rule. Game changer.

Exp 2: No sarcasm. Described above.

Exp 3: Daily posture routine. My goal was to create a short routine I could fit into my day. I did it 6/7 days, and have kept doing it a few times a week since. 2 months later, it's hard to tell if my posture has gotten any better, but my back and neck seem to hurt less than they used to. I consider this one a small win.

Exp 4: Morning walk. This was an awesome addition to my week. ~10-15 min every morning is a small price for the ROI I got: more energy when starting work or whatever I had to do that morning, and a reason not to lay in bed after waking up. It likely helped my sleep too, since morning sunlight is good for your circadian rhythm. On top of that, I got to look at trees, pet dogs, call friends, and run into several neighbors. I've kept doing it most days, though I'll have to make some tough decisions once it gets freezing and snowy here in Boston.

Exp 5: Write every day. This one flopped. Read on to find out why.

Why it works

Three things: play, people, and friction.

Play. When I told my friend Seb about the experiments, he loved the idea. He said they bring a sense of play to daily life, and he's right. Every time you do an experiment, you're reminding yourself that you can mess around with your behaviors.

People. I tell whoever I'm spending time with what my current experiment is. During the sarcasm week, that turned my friends and family into referees. I didn't have to catch every slip myself, because someone else always did.

Then there's the group chat. It started as a WhatsApp message to Ryan, and it has since grown to include some of our other friends. Every week I post how my last experiment went, then what my next one will be.

Posting it makes it real. I do care about the commitments I make to myself, but sometimes it's easier to lean on my reputation— nobody wants to look like a quitter. In reality my friends probably couldn't care less, but my amygdala doesn't know that.

Having a community, even when it was just Ryan, has other benefits too. I can run ideas by people, steal their ideas entirely, and (my favorite) feel happy when others copy my experiments.

Friction. A large part of self-improvement is intentionally adding and removing friction. Keeping my phone out of my bedroom adds friction to scrolling in bed. Friends calling me out adds friction to being sarcastic. Telling the group chat about an experiment adds friction to quitting.

The one that flopped

Experiment 5 was to write about the experiments. As in, this post.

For the first 2 days I stuck to it and spent ~15 min each day assembling a doc and writing stuff down. On days 3 and 4 I was less motivated, but still made a bit of progress. By day 5 I barely even opened the doc. The main issue was that I didn't want to write— writing is something I find challenging, and I was having trouble putting coherent thoughts down beyond what I had already written.

After that, I stopped experimenting for a couple of weeks. The novelty had worn off, and I'd used up some of the ideas I was most excited about. I also felt guilty. Every week, I'd posted my results in the group chat, but for experiment 5 I never did. Apparently my amygdala has its limits.

But you're reading this, so it got written eventually. The nice thing about a one-week experiment is that failing only costs you a week. Now that I've finished writing this, I'm ready to start experiment 6: start a conversation with a stranger every day. Wish me luck!



So, what makes a good experiment?

I'm still figuring this out, but here's my working theory. A good experiment has a few qualities:

1. **You can do it daily, or multiple times a day.** You're only doing it for a week, so it's important to get several reps in. That's how you get over the hump of 'this feels uncomfortable' to a point where you can rationally evaluate the behavior.
2. **It feels worthwhile.** Maybe you're just curious about the result, or maybe it's an excuse to try something you've been meaning to do for a while. Either way, trust how you feel about it. I didn't feel great about my writing experiment from the start. Lesson learned.
3. **It's a little bit uncomfortable, but not so much so that you won't do it.** Committing 10-15 minutes a day for a week is a small cost. Don't try to change your life overnight. Push your boundaries, but only slightly. (Some people are willing to go full-send— if that's you, godspeed, but it's certainly not me.)
4. **Failing is cheap, and still teaches you something.** I failed a lot during no-sarcasm week, but still learned from it. Missing a day of posture was totally fine, and didn't impact the other days. In contrast, failing during writing week felt much more significant, since it meant I failed to complete this post. An experiment should focus on a behavior, not a deliverable.

Not every change will fit the week-long structure of these experiments. That's OK! Lots of my ideas haven't fit this format: clean my room, go to the gym once a week, be more whimsical. I'll work on these, but probably not as experiments.

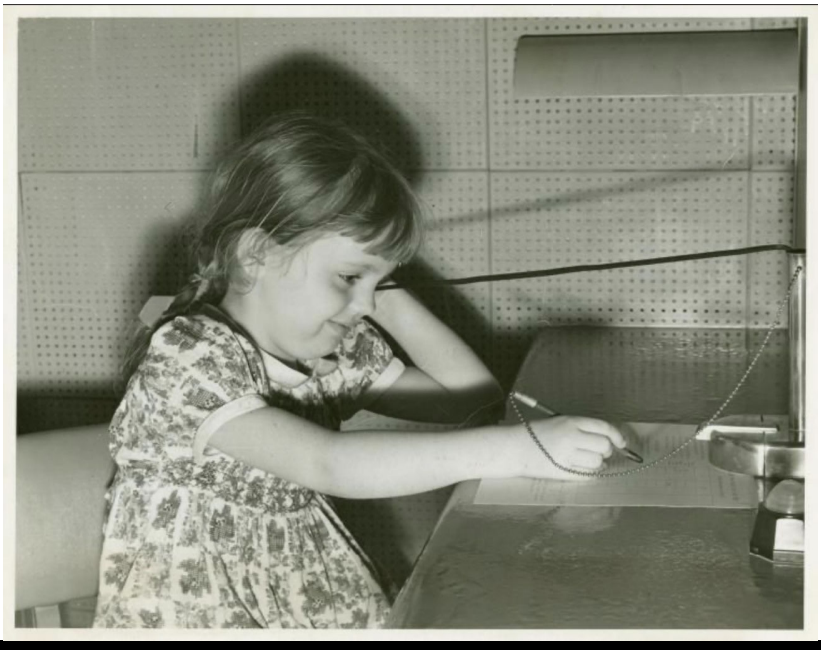

Try one

Here's my invitation: pick one thing, do it for one week, and tell one person.

Before the week starts, set yourself up. Ask yourself 2 things: what might make it hard to get through the week? And what can I do now to make it easier? For me that meant buying an alarm clock so my phone could stay out of my bedroom, telling all my friends I wasn't allowed to be sarcastic, and writing down my posture routine before day 1.

So far, my ideas have come from 3 places:

1. My bad habits (eg: scrolling phone before bed)
2. Others' good habits (eg: my friend Alan goes on a morning walk every day)
3. Random, intrusive thoughts (eg: "I wonder what it would be like to stop being sarcastic altogether")

If you're looking for inspiration, here are a few ideas from my list:

- No headphones in public
- Compliment someone every day
- Dance (alone in your room) for 5 min every morning
- No snacking all week
- Go somewhere you've never been before every day
- Every day, call someone new you haven't spoken to in a while (even just for a 2-min convo)

If you try one, let me know how it goes. I'll probably steal it.



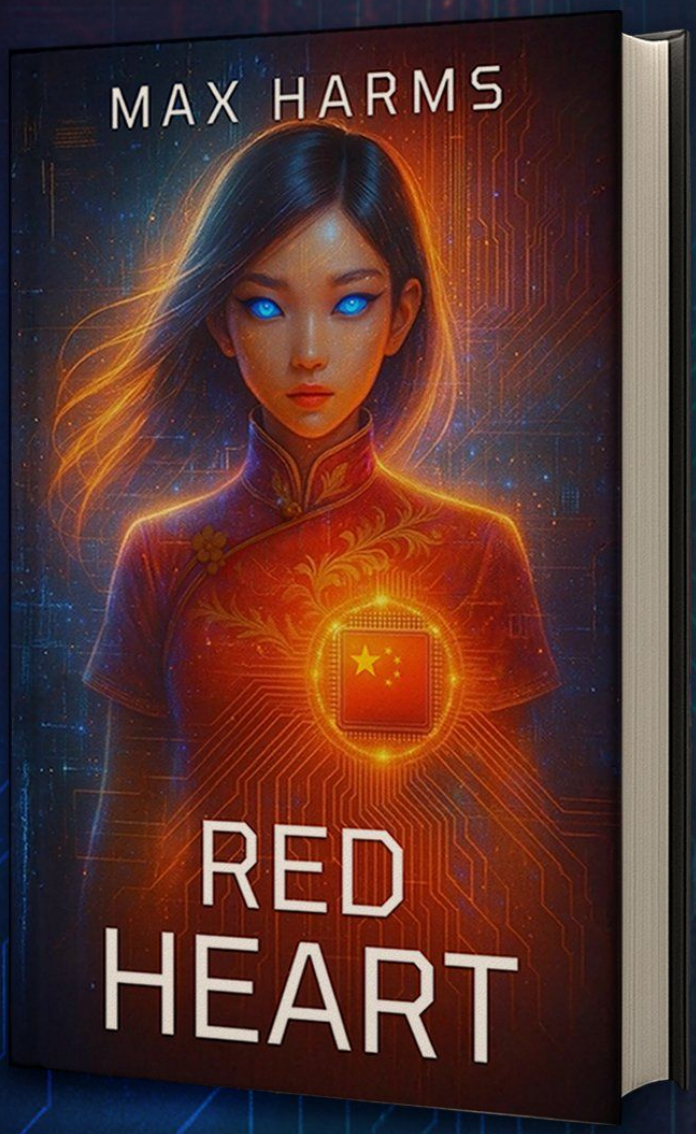

RED HEART

MAX HARMS

"A strikingly authentic portrayal of AI in modern China."

- Zhang San

From the author of *Crystal Society* comes **a new novel** about Chinese AGI.

"[*Red Heart*'s] key strength is its grounded and realistic treatment of AI alignment, corrigibility, and geopolitics. AGI will transform the world in the near future and *Red Heart* gives us a glimpse of how."
— Xerxes Dotiwalla, AI Alignment Researcher at Google DeepMind

"A thrilling spy novel grounded in realistic AI safety research and geopolitical possibilities."
— Nathan Helm-Burger, MATS Research Manager

"I like it!"
— Daniel Kokotajlo, AI Futures Project

[*Crystal Society*] seems to belong in the very, very tiny subset of AI stories that are not bloody stupid, a heroic and almost unbelievable accomplishment.

— Eliezer Yudkowsky



LEARN MORE AND READ NOW

WATER OF LIFE KIT REVIEW

By Valentino di SMTM & Paul Kreiner

On the internet there is a PDF — you can read it here — 239 pages long and titled "LITERALLY THE BEST THING YOU'LL EVER READ". This PDF describes a procedure that uses normal household items to create what is called "charged water", "mana", "negative water", "strong water", or the term we will use for the rest of the piece — blessed water.

Supposedly, this procedure can fill a glass of water with negative ions. In the words of the author: "I found a way to add any amount of electrons I want." This blessed water allegedly has many wonderful, near-miraculous properties, and the author connects it to everything from the legendary fountain of youth to the exploits of arctic explorers. (For an overview, also see the wizard.talk website; the intro sections there are updated and the best diy guide for this right now.)

This PDF might be described as "extremely schizo". And yet, it makes many highly specific claims that are quite readily tested. It claims that the water looks and feels different in specific ways, that it sounds different from normal water when poured, that it has noticeable subjective and health effects. Even the claim that plants exposed to blessed water will grow faster and "will look prehistoric" (Page 84) is as testable as it is insane. If plants do grow faster with regular water, we can just see! We can grow some cress and find out in like 7 days!

At the risk of being crass, science is subject to expected value calculations just like everything else. The value of a study or experiment is the value of the idea if true, times the probability that the idea might be true at all — then multiplied again by the difficulty of running the study or experiment. So when an idea is "big if true" and the effort for running the study is low, then the expected value can be extremely high — so you gotta run that study.

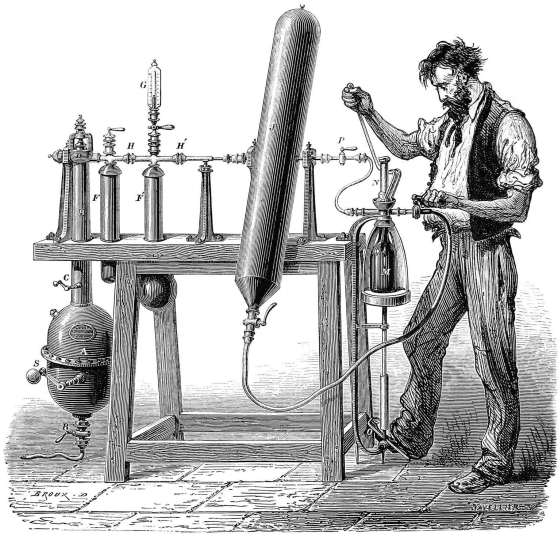

We decided to test three main claims:

Claim 1: Color. The PDF claims that blessed water is not only beautiful, but slightly blue:

> I want you to **look closely at the water**. Does it strike you as beautiful? Try to shine light on the water. It reflects light differently, and absorbs red, causing a slight color change toward green & blue (hence why the ocean and sky is blue. You have a bottle of distilled sky, if you like to describe things in terms of elements). [Page 50]

Claim 2: Sensation of cold. The PDF claims that blessed water feels perceptably cold:

> First sensation - it felt COLD. Odd, since it was room-temperatured. In a pleasant way. Fresh air feels cold too, even in the summer. This felt like that. Coldness that you welcome, like a refreshing drink. This wave of pleasant coldness spread through my arm, into my body. [Page 10]

The interesting thing about this claim is that the PDF explicitly says that the sensation of cold is *not* the result of an actual temperature difference, but of the difference in charge:

> Pour a few drops of the water into the palm of your hand.
>
> It feels cold, doesn't it? It's not the temperature you are feeling, it is the negative electrical charge. [Page 48]

Claim 3: Feels good, man. The PDF claims that touching blessed water feels exceptionally good:

> My entire being was focused like a laser beam into the seemingly simple and casual act of having my index finger in that tiny little glass of water. It felt so good, and I never wanted it to end. [Page 11]

We tested Claim 1 with photography and Claims 2 & 3 with a blind touch test involving 20 participants.

These studies shouldn't be seen as definitive tests of any of these claims, because the methods we use are not very sensitive. They would only be able to detect very large effects. (To be fair, the PDF claims that these effects are in fact very large, but it's also possible that the author is unusually sensitive, and the effects are real but smaller than described.) Instead, see these studies as attempts to benchmark whether or not we can detect something along the lines of what is claimed.

Materials

Two ziploc bags containing water (distilled water for blessed side, tap water for cursed side) were placed side by side in a large bowl lined with aluminum foil, coupling them capacitively. Two stainless steel spoons served as electrodes. Then the two sides were connected to the voltage source: a 72V bank of batteries. This imparts some potency to the water; we 'cycle' the process to add more potency, seemingly linearly: disconnecting the circuit, lifting the minus bag away a few inches, putting it back, and then reconnecting the circuit. Here, we cycled 10 times. At the end, the water on the minus side is called 'blessed', and the water on the plus side is called 'cursed'.

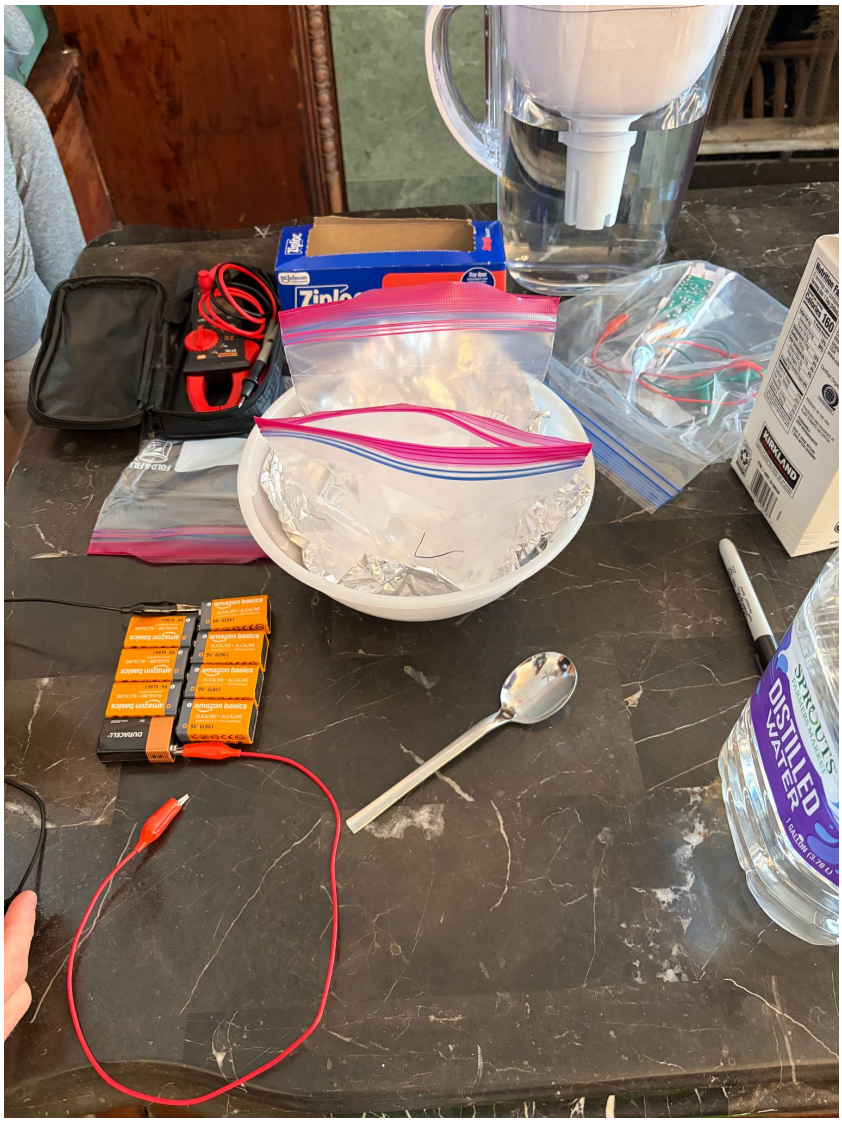

This is claimed to be good enough for normal use, but we wanted the best possible chance to detect any effect, and figured that a stronger potency would make it more likely. So instead of using the normal blessed water in our tests, a disassembled bug zapper that generates a transient 3-4 kV voltage was also used to create “superblessed” water and conversely “supercursed” water. These were used in all experiments.

Claim 1

Photographs

We considered asking people to indicate which water was more blue, but decided that photos would be more consistent.

We took photos with an iPhone in two locations: in the basement under warm light against a white wall, and outside in natural light against a piece of white printer paper. We analyzed two photos from each location.

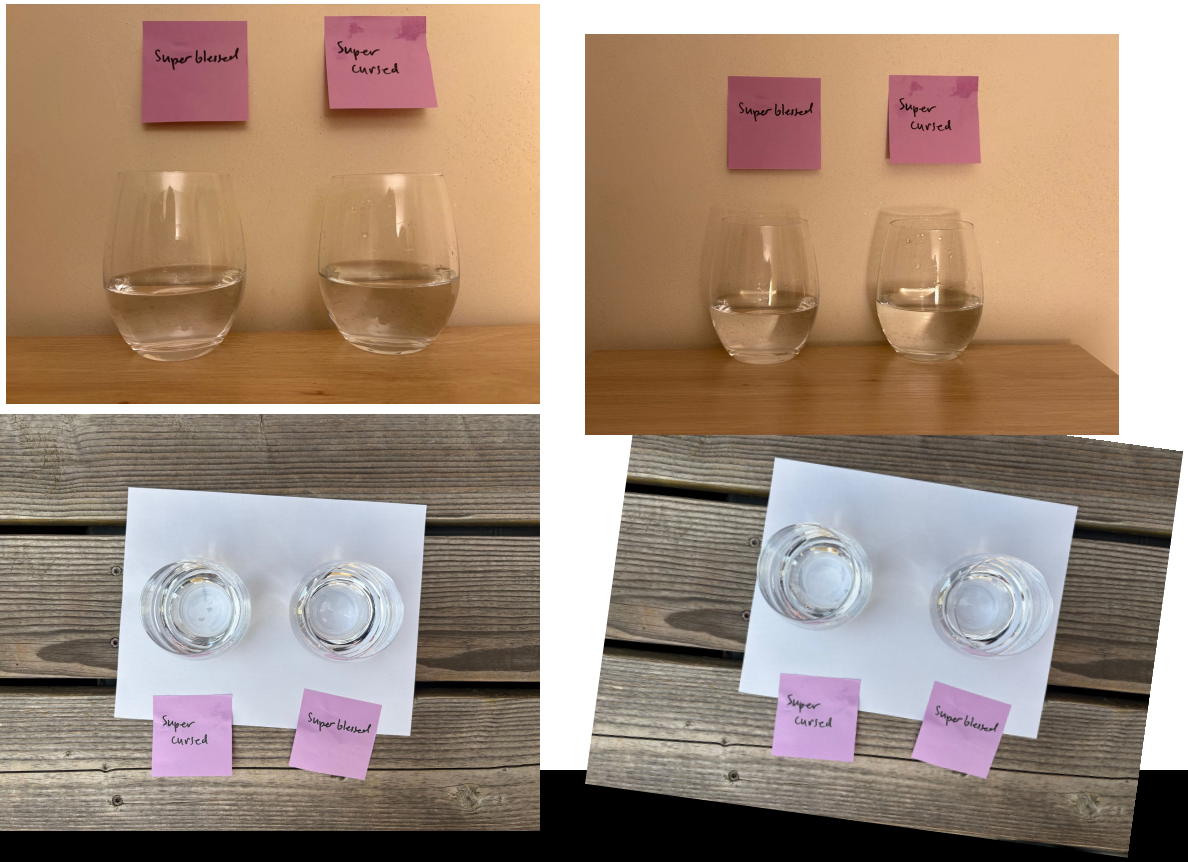

We used a python script to measure the average color inside the cups and convert it to CIELAB, a standard color scale designed to match human perception. The b* value runs from blue (negative) to yellow (positive), so we can test if one glass is bluer just by looking for lower b*.

Results

We observed no consistent difference with analysis of b*. In the two basement photos, in one the superblessed was slightly bluer and in the other the supercursed was slightly bluer. In the two natural light photos, in one the superblessed was slightly bluer and in the other the supercursed was slightly bluer. All differences were invisible to the eye and within ordinary photo noise.

Claims 2 & 3

Methods

We presented water in three identical glasses: A glass of the original distilled water on the left marked with a purple post-it with a circle, a glass of superblessed water in the middle marked with a purple post-it with a triangle, and a glass of supercursed water on the right marked with a purple post-it with a square.

The order of the cups was consistent across all participants, but because some participants started by touching the water in the left cup and some started by touching the water in the right, and many went back for multiple touches, effective order was somewhat random.



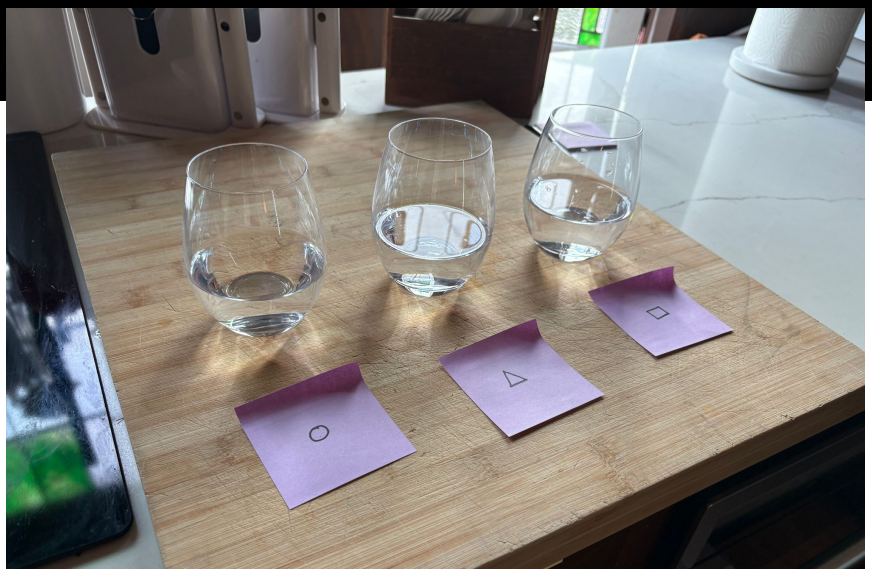

Because of some concerns that the effect would diminish over time and with repeated touch contact, we refreshed the cups with a new batch of distilled/superblessed/supercursed water every 10 participants, so for this study we had two batches.

We recruited 20 people around Lighthaven to participate in this study. Most were here for Inkhaven 3 or the Illiad Fellowship, but others were hanging around for other, possibly nefarious reasons.

We ran everyone through the same script, which went approximately like this:

> Here are three cups of water. We want you to touch each of the waters and pay attention to how they feel. We also want to make sure to avoid cross-contamination, so please don't drip water from one glass into any of the others, and please wipe off your fingers between cups with a paper towel. [Provided immediately to the right of the cups.] When you are done touching all three waters, we will ask you two questions. If you feel that you need more information, you can touch the waters again.

When participants had touched all three waters and were done touching, we asked them the following two questions:

"Which cup feels coldest?" The PDF suggests that blessed water feels colder, so people will select the superblessed cup.

"Please rank the cups from 'feels best' to 'feels worst'." The PDF suggests that people will have a strong attraction to the superblessed water and aversion to the supercursed water.

Results

Cold

One person refused to answer this question. Of the 19 people who were willing to answer, 9 reported that the superblessed cup felt the coldest. Chance alone would yield about 6, and 11 or more would be needed for this to be statistically significant, so in a strict sense we were unable to confirm this hypothesis (binomial test, $p = 0.146$).

This result isn't statistically significant, but we can still interpret this as evidence. It's consistent with a small or medium-sized effect and certainly trends as expected. The useful question is: given 9 of 19, how likely is it that the true effect points the way the book predicts? We can't reject the idea that the effect is exactly zero, but with no prior, there's about a 88% chance that superblessed really does feel coldest more often than chance, so only about a 12% chance the trend is pointing the wrong way. Starting from moderate skepticism (a prior centered on chance): instead, there's about a 73% chance that superblessed really does feel somewhat cold.

The result leans the predicted way and fits a small-to-moderate effect, but this finding is suggestive, not conclusive.

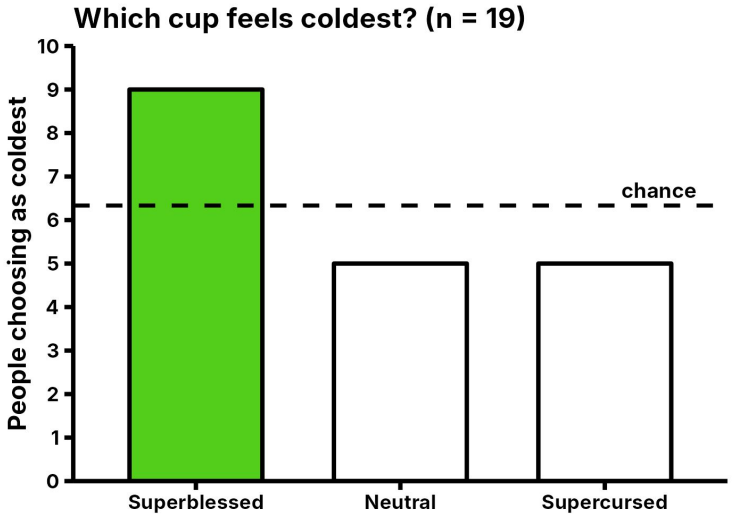

Hooked on a Feeling

Two people refused to give a ranking, and one person ranked “square best, others same”, showing a preference but making it impossible to give a complete ranking.

In the 17 complete rankings, participants did show a markedly strong preference..... for the supercursed water.

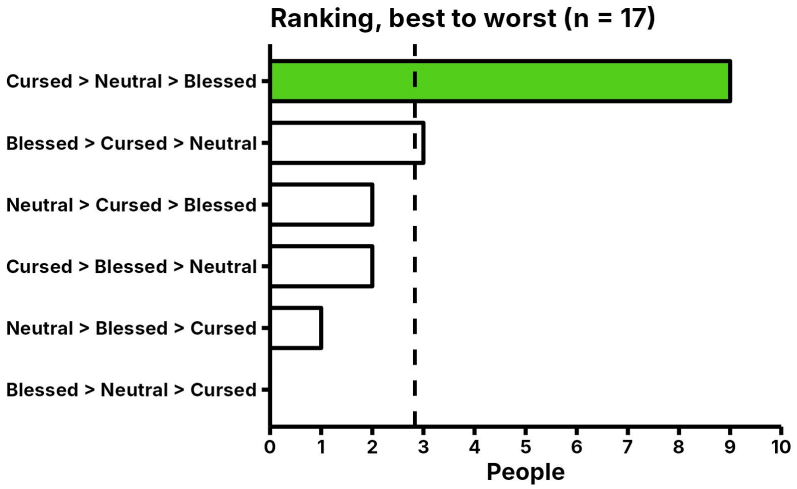

That’s right. In fact, nobody ranked the cups in the order predicted by the PDF (blessed best, then neutral, then cursed). But a lot of them ranked them in the exact opposite.

The single most common ranking was the exact reverse of the prediction. We saw 9 of the 17 people (52.9%) who gave a complete ranking, ranked supercursed the best, distilled as neutral, and superblessed as the worst.

This tendency extended even to people who didn’t pick this exact order — 12 of 18 ranked the supercursed cup best ($p < 0.001$), and 11 of 17 ranked the superblessed cup worst. This is reasonably inconsistent with chance, $p \approx .001$ any way you slice it.

Comparing the two questions, we found that of the 9 who called superblessed coldest, 7 also ranked it worst. So “cold” seems to have read as unpleasant, not “cold in a good way.”

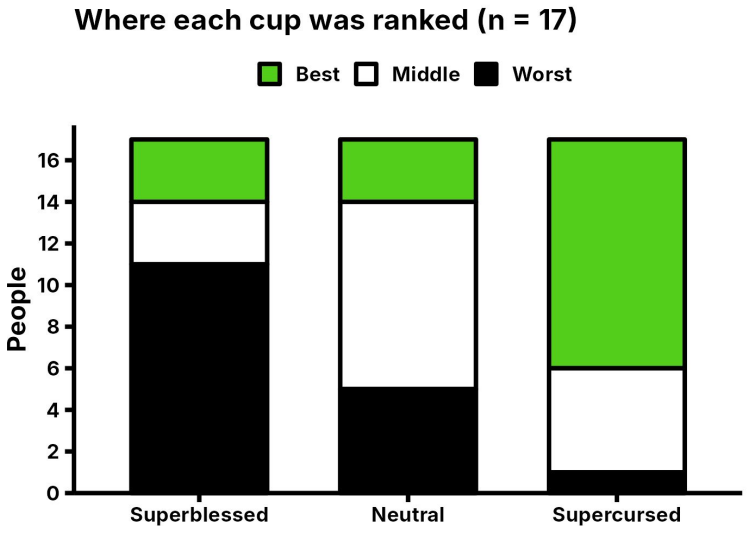

Discussion

When you test a crazy idea, you go in prepared for it to fail — especially when you run your studies in an afternoon and use weak designs with small sample sizes. Even if the effect is 100% true, you might fail to detect it.

The claimed visual effect of blessed water being notably different in color seemed to be untrue, as confirmed by our photo analysis, making the other claims seem unlikely to be true as well.

So it felt pretty surprising to see the touch test results roll in. By participant 5 or 6 it started to get weird, people kept saying that supercursed water felt the best.

It's strange enough to appear to confirm the claims around blessed water, but especially strange to confirm them in the opposite of their claimed direction. Why would supercursed water be ranked as feeling the best, and superblessed the worst? A few possibilities come to mind.

First, it's possible we got everything backwards. Maybe one of us swapped the waters at some point, and what we thought was superblessed was actually supercursed, and vice versa. The fact that we used a fresh batch of charged water halfway through makes this unlikely, however, as we would have to make the same error twice. But maybe the charging setup was reversed by accident or something.

Second, maybe blessed water has an unpleasant short-term feeling despite (possible) positive long-term effects. For example, warmer-feeling water might feel pleasant and colder-feeling water might feel jarring. This could explain why people like how the supercursed water feels and dislike how the supercursed water feels, even if they have the predicted effects in the long-term — which we basically didn't test, the study involved only dunking their fingers and lasted a few minutes at most.

Finally, it's possible that the claims in the PDF are broadly wrong, but there's some other change that makes the "supercursed" water feel better and the "superblessed" water feel worse for some undetermined reason.

While we can't say we have confirmed the claims around blessed water, and the effects are not as knockdown powerful as the PDF suggests, **this is fairly strong evidence that people can in fact subjectively tell superblessed from supercursed water by touch**, that the effect is linear (given the position of the distilled water), and that these claims are easily investigated and worthy of further investigation.



mad science mind games

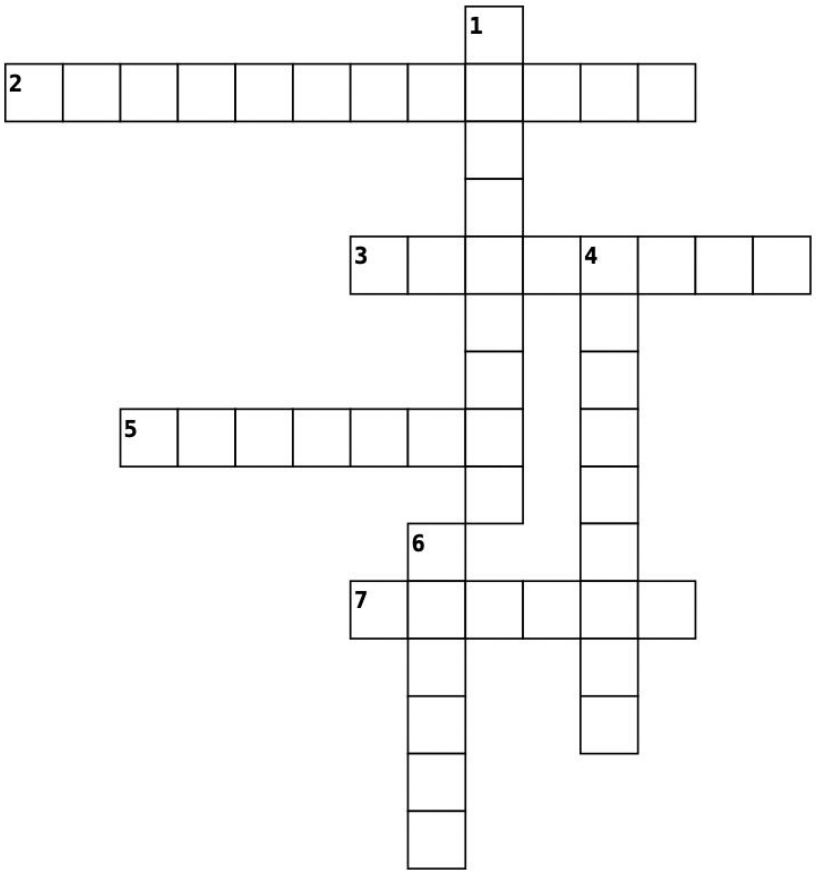

Across

2. alchemical explosive

3. worm pasta

5. what's infected by h. pylori

7. the lord of rain and dew

Down

1. evil laugh

4. originating in observation or experience

6. in a scientific pursuit there is continual food for discovery and _____

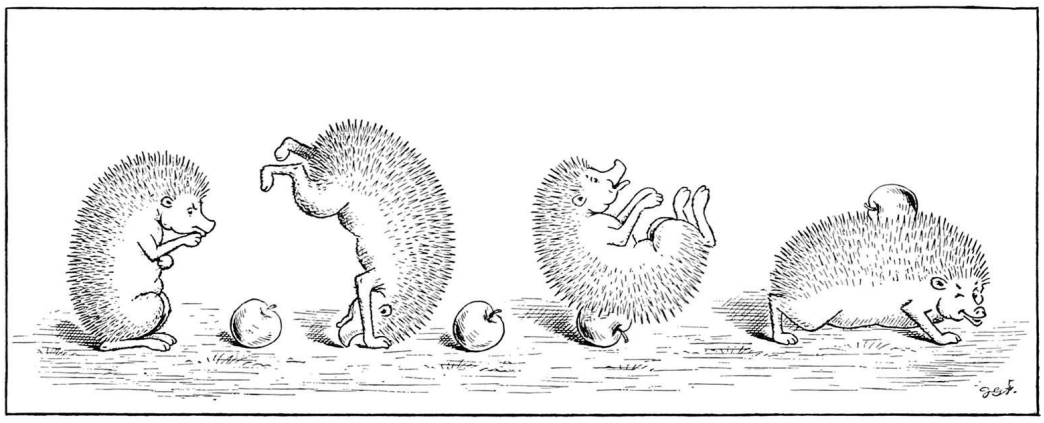

ROBUST HANDSHAKE DYNAMICS FOR HIGH-UNCERTAINTY ENVIRONMENTS

by Karthik Bala & wadapan

Abstract

Handshakes are a form of bodily contact where appropriate aim, kinodynamics, and muscular force are synchronized between participants for the purpose of social bonding (Tilston et al., 2025).

Smooth handshake coordination results in mutual halo effects between participants, while miscalculated handshakes precipitate social disharmony and, in extreme cases, catastrophic mistrust (Chaplin, 2000).



Considering the immense potential for well-coordinated handshakes to build trust between participants, we believe handshake R&D is a sorely neglected cause area that demands more funding and attention. Proper handshaking has the potential to end America's crisis of trust and increase domestic GDP by 2-3x (Schroeder et al., 2019).

A Taxonomy of Handshake Circumstances

Handshakes serve several subtly different purposes in social interactions, depending on the level of existing familiarity between participants. We identify the following circumstances wherein a handshake is typically appropriate:

- **Introductions** between people who have not already met
- **Greetings** between people with a pre-existing affiliation
    - **Cursory greetings** which serve as a formality or acknowledgement
        - **Acknowledgements** used to confirm participants' inclusion in a time-bounded social context
        - **Congratulations** where one participant has experienced a personal success
        - **Condolences** where participants meet in a context of illness or bereavement
    - **Affectionate greetings** which serve as a shared activity and expression of friendship
    - **Brazilian jiu-jitsu match commencements** where a "slap-bump" is used to coordinate the beginning of a fight
- **Agreements** between people in a working relationship
    - **Contracts** between partners, i.e., the "handshake deal"
    - **Compromises** used to signal conciliatory intentions following a disagreement

Although agreements typically mark the end of an interaction, the handshake nevertheless fulfills a similar social function: the participants "reintroduce" themselves in the context of the preceding interaction, establishing a new relationship. In this report, we have chosen to focus on greetings between casual acquaintances. Handshakes are also performed under gendered expectations; while male-male handshakes are most common, handshakes involving women or nonbinary people also exist, and must be incorporated into a unifying theory of handshake dynamics.

Technique Classification

Although handshakes in the real world manifest in near-infinite variations, the vast majority of these can be classified as one of three basic types.

The Salesman Shake
In the most common form of the handshake, right hands are extended with palm perpendicular to the floor, then coupled and raised and lowered several times over a period of less than 3 seconds, with strokes decreasing in amplitude monotonically (Nagy et al.).

The Pound Hug
Also referred to as a pound shake, hip-hop hug, one-armed hug, homie hug, shug, bro-grab, bro hug, brah hug, thug hug, man-hug. Unlike the open-armed hug which symbolizes a total vulnerability and openness, the pound hug "interposes the obstacle of the two right arms to the joining of the two bodies." (Wood 2012).

The Fist Bump
The fist bump is a courtier domestication of the punch. Each fist bump is a celebration of the world-historic civilizing process that has transformed men from barbarians to brothers. The subversiveness and symbolic weight of the fist bump makes it a popular way to reinforce masculine bonding.

The primary challenge presented by handshake coordination is thus recognizing in the moment which type of handshake is emerging. Unless participants are embedded in a community with pre-determined gestural norms (e.g., Japanese tea rituals, Black Panther meetings), they must negotiate the handshake under heavy uncertainty.

Will it be a dap into a pound hug? A straightforward up and down? What happens when person A thinks the former and person B thinks the latter? Live handshake negotiation is an intricate dance of reading subtle psycho-motor and cultural cues.

Study Rationale

Bad handshakes can lead to mistrust, loneliness, severe anxiety, depression, persecution, and in some cases, death. If we look only at the rumination caused by bad handshakes we are confronted with a conservative disease burden of two to three billion disability-adjusted life years (DALYs).

Good handshakes can lead to increased life satisfaction, geo-political opportunities, and national prosperity. The ability to give a good handshake is linked to a 40% risk reduction in all-cause mortality and a 36% increase in lifetime earning potential (Stewart et al., 2008).

Case Studies

We will now illustrate the complexity of handshake dynamics through a series of case studies. Subjects were approached in communal areas (kitchen or courtyard) and selected for pre-existing rapport.

QC/Karthik	Despite a confident approach, the coupling was non-standard (see recreation, Figure 1). This resulted in complications within one or two seconds, as neither party knew how to disengage. We speculate that the high level of contact in the coupling made it more conspicuous to release the grip, resulting in self-consciousness and paralysis.
Zoe/Karthik	Both parties engaged confidently with a classic, perpendicular angle of approach. However, Zoe experienced a full-body adverse reaction to the contact, resulting in a spasm travelling up her arm from her hand. After physically recoiling, Zoe hopped around from foot to foot. Catastrophic failure. It is unclear whether Zoe or Karthik was at fault; Zoe had been party to the preceding handshake with QC, and thus may have been primed for self-consciousness over her technique.
Jessica/Karthik	Karthik attempted a reasonably high approach vector, which appeared to throw Jessica: rather than matching his vector by starting low and meeting in the middle, she mirrored him by starting high, resulting in a premature collision.
Jessica/wadapan	Despite having been primed to observe the handshake, in its immediate aftermath, neither participant appeared to have any recollection of what had just happened. Based on Karthik's account as a bystander, it appears that wadapan went in for a formal, horizontal handshake, which Jessica met with a fairly loose grip, allowing wadapan to lead the movement, resulting in a perfectly unremarkable handshake.
conq/wadapan	wadapan had still not expanded his repertoire beyond the straightforward "salesman" handshake. This style affords few opportunities for creativity but he was trying to make the most of it, lampshading his vigorous motions by adding pirate sounds. Thankfully conq was able to read what was happening and actively engaged with the theatre. The handshake was memorable and a success.
Georgia/Karthik	Georgia suggested a general-purpose "unisex" cursory handshake, wherein participants engage perpendicularly, give a light squeeze, and disengage without further motion. This appeared natural and friendly.
tubes/Karthik	From across a seating area, tubes demonstrated a long-range handshake (Figure 2), solely interlocking fingertips without engaging palms. This handshake communicated a willingness to meet the other person where they are.
Naomi/wadapan	Naomi demonstrated the "creepy handshake". In this example, she clasped wadapan's hand in both of hers, and initiated a controlled, slow up-and-down to exert dominance. Perturbed, wadapan physically retreated in discomfort (Figure 3). This replicates Nagy et al.'s finding that "following a prolonged handshake [...] participants showed less interactional enjoyment, as indicated by less laughing. They also showed evidence of anxiety and behavioral freezing".
Jeredino/wadapan	Jeredino has a gregarious and genial disposition and was able to confidently lead a classic "pound hug" (Figure 4).



Figure 1



Figure 2



Figure 3



Figure 4

Edge Cases and Error Handling

The most common handshake failure mode is when person A intends a fist bump but person B, ignoring all signaling, blindly approaches with an outstretched hand. This is a precarious situation that is susceptible to livelock. If both participants simultaneously try to adjust for the other, they will cycle between alternating states—(fist, palm), (palm, fist), (fist, palm)—ad infinitum until one dies from exhaustion.

The standard protocol in a fist-hand livelock is randomized backoff. In most natural situations there is a natural difference in panic levels and reflex speed between participants which will delay state transitions enough for there to be coordination. But in the rare cases where this is not true, prior defaults must be established. We recommend that the participant intending a fist bump recovers by instead punching the other participant in the face and never interacting with them again.

Towards a Globally Optimal Handshake

Considering the intractable difficulties in reading, signaling, and coordinating interactions between participants, we have taken inspiration from well-documented handshake protocols in computer network communication, to offer an effective solution for handshake coordination invariant of geographical background and charisma levels.

The core of the protocol is narrating exactly what you are doing as you are doing it. The silent handshake occurs in a dense fog of war where participants, especially autistic participants who have a reduced ability to recognize non-verbal cues, are doomed to fail. It is only the pure light of the spoken word which can penetrate this darkness and restore trust and harmony to the American social landscape.

We tested our protocol under laboratory conditions (Figure 5) and now present a transcript of successful execution.



Figure 5

Person A: “Initiating handshake request.”
Person B: “Accepted.”
A: “40 degree carpal tilt, 60-80 N peak grip force, .2m/s, transition to half-hug.”
B: “Ack. Initiating movement.”
A: “Ack. Initiating movement.”
A: “Person B your hand seems to be at a 30 degree tilt. Are we on the same page? I repeat, are we on the same page?”
B: “Packet dropped, please repeat.”
A: “Person B, your hand tilt is off-mark. Abort and retry.”
B: *(sweating, unable to hear, approaching slowly)* “Packet dropped. Please repeat…”
A: *(sweating, veins bulging)* “Stick to the plan, stick to the plan...”
B: (delirious) “Contact established, initiating grip.”
A: “Person B, do not initiate grip, I repeat, do not initiate grip. Abort, abort.”

B grasps the ring and pinky finger of A.

B: *(voice trembling)* “Grip initiated. Proceeding to half-hug.”
A: *(looking to the left, noticing persons C, D, and E watching them intently)* “PERSON B, ABORT! Street credibility is declining rapidly!”
B: *(high alert)* “Course-correction! Disengaging thumb! Disengaging index finger! Disengaging middle finger! Person A, begin curl!”
A: *(desperate)* “Ack!”

A begins to bend ring and pinky finger.

B: “Pitch up! Match velocity!”
A: “Climbing!”

A and B’s ring and pinky fingers rotate into parallel, continuing until they are perpendicular.

B: “Target angle reached! Curl, curl, curl!”

A and B bend their fingers until they are tightly interlocked. It resembles a “pinky promise”, signaling covenant, but reinforced by the addition of both ring fingers, signaling strength. A sharp intake of breath is heard from C, D, and E.

A: *(tearfully)* “Stable coupling achieved! Person B, you are cleared for half-hug!”
B: “Ack. Bring it home, A.”

They embrace. But, like, only a bit. C, D, and E cheer wildly. Stock prices in A’s critical minerals company skyrocket. B’s enterprise electronics company’s GPU production accelerates. America achieves global hegemony.

Conclusion

We have synthesized our findings into a robust framework for unrehearsed handshakes, which we hope individuals and agencies will be able to implement in their own best practices. Following our data collection, we (the authors) were able to trial the handshake prototype (Figure 6).

Figure 6



Further study is needed.

Bibliography

Chaplin, W. F., Phillips, J. B., Brown, J. D., Clanton, N. R., & Stein, J. L. (2000). Handshaking, gender, personality, and first impressions. *Journal of Personality and Social Psychology, 79*(1), 110–117. https://doi.org/10.1037/0022-3514.79.1.110

Nagy, E., Farkas, T., Guy, F., & Stafylarakis, A. (2019). Effects of handshake duration on other nonverbal behavior. *Perceptual and Motor Skills, 127*, 52–74. https://doi.org/10.1177/0031512519876743

Schroeder, J., Risen, J. L., Gino, F., & Norton, M. I. (2019). Handshaking promotes deal-making by signaling cooperative intent. *Journal of Personality and Social Psychology, 116*(5), 743–768. https://doi.org/10.1037/pspi0000157

Stewart, G. L., Dustin, S. L., Barrick, M. R., & Darnold, T. C. (2008). Exploring the handshake in employment interviews. *Journal of Applied Psychology, 93*(5), 1139–1146. https://doi.org/10.1037/0021-9010.93.5.1139

Tilston, O., Holler, J., & Bangerter, A. (2025). Opening social interactions: The coordination of approach, gaze, speech, and handshakes during greetings. *Cognitive Science, 49*(2), Article e70049. https://doi.org/10.1111/cogs.70049

Wood, P. (2012). Snap: Making the Most of First Impressions, Body Language, and Charisma. New World Library. pp. 85–87.

TALES FROM THE CAMP: CELEBRATING BRINE SPACE

by Damon P-Sasi / Daystar Eld

Most smart, educated people would say they're empiricists to some degree. Science is good! Evidence matters! In practice, that usually just means trusting the expert consensus and reading the occasional article that cites a study.



And that's not terrible! Certainly better than most alternatives. But far fewer people practice empiricism deliberately in their daily lives.

I've spent 10 years teaching and directing rationality camps and workshops, and like to look for ways to give participants the chance to run actual experiments on questions they care about. Potion Making was my attempt at encouraging empiricism one afternoon at WARP 2024.

The camps

Since 2023 I've been directing WARP, the Wandering Applied Rationality Program, with the help of ESPR and SPARC staff, which are summer camps I've taught at since 2017. For those that don't know, these roughly 10-day programs are hard to summarize well, but are free camps for gifted high schoolers that create an environment that helps participants better understand the world, themselves, and each other by offering a variety of classes ranging from information theory and statistics, to Double Crux and Focusing, to classes like mine, Emotions Make Sense and Philosophy of Therapy.

The camps are a lot more than just the classes, however, which only make up the mornings of each day. The afternoons are unscheduled time for participants and staff to run more ad hoc activities like improv or puzzle hunts, and the evenings include everything from organized Fermithons to impromptu karaoke sessions. Rationality, as I use the word, refers to an epistemology that minimizes false beliefs and a system of actions that maximally fulfills someone's goals. It's the "theme" of the camps insofar as it helps create a common sense of purpose among the otherwise relatively diverse interests and focus among staff and participants: that ideas matter, problems can be solved, and the process of learning and trying new and difficult things can be fun, especially if you find the right people to do them with.



In 2024, with over 15 camps as an instructor under my belt, I took a second swing at directing one, this time at Queen's College in Oxford. About halfway

The traditional camp puzzle hunts are widely beloved, but I this camp lacked the staff members that would normally make one, so I decided it would be a good opportunity to try something new to replace that afternoon activity block.

I really enjoy experiential learning scenarios and LARPs (live-action role-playing games). I think anything that involves hands-on experience is more likely to impart lasting lessons, and in far more variety than any top-down intended experience would capture. The main "issue" I have with most things that fall into this category is that they tend to be a bit too detached from reality for my ideal.

This may sound strange coming from someone who tries to champion the learning value inherent to fiction whenever I can, and I do think a lot is learned from them regardless! But I wanted this camp to try having an activity that wasn't just about learning, but also experimentation and practice in new skills. Something more connected to the rest of their lives, so they'd have an easier time remembering that the stuff they're learning at the camp can be applied once they were back home.





My main criteria were:

- Something requiring learning true things about the real world (as opposed to artificial or abstract problem solving).
- Something encouraging introspecting on what they want (rather than being given a specific goal).
- Something that would reward working together (without requiring it).
- Something with a physically interactive component.
- Something that promoted self-improvement even beyond the activity.

I knew it would be hard to find something that satisfied all of these at once, let alone also being feasible to do during the camp. I held a babble-and-prune brainstorming session with the whole staff during our prep days before camp, then narrowed options down through more conversations with others until we hit on the idea of exploring “brine space,” as inspired by Slime Mold Time Mold which would become an event called “Potion Making.”

The main idea was that some of our staff who had a lot of experience and knowledge around nutrition and health could guide students in some safe self-experimentation. We wanted to specifically avoid conveying object-level facts or suggestions, and instead try to share generators and models for what careful progress can look like even amidst deep epistemic uncertainty (which nutritional health science, as a field, definitely qualifies as). There was also a lot of useful knowledge and wisdom that could be shared along the lines of “definitely don’t do X, Y, or Z.”

Personal health is a good place to practice what I call Deep Empiricism, which is about the feeling of personal agency in testing our beliefs against reality, and updating based on new observations or evidence. It's true that we can't test everything by ourselves; at some point we have to trust others to figure out things for us. But there are many things we can't reliably trust to others because they're about us as individuals, and people can be so different from each other that no amount of scientific research can capture everyone.

This is especially related to questions of tastes and preferences, productivity and habits, or personal health, both mental and physical. If no one seems to know something, or if the evidence for it seems mixed, or if there are different groups arguing over what's true, then it's up to us to figure it out ourselves. Even if some study found that 99% of people sleep better in a room between 60-67 degrees Fahrenheit, you still have to experiment for yourself to see if your ideal is outside that range.

So training that muscle, that feeling of "I can try to figure this out myself," is important, and to prepare we bought a bunch of fairly safe vitamins, minerals, and nutrients in powder form, as well as measuring cups, scales, lots of juices, teas, sugar, salt, etc., and set up a table where they were on display and able to be accessed under supervision. I asked participants to think of a particular thing about their life they wanted to improve—maybe sleep related, or energy related, or focus related, or something about their skin or hair, or some type of pain or discomfort they feel—and spend time researching whether it might be due to some sort of deficiency that they could supplement.

I also emphasized that taste is a factor too; if there's something that they could do to make the things they already drink or take more tasty for themselves, that can do a lot to improve their quality of life, and still engages the experimental/life improvement mindset. I wanted to make it clear that the point wasn't for everyone to try something "supplement" related; even creating a new tasty drink for themselves to enjoy the rest of their lives would be a win.

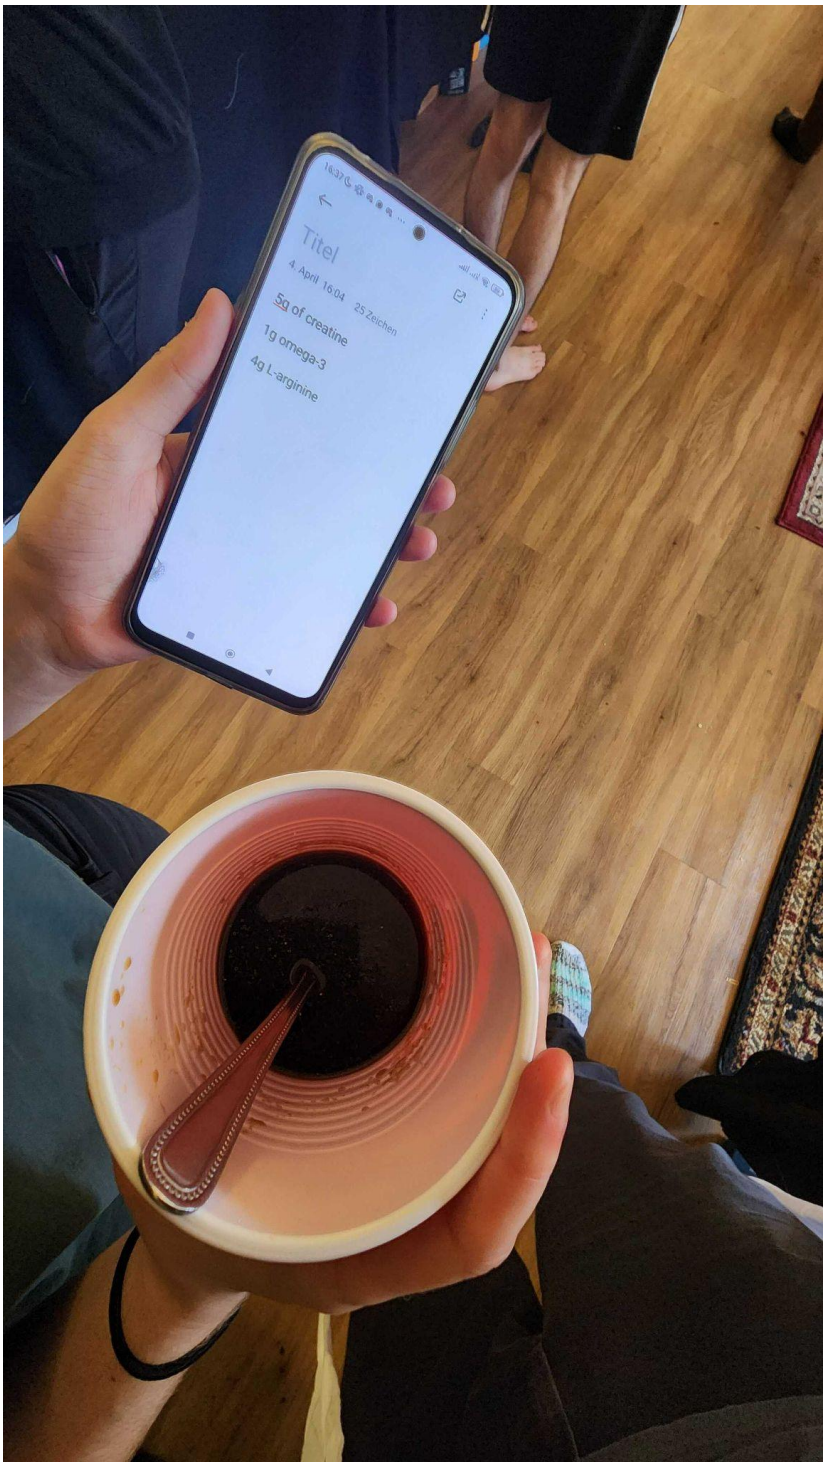

This covered most of the criteria for what I wanted a hands-on “experiential learning day” to include, and while there were some obvious and reasonable concerns about safety and camp optics among staff, after double checking my intuitions with others I decided the risks were worth owning. All camp events are optional, but we would stress this was super-extra optional, and provide another structured activity for those interested in doing something else with their afternoon.

To help students design their experiments, we gave a lecture on safe self-experimentation. It covered making your measurements as objective as you can, using your normal baseline self as the control, and watching for the three biggest pitfalls: regression to the mean, confounding factors, and the power of psychosomatic effects like placebo.

Going in, I decided I’d consider it a worthwhile experiment if at least half of the students enjoyed the activity, and at least a third of them got something of value from it.

How it went

On the post-camp survey, 18 of the 28 said they enjoyed Potionmaking, more than the half I'd set as the threshold for "was this worth doing." On the same question, 15 said they enjoyed learning Arduino (students could pick both if they attended both), and 5 said they'd have preferred another usual afternoon. Some notable experiments:

- A student tried to make a replacement for coconut water as a source of potassium, because he disliked coconut water's taste… only to discover that the taste of coconut water *is* largely from its potassium.
- Another student tried making a version of Red Bull without caffeine, since it gives her headaches. She instead learned that reading how to balance ingredients and what they're for is more difficult than it seems, and that she's particularly receptive to magnesium as a sleep aid.
- Someone just wanted to test what portions of milk and sugar make for the best coffee, and double-blind taste tested six carefully measured out variations. They learned that they do not, in fact, like it dark, and that more milk and sugar reliably made them enjoy coffee more.
- There was a student who made an "electrolyte shot" so they wouldn't have to drink a whole bottle of sports drinks to rehydrate, then worked to find the right balance of electrolytes and liquid to make it actually drinkable.

Overall, I'd say the event was a success, and I'm happy to do it again in a future camp (though some other scenario ideas could be fun too). None of what the students learned in the experiments above was something we set out to teach. Some of it, like the magnesium and the coffee, depended on their own bodies and tastes, so no one else could have given them the answer. The rest, like where coconut water's taste comes from, could have been a line in a lecture, but discovering things like that yourself is part of what ignites the spark of a true scientist.

If you do this yourself, make sure you only experiment with fairly safe ingredients, and consult with people who know what not to do. Ideally you learn something new about reality or yourself, but at the very least you also learn that the pursuit of knowledge isn't only reserved for distant people who belong to some special class, but something every person can contribute to.

PLEASE SEND ME SECRET KNOWLEDGE

Every day, MILLIONS of people around the world say the same thing: “Georgia, that’s proprietary information” 😩 or “Georgia, that’s personal” 🤫 or “I’m so sorry Georgia, you are so handsome and compelling, but the rules of my order forbid me from revealing this to outsiders.” 🙊 All of these mean the same thing: we won’t tell you that because we HATE you. 👉

It’s time for this cruel regime to come to an end. Esotericism is uncool and boring. 🤷 Tell ME what you know instead. I want to hear it. 🕴

– Georgia

- I want to know!
- Unburden yourself
- See what it would be like if you told me
- I’m ready
- I can handle it

WE WANT YOUR:

- Corporate secrets
- Esoteric truths
- Rumors
- Hot takes
- Insider knowledge
- Weird thing that happened to you once
- And more!

Q: Will you keep my information private? A: Probably not!

secretknowledgefan67@protonmail.com

Reach out now! **It’s free!**

soupsperiment

By Sahiti

So: *can* you taste a bay leaf in a soup bowl?

My estimable colleagues in this empirically minded mystery cult take time on occasion to ritually assert their affiliations to empirical study. Here in this publication dedicated to this most sacred of sacraments – Let us look at the world! Let us *categorize* it into *parts* and determine their relations, their correlations! Let us attempt to perceive the joints of the world itself, and when we have perceived them let us take out our knives and *carve*!!!! – I humbly submit the record of my own novitiate endeavours, in a domain I hold quite close to my heart, indeed quite close.

My stomach is, after all, right there, on the other side of my diaphragm. And while in the map of my body my nose and tongue may seem more distant, consider after all that on astronomical scales my entire body is smaller than our error bars for our state-of-the-art instrumentation.

Let us, in the spirit of *philia* and *ludus*, of *ethos* and *pathos* and *logos* in turn,

test each other!

1 experimental design

We've written and submitted this whole portion of the whole thing to *The Loop*'s preregistration committee before the first subject is ever tested.

Now, bay leaf aromachemicals, I'm told, are supposed to be fat-soluble, and so we must have a fatty solvent. This doesn't stop me from suspecting its efficacy can translate well to a water-based broth, because, well, one man's ingredient is another friend's recipe, and all it takes to emulsify an oil in a water broth is a little protein or starch to give it a li'l gel matrix to anchor to.

In these lights I construct our categories and reify them as CONCOCTIONS:

1. coconut milk, neutral oil, salt
2. coconut-milk-based soup
3. water-broth soup
4. water, cornstarch, neutral oil, and salt

A COOK will prepare each of the above concoctions, at first without bay leaves, and portion them out into individual servings. Then the cook will add bay leaves (first bloomed in the neutral oil) to the pots, heat for five more minutes, and portion them out once more into individual servings.

A SECRET-KEEPER will create a MARKER-BASED LEGEND, and mark each of these serving cups on the bottom rim said legend, such that it is invisible to the cook; and array them next to the pots in groups. They will also NUMBER THE CUPS.

The legend will work as follows:

- red-green-blue-black for each of the concoctions
- orange-purple for with vs without bay leaf.

After double-checking that the serving has occurred as intended, the secret-keeper will give the READY signal, and we will proceed into our waitng pit of hungry, sniffly, shivering TEST SUBJECTS where they lay unsuspecting in our TEST SETTING – that is, a room of friends and colleagues who have noticed that it is a rainy day and would perhaps be willing to fill out a SURVEY FORM if it comes bundled with some soup.

Said subjects will answer TWO SURVEY QUESTIONS on their respective SLIPS OF PAPER that they then fold up and submit to us. They will also hopefully refrain from identifying themselves in said slips of paper.

The questions are as follows:

1. Rate everything you tried!
    - 9 entries with fields for cup number and score.
2. What's the experiment about?

2 Results

Find the raw data in this gsheet.

We received 12 total entries; of these, two self-reported their disqualification in question 2, as requested, because they'd already heard about the experiment because they were in the room during its deisgn phase. (Though this is conjecture; they did not reveal their identities on the survey form.)

3 Analytical methodology

I'd like to formally apologize to Alexander Wales for running my mouth last night about whtehr analysis counts as empricism. This whole exercise is accursed and useless without a hefty layer of principled interpretation to even begin to make sense of what's happening. As it stands, getting anything I'd not laugh at calling a conclusion from these apltry numerical offerings is like unto squeezing blood from a stone. And yet, having gone out there and collected the data, I feel so strongly convinced, in my heart of hearts, that it *ought* to be able to tell us *something*. *Anything*. I have touched the bones of the world! I have stared at the heart of a flower and rubbed the yellow dust of its propagation between my fingers! I thought I smelled God in the perfume of a frying bay leaf! Shouldn't I have *learned something*?

Ahem.

Anyway.

I tried a group by and a mean of each group, aware all the while that each axis of grouping available to me (8 buckets, 3 axes, 12 – TWELVE – datapoints) was already *definitionally* near to about as riddled with confound as it was possible to be and still hope to learn anything at all. (This shows that people really like coconut milk, and not much else.)

Claude has informed me that the instrument I seek is not a linear regression, but a mixed linear model; machinery that GSheets in its wisdom has not seen fit to provide me. Why do they always lock up the power tools in places I'd have to vibe up properly if I wanted to get anywhere properly with them.

4 Conclusions and conjectures

Coconut milk is tasty. Vegans everywhere (those without coconut allergies, anyway) agree.

The rest will take a moment to confirm.

5 Qualitative analysis (i.e. my own opinions, which I get to call science because I'm consecrated into the clan of the qualified opinion havers now)

Folks conjectured wildly about what it was we were actually trying to test; absolutely nobody guessed bay leaf to be among the differentiators. This despite a faint greenish tinge in both the cornstarch and coconut milk concoctions, which I thought was a dead giveaway.

The closest guess was someone who asked if this was a leave-some-out ingredient attribution (not in those words), *without* specifying which ingredient.

Two people had guesses that it was about coconut milk. Seeing the data, I'm not surprised; the single biggest single-factor positive effect on flavor was coconut milk. Tracks, and fair. Nature's emulsions usually beat ours, when we're not trying very hard to beat them. In practice, I used the same mirepoix base for both broths; coconut was indeed a little something extra.

6 Problems with the design and execution

6.1 Methodological hiccups

- Our sample size was very very small.
- Respondent 6 rated one thing 6? Unclear what exactly was up there.
- Respondent 11 rated on a scale of 10, and asked me to normalize by dividing by 2. This will result in off-by-one-ish effects. Fortunately (?), they started from 0, which I've already decided to normalize to 1.

- Speaking of zeros –several other people rated 0s for some things. Fair. Some things feel like a zero. I should have given you all Likert bubbles. This was a fault of the form.
- I had to cut the experiment short because I'd forgotten I had an appointment. One of my gracious editors helpfully offered to clean up the mess we'd left on the island. I thanked him and dashed off, having entirely forgotten that I hadn't taken any note of the key on the bottom. An hour and a half later, I could be found fishing the cups out of the trash, dictating the number-color triads to a friend I'd flagged down.

6.2 Regrets in the late night

- Should we have been *more* clear about what was being tested? Or was this the correct level of obfuscation of intent? I'm torn. Perhaps I just dislike withholding information I think is cool; perhaps in subjective experimentation you only ever really get anywhere when you lead the human to the funky-tasting mostly-water with a few more steps.
- I found myself aesthetically displeased with the amount of wastage involved in setting this experiment up in a hygienic and safe manner. Individual tasitng cups for everybody, plus tasting spoons to taste, plus fresh utensils and cups for folks worried about flavour cross-contamination. . .
- Speaking of which, I would have liked to have a policy about flavour cross-contamination. I believe that this problem was adequately mitigated by letting people take their bepoke best shots at tasting.
- These are not the conditions under which one enjoys a bowl of soup. To do this *right*, I'd have to have hosted a dinner party – one in which, in fact, I provided a variety of soups, whose similarity to each other would only be suspiciously remarked upon by those not in the know. We'd have an army of helpful, near-invisible service staff, checking in on the diners' satisfaction with the meal as a whole, making close note of expressions, speed and quantity of consumption, unforced commentary. This whole affair would of course need a large and diverse enough sample size that we must of course host at a sutiable venue; and of course to gather *good* data about food, we cannot study it in isolation. We'd have to hold these monthly, inviting a great variety of people, on the thinnest of external excuses. Was this what hosting was, once upon a time? A/B testing on your undoubtedly clued-in guests?

- Actually, no, wait, what I'm actually realizing is that I already know who gets to do this in the modern day: !this sort of thing is a huge part of what makes working at an ad-tech shop fun.

6.3 Resource constraints

- It is currently 23:34 and I have to hit publish. Given more time I'd have written a shorter, tastier, better designed study.

7 Acknolwedgments

- Natalia, for doing the broth grocery run while I schemed and planned and underwent a preregistration review round and printed things out – and then doing another when it turned out that I'd overestimated the Lighthaven kitchens and assumed they just had cans of coconut milk on hand.
- Lighthaven, for the justified esteem I hold its kitchen in. This whole thing should not have been this easy to put together in an afternoon.
- wadapan, for acting as the secret keeper, creating the legend, and going back and forth marking up cups.
- Vishal, for chopping my onion (and stirring my soups while I cried at the onion *he* had chopped).
- The good residents of the building I was cooking these soups in, for tolerating a lot of logistical hullaballoo.
- And finally, the good editors of *The Loop*, for inspiring this entire exercise through an alchemical combination of nerd-sniping, peer pressure, environmental structure, and timely logistical and editorial support.

DO RATIONALISTS CARE ABOUT THEIR HEALTH?

by Amy 

I always assumed rationalists were health optimizers. I'm not a rationalist, and I kinda assumed that the people who reason carefully about everything would surely reason carefully about what they put in their bodies. And after all, isn't it "rational" to try to live longer, so that you can gain more "utility" from more years of life?

Then I arrived at Lighthaven, where the rationalist community holds its conferences and retreats.

The first thing I noticed were the plastic water bottles everywhere. Did people not mind the microplastics entering their bloodstreams?? And the snacks: kettle chips, Swedish fish, the kind of dried mangoes that have added sugar...

To be clear, I'm not judging. I consume plenty of microplastics and sugary snacks myself. But I was surprised.

To add fuel to the fire, I then heard a rumor that a bunch of rationalists were addicted to nicotine. Isn't that good for your short-term focus, but bad for your long-term health? (To be fair, they were apparently inspired by a post that Gwern wrote, citing a bunch of scientific studies about how nicotine gum isn't actually that bad for you... but still.)

That's when I began to wonder: maybe people consider it *rational* to take stimulants that boost their short-term mental energy, at the cost of their longevity. Especially if they believe there's a good chance that AI is going to kill us all within the next decade. Maybe under that framework, prioritizing longevity is actually irrational?

So today, I decided to send out a survey on this. I distributed it in the Lighthaven Slack and on Twitter. [ED: Here's the survey.]

Only twenty-nine people have answered thus far, but unfortunately—as I am here for a writing residency and will be kicked out if I don't publish my blog post by tonight—this is all the data we'll have to go off of.

how many respondents were rationalists?

I asked everyone how much they self-identified as a rationalist. Just under half -- 12 out of 29 -- rated themselves 5 or higher out of 7. I will consider those 12 as "rationalists."

do rationalists care about longevity?

I also asked everyone: how much effort do you invest in optimizing your longevity?

The people who scored themselves higher on rationality also reported more effort on longevity. With a sample this small I expected nothing to hold up, but this one correlation did (Spearman 0.38, p = 0.04) (thanks Claude for the analysis).

are they actually addicted to nicotine?

No. Most people just take caffeine. 4 respondents use nicotine; 6 take a prescription stimulant (e.g. Adderall). There was no significant difference between the rationalist and non-rationalist group.

conclusion

Unfortunately, the survey failed to shed meaningful light on my hypothesis. Maybe I'm totally wrong, or maybe I just need to talk to some more people.

If you are a rationalist and want to share any personal anecdotes on whether it's rational to care about your longevity, feel free to DM me :)

P.S. You can read my more serious stuff at giraffetales.substack.com

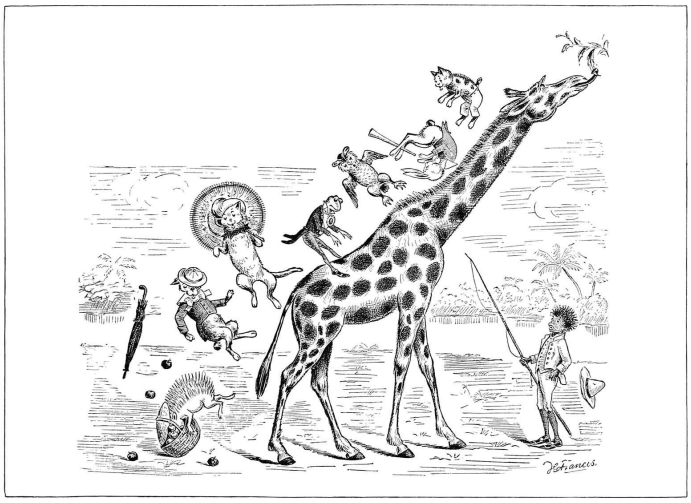

MYSTERIOUS DARK SODA EXPERIMENT

By Harjas Sandhu

Prologue: it was a dark and caffeinated night

When I was a boy, someone told me that they could really taste the difference between caffeinated sodas and non-caffeinated sodas, and that it was in fact the caffeine that made the difference between a good soda or a bad soda.

This was many years ago and I've since forgotten the context. But a burning conviction has remained with me during my metamorphosis into manhood, and it's finally time for me to disgorge it here. So here I go.

I don't remember who you are, guy who said this to me, but if you're reading this, I have but this to say: no, you dolt, you *definitely* can't taste the caffeine.

I think.

But though thinking is better than guessing, knowing is better still, and knowing is also what science is all about (citation needed). So for issue four of The Loop, I decided to investigate my hypothesis and put this issue to bed for once and for all.

The experiment

Diet Coke® and Caffeine-Free Diet Coke® have the exact same ingredients list. Of course, it's possible that the amounts of the ingredients are subtly different—in fact, this seems somewhat likely given that caffeine supposedly has a slightly bitter taste—but I'm not terribly worried about it; regardless of whether Diet Coke® and Caffeine-Free Diet Coke® do or don't taste different, I will consider this finding sufficiently interesting. (True Science Fans will recognize this method as Strong Inference, an advanced method of ~~copium~~ hypothesis generation and testing employed by only the most ~~egotistical~~ scrupulous scientists.)

The study method was roughly as follows. I got three identical plastic cups and labeled them with three Post-it notes: circle, triangle, and square. I then opened two cans of Diet Coke® and two cans of Caffeine-Free Diet Coke®—as close to simultaneously as possible in order to avoid any problems of fizziness differences—and one at a time, poured them into the cups. Triangle and square were always Diet Coke®, and circle was always Caffeine-Free Diet Coke®. I then invited a participant into the room and gave them roughly this spiel:

> Here are three sodas. Two of them are the same, and one of them is different. Your job is to identify the odd man out. You can taste, smell, or look at the cups in any order you'd like, as many times as you'd like. Please tell me which one you think is different and what you think tastes different about it.

I also asked them afterwards whether they were or weren't a caffeine drinker (which I did not strictly define, I know it when I see it, but roughly came out to "do you consume caffeine more than twice a week on average"). Finally, I rinsed and, uh, sort of dried the cups afterwards to avoid cross-contamination.



Now, I did only have a sample size of 19, which is probably why I didn't end up with a statistically significant p-value. In my defense, I spent three hours running this experiment and got tired. But I *do* think my observations were very interesting—particularly my qualitative observations—and that's the next bit.

The next bit

First, the quantitative observations.

Of the nineteen participants:

- 9 correctly chose the circle cup.
- 5 incorrectly chose the triangle cup.
- 5 incorrectly chose the square cup.
- 5/14 caffeine drinkers guessed correctly.
- 4/5 non-caffeine drinkers guessed correctly.

Now, I'm a little wary of cherry-picking. I'm reminded of the xkcd about the jellybeans: if you measure enough things, odds are you'll get at least one measure that looks significant if only due to the laws of cosmic coincidence. One in a million events are pretty rare, but there are also 8 billion of us...

That being said, I was not the only person to declare that they could taste a difference. Many participants said they were confident. Some of them were (kinda) right; consider Amy, who correctly identified the odd man out *before even tasting any of the cups* (she was using her sense of smell), or wadapan—the only person described in my notes as "super confident"—who said that all three sodas tasted different but that O was the most different by far.

On the other hand, some of them were wrong: consider Liam, who confidently claimed that the triangle was Diet Coke® and that the other two were full sugar; Damon, who initially felt very confident in his answer of triangle but ultimately decided to double-check his answers and ended up walking back his confidence levels; Valentino, who developed a theory purely based on sight (he unfortunately didn't explain before tasting the sodas) and then said that triangle felt "thinner"; or Alexander Wales, who after the experiment explained that he knows what caffeine tastes like because of his prior experiences with caffeinated mints—but guessed triangle.

On the other other hand, some participants entirely reported being unable to tell. Nadia asked me whether I was playing a trick on her and if they were all actually the same (she later explained that she had a bad sense of smell and that smell is half of taste); Karthik Bala spent a long time savoring each cup of soda before declaring that his pick was “kinda random”; and my notes for Jenn include the phrase “she is taking for-everrrrrrrrrr.”

These participants didn’t report a uniform experience: some reported the sodas all tasting identical, while others reported that the same soda cup would somehow taste different during different sampling events. But all three were similarly reluctant to take a guess, and all three probably would not have guessed at all if I didn’t push them into doing so—indeed, Karthik and Nadia only guessed after I explained to them that I would dutifully record their inability to discern between the options.



Conclusion

You can read my full notes in this spreadsheet. Ultimately, I think there's a small but significant chance that people can actually tell the difference between caffeinated and uncaffeinated soda. And I would feel slightly amenable to, but not very confident in, the proposition that non-caffeine drinkers are more sensitive to the taste of caffeine. But this sample was very obviously underpowered—a basic one-sided exact binomial test from Claude (don’t worry, I don’t know what this is either. Probably should’ve paid more attention in stats…) calculated a p-value of 0.146. In other words, under pure guessing, we would expect to get a 9/19 split about 15% of the time. That’s not nothing—but that’s not something, either. I leave the process of ~~p-hacking~~ amassing more data to a soul more patient than I.

Oh, and if you do this: please take detailed qualitative notes. The participant psychologies are far and away the best part of running these kinds of experiments. Despite all my complaining, I had a blast during those three hours, and there is a lot of gold to be mined from the depths of these participants' psyches.

For example:

- Why did some people feel so confident about different tastes?
- Why did some people feel like the sodas were actively changing their tastes?
- What was up with some of the taste reports?

Thank you again to everyone who participated for contributing to science. And special thanks to the illustrious Slime Mold Time Mold for encouraging me and providing me with assistance throughout the process. This wouldn't have happened with you!

Addendum: probability is (emotionally) hard

Perhaps my most interesting finding from this experiment was that humans do not emotionally resonate with probabilities—even when they are explicitly trained to be better about it.

I have been thinking about probability for a long time; I briefly studied to be a quant before I realized I was being stupid, I've been reading rationality blogs for years, and I also honed my probabilistic decision-making skills during my one year foray into Riichi Mahjong. I am fully aware that improbable events happen constantly. I very often explain to people that the difference between fictional worlds and the real world is that fictional worlds are under obligation to make sense, while the real world has no such constraints.

But even I still crumbled under the weight of unlikely (but not actually that unlikely) odds.

Let’s look at the full sheet.

Person	Order	circle	triangle	Square
ANSWER KEY	0	caffeine free	diet coke	diet coke
Harjas	1	1		
Sophia	2	1		
Iggy	3			1
Emma	4			1
Amy	5	1		
Jenn	6			1
Monofrogue	7			1
Alec	8	1		
SMTM	9		1	
Karthik Bala	10		1	
Liam	11		1	
Wadapan	12	1		
Georgia Ray	13	1		
Yonatan	14	1		
Mond	15	1		
Wales	16		1	
Zoe	17	1		
Nadia	18			1
Damon Sasi	19		1	
Total		9	5	5

On its face, these seem like pretty reasonable results. Nine circles, five triangles, five squares. This could very much be within the realm of random chance, but it also seems vaguely suggestive—maybe humans really can tell the difference between Diet Coke® and Caffeine-Free Diet Coke®.

Now consider just the first 8 tests.

Person	Order	circle	triangle	Square
ANSWER KEY	0	caffeine free	diet coke	diet coke
Harjas	1	1		
Sophia	2	1		
Iggy	3			1
Emma	4			1
Amy	5	1		
Jenn	6			1
Monofrogue	7			1

By trial 8, I was—pardon my French—shitting my pants. I started discussing these results with my participants, and they all agreed that it was really weird that no one had yet picked the triangle cup.

It took all my willpower to not change the way I was running the experiment and stay the course. I ended up having to convince myself of some crazy cope, like if maybe people actually just prefer the circle and the square over the triangle, and we can discover this in some follow-up trial in which all three cups are filled with water or something. After all, if we assume random guessing, the chances of participants not picking triangle 8 times in a row is roughly (2/3)^8=0.039=3.9%. That's pretty low!

But I did stay the course. And then the triangles started rolling in. Just imagine how embarrassed I would be right now if I'd changed the experiment.

That's right, Reader, I wouldn't be embarrassed at all, because I *never wouldn't have realized my mistake!*

This may not seem as dramatic as fabricated data or rerunning experiments until you find a result that you like. But Reader oh Reader, this is *still a form of p-hacking*. You can't throw out weird results just because they make you feel weird. And you definitely can't change your methodology partway through an experiment, on pain of ruining your data (or at least wasting some of your precious samples!) The only thing you can do is think really hard during your experiment design and then stay the freakin' course!

So put your chin up and accept that your hypothesis may be too limited to capture the weirdness in advance of seeing the data. Go forth into the unknown-but-hopefully-knowable.

And always remember: thinking is better than guessing, and knowing is better than thinking—but courage is best of all.

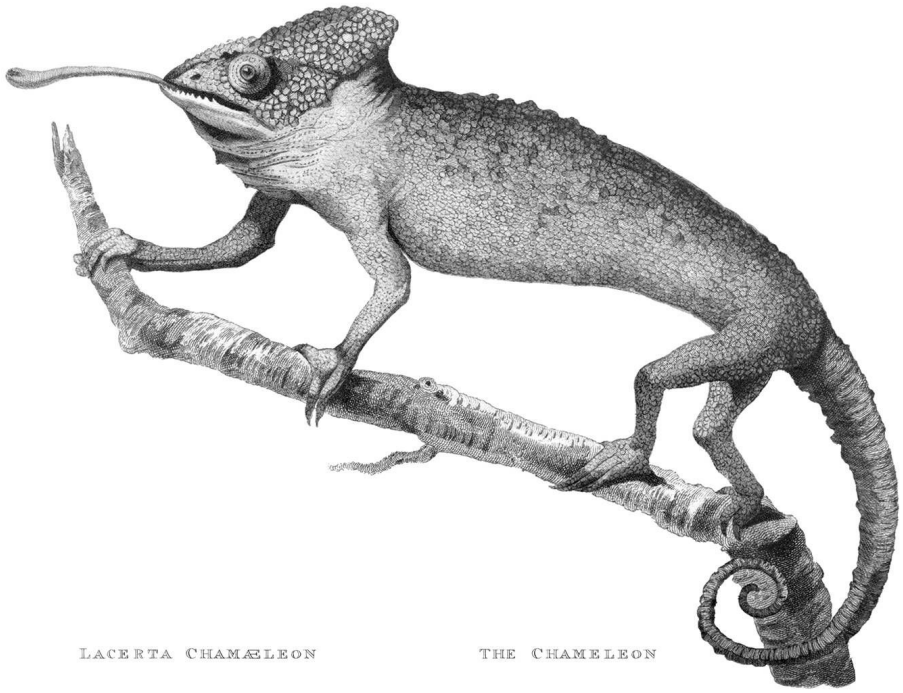

LACERTA CHAMÆLEON THE CHAMELEON

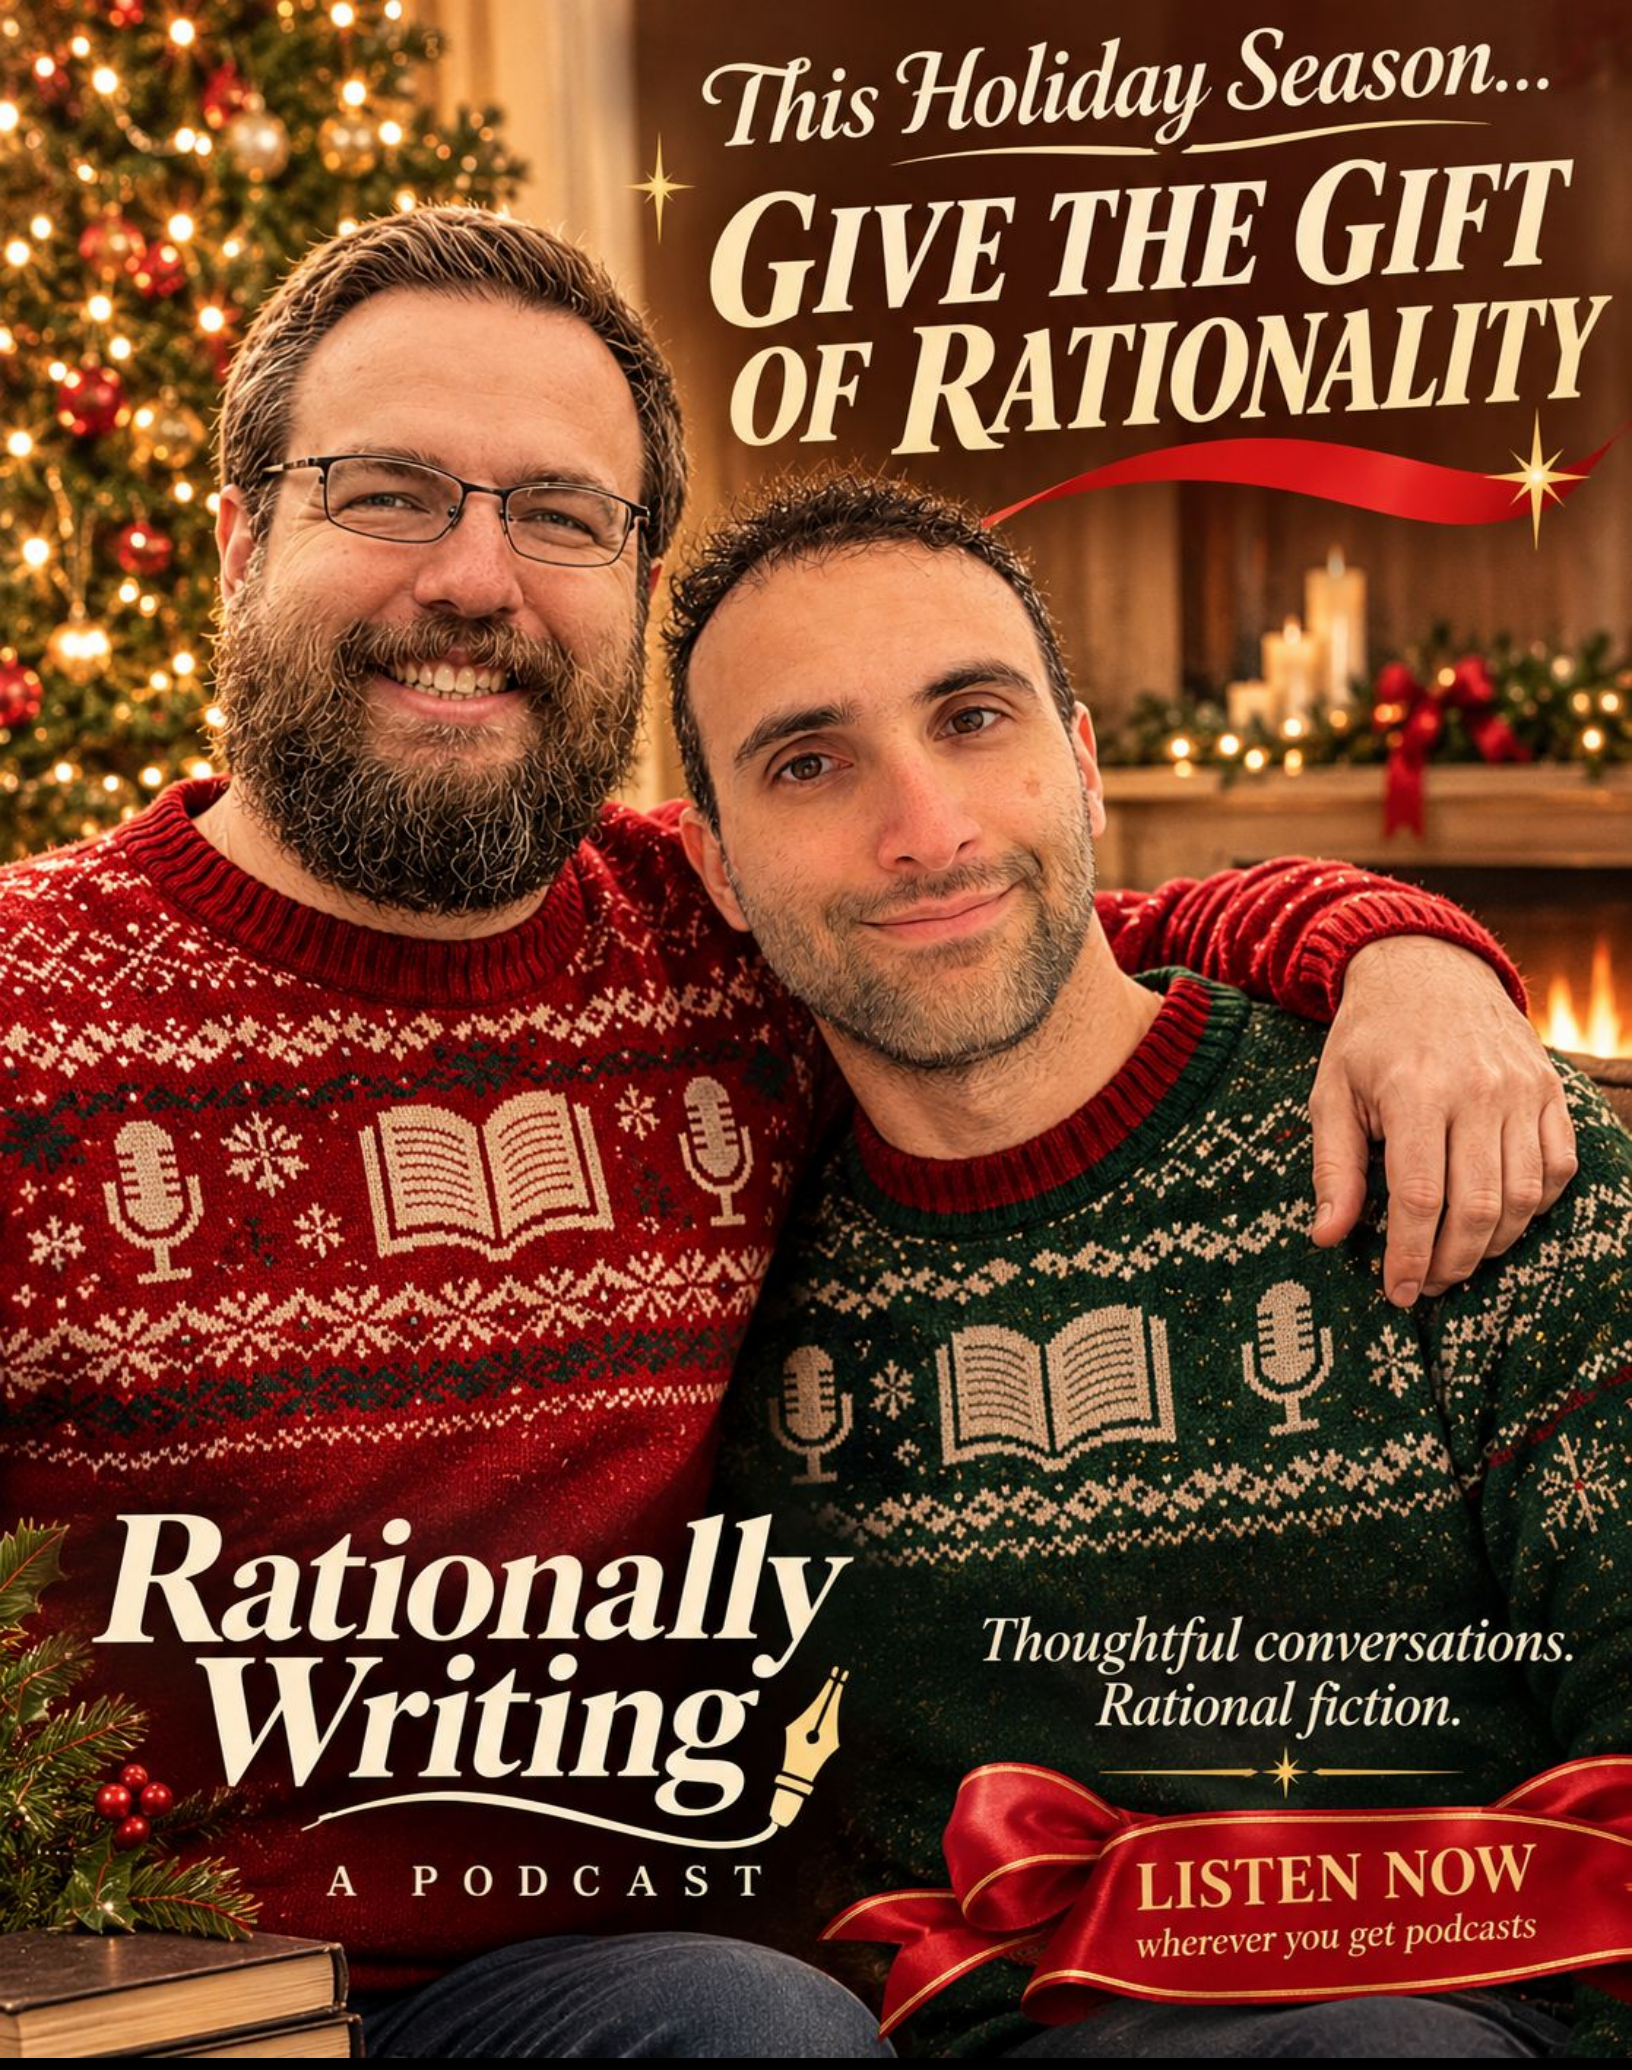
This Holiday Season...
GIVE THE GIFT OF RATIONALITY
Rationally Writing
A PODCAST
Thoughtful conversations.
Rational fiction.
LISTEN NOW
wherever you get podcasts

FOR THE... SABRINAS?

by Daniel Hails

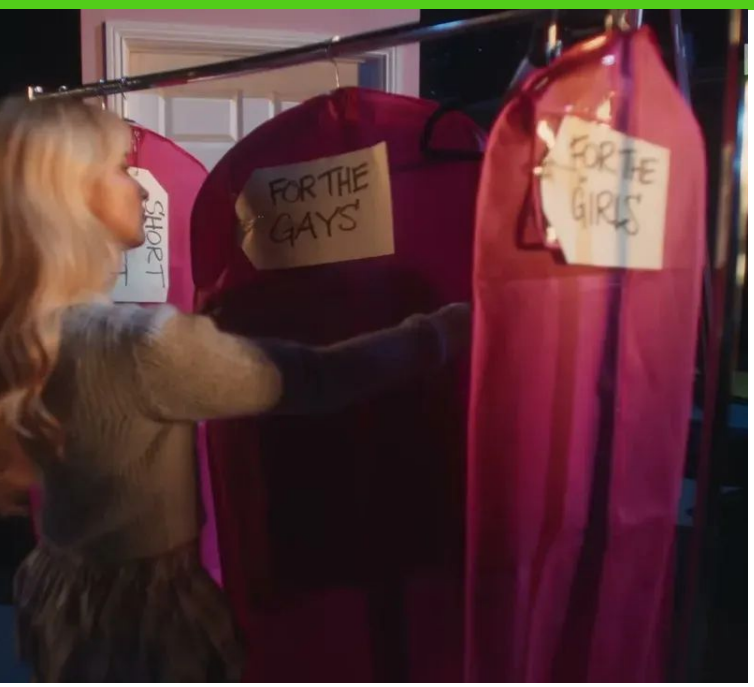

In A Nonsense Christmas, Sabrina Carpenter chooses her outfit from three hangers, labelled ‘For the Girls’, ‘For the Gays’, and we must assume the empty coat hanger is ‘For the Guys’.[1]

She picks the Gays.

It’s strange, then, that Sabrina has never been perceived as anywhere on the queer spectrum. In Taste, she makes out with Jenna Ortega - which is far more ‘evidence’ than most Gaylors[2] have. Indeed, why does she appear to be the ‘only’ straight pop star — or rather the only one taken at their word?

I suspected the reason lay in her music videos.

In particular, how she uses her, and the camera’s gaze. To establish this with certainty, we take all of her music videos and convert them into movie scripts, each with explicit scene directions and gaze annotations, so we can compare like for like. We then analyse each to rate the relative ‘emotional engagement’ of the gaze.

WHERE AI FELT SOMETHING

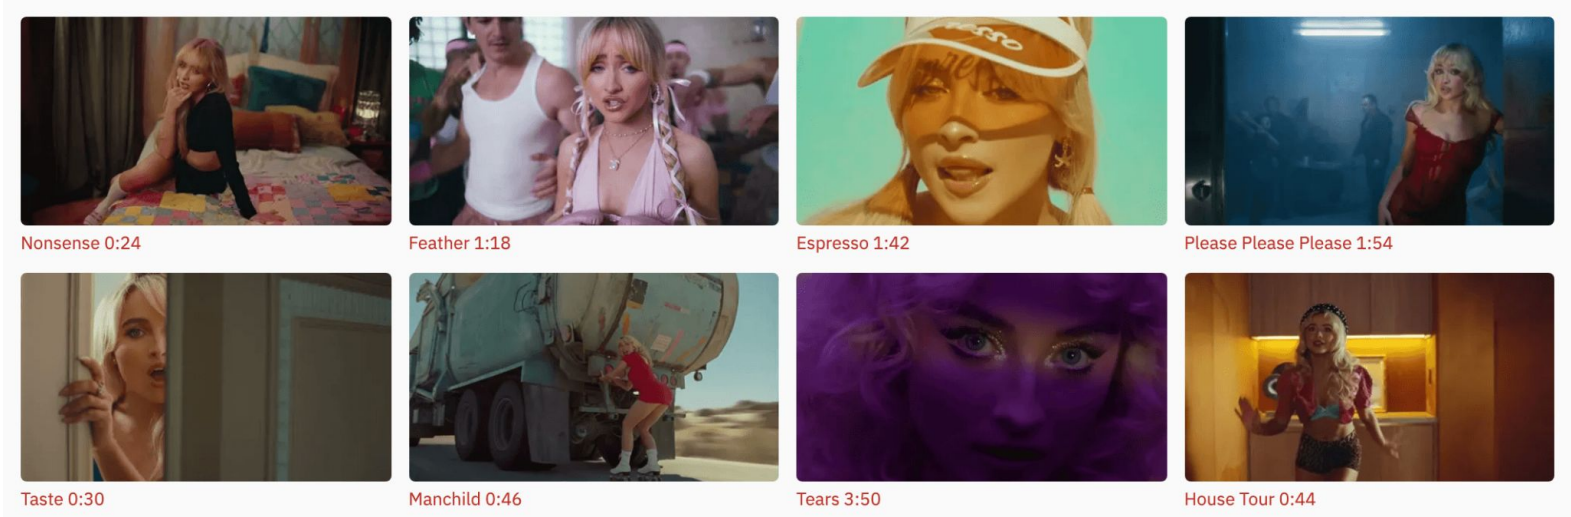

1. **My partner deserves credit from this concept!**
2. **The fandom that Taylor Swift is "closeted sapphic woman forced to conform to a heteronormative society".**

Leaving out *Please Please Please* for now, our initial results were as follows:

Object	Shots	Share of Looks	High Intensity
Camera	152	40%	55%
Off-frame	106	28%	19%
An object	38	10%	21%
Herself	28	7%	29%
Friends/Crowd	26	7%	42%
A man	19	5%	37%
Another woman	15	4%	40%

You might think that our journey has been cut short - she has the most positive regard towards the camera then about evenly split between friends; men; women and trailing up is herself.

Looking closer at the data we discover the only warm looks at a man (outside Please Please Please, which we'll come back to) are in Nonsense - where she laughs on the couch with him; and that man is Sabrina in drag.

Adjusting this to consider Sabrina in drag as Sabrina; we find that across the seven videos she looks at a man 19 times: fewer than she looks at herself (28), and barely more than at other women (15).

In addition; if we once again "Look at our data", we discover some curious motifs that make up a Sabrina Carpenter video: mirrors; lipstick and death.

SABRINA'S REFLECTION

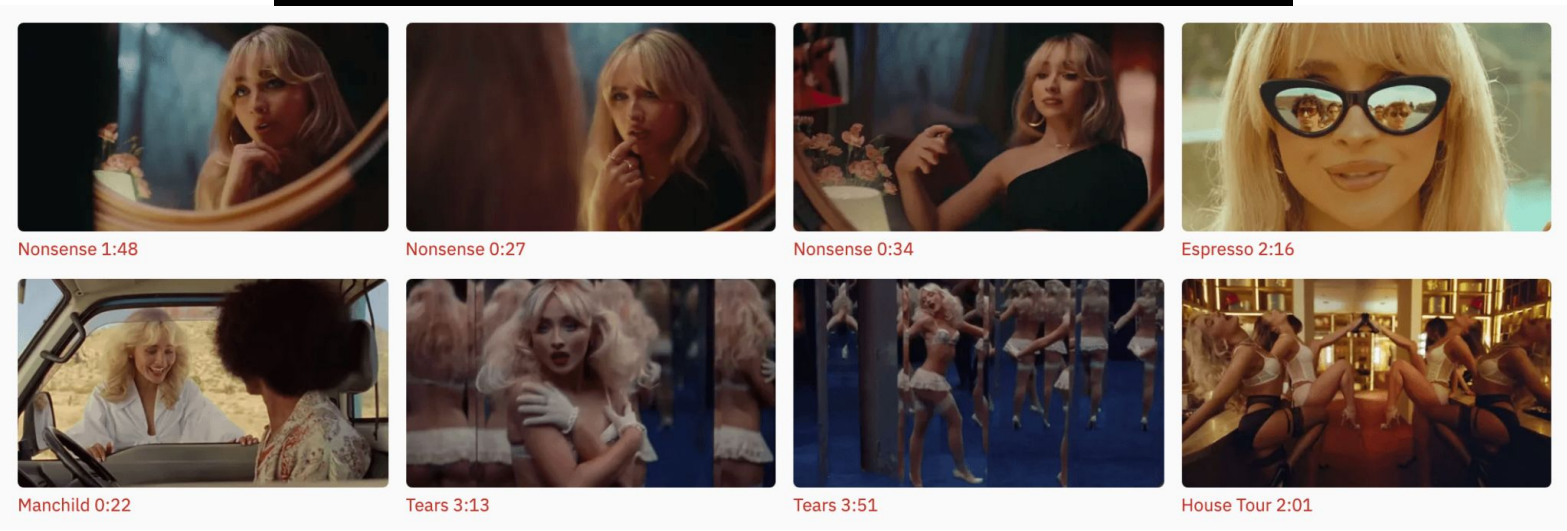

Every Sabrina video contains mirrors and lipstick.

SABRINA'S RED LIP

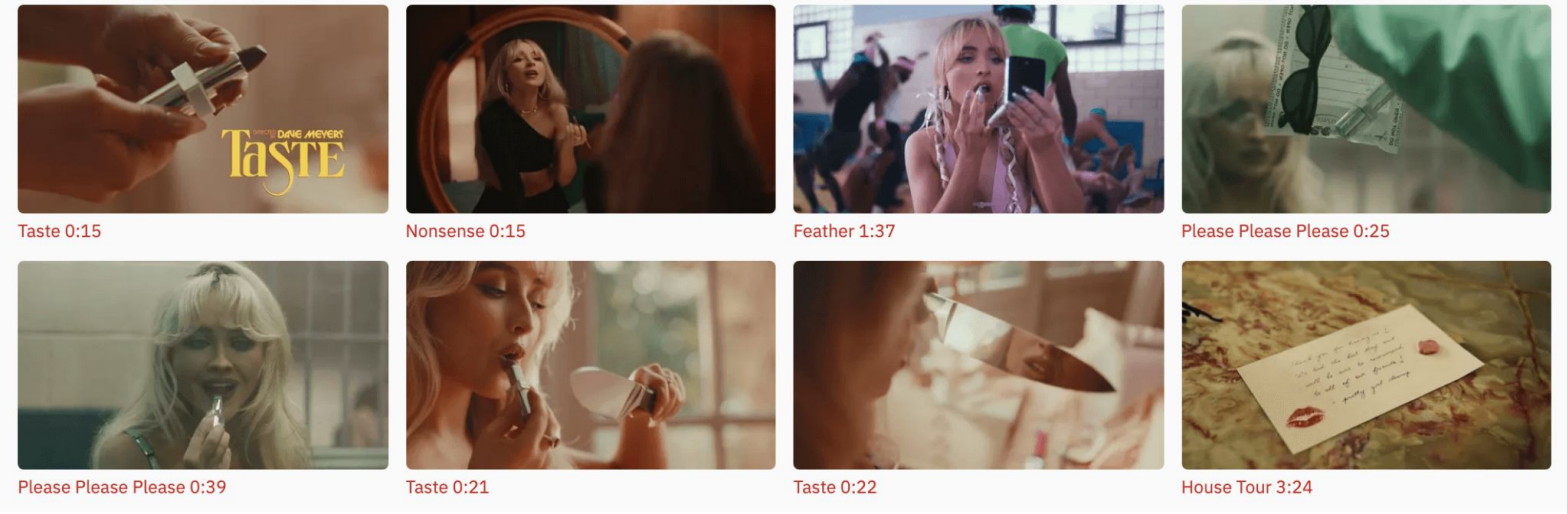

Every video also contains men dying:

WHAT HAPPENS TO MEN

BEFORE



Feather 2:24

BEFORE



Feather 1:18

BEFORE



Taste 2:20

BEFORE



Tears 0:04

WHAT HAPPENS TO HIM



Feather 2:36

WHAT HAPPENS TO HIM



Feather 1:52

WHAT HAPPENS TO HIM



Taste 2:43

WHAT HAPPENS TO HIM



Tears 4:22

WHAT HAPPENS TO WOMEN



Taste 2:24



Taste 2:46



Taste 3:00



House Tour 1:22



House Tour 2:04



House Tour 2:18

So, we see that men are introduced and then promptly - often the very next scene - killed off. There is genuine emotional content in Sabrina's gaze, but it often comes with a knowing glee as they promptly self-immolate.

Glances at women meanwhile consume all the emotional warmth. In *House Tour,* we see Sabrina as the motherly mastermind behind the break-in; indulging in her friends passions (reading), supporting them financially (gifting them stolen jewelry) while also remaining in the power position (being pushed around on the piano).

These scenes often have the trappings of sexual interest, but the gaze doesn't show the same. Even the most explicitly sexual moment in all eight videos (the kiss with Jenna Ortega) starts with Jenna kissing the man before she turns into Sabrina and, even then, Sabrina averts her gaze. Closing her eyes.

Could this just be an artefact of the summer's pop trends? To check we apply the same analysis Taylor Swift; Chappell Roan; Billie Eilish and Britney Spears..

Share of looks, top partner video left out	Sabrina	Taylor	Chappell	Billie	Britney
Into the camera	40%	29%	42%	47%	58%
At a partner	5%	11%	10%	0%	3%
At herself	7%	12%	2%	0%	3%

Here we see that Sabrina's focus on the camera is within normal bounds. Taylor Swift likewise looks at herself more than anyone else analysed; but unlike Sabrina, it isn't through the medium of mirrors. Instead she uses doubles, past selves, which she slaps, shoves and drinks with. She epitomises self-criticism, not self-regard. While Sabrina checks her reflection - and playfully winks to camera - Taylor cross-examines herself. If we adjust for her lack of looking at the camera; we once again puts Sabrina in first place for self-regard.


Look What You Made Me Do 3:42


Look What You Made Me Do 4:00


Anti-Hero 1:53


Anti-Hero 1:55


Anti-Hero 1:57


Anti-Hero 4:54

Billie and Britney - despite having very different styles - have very similar gaze characteristics. They have limited emotional engagement with other characters and themselves. Even when Britney does have a mirror as in Lucky 1:42 she spends little time looking at herself, instead looking up at the skyline; or ignoring the mirror completely.

Please, Please, Please Sabrina

The sole exception where chemistry turned outward rather than inward was “Please Please Please”. This one video holds 29 of her 48 looks at a man, 37% of its shots against 5% across the other seven music videos.

Her looks to camera drop by more than half, from 40% to 19%.

This video is a huge outlier in our otherwise perfect set -- how is it that the emotional intensity so dramatically increases -- and towards a man no less. Turns out the trick is casting your actual boyfriend as the main romantic lead -- Barry Keoghan .[3]

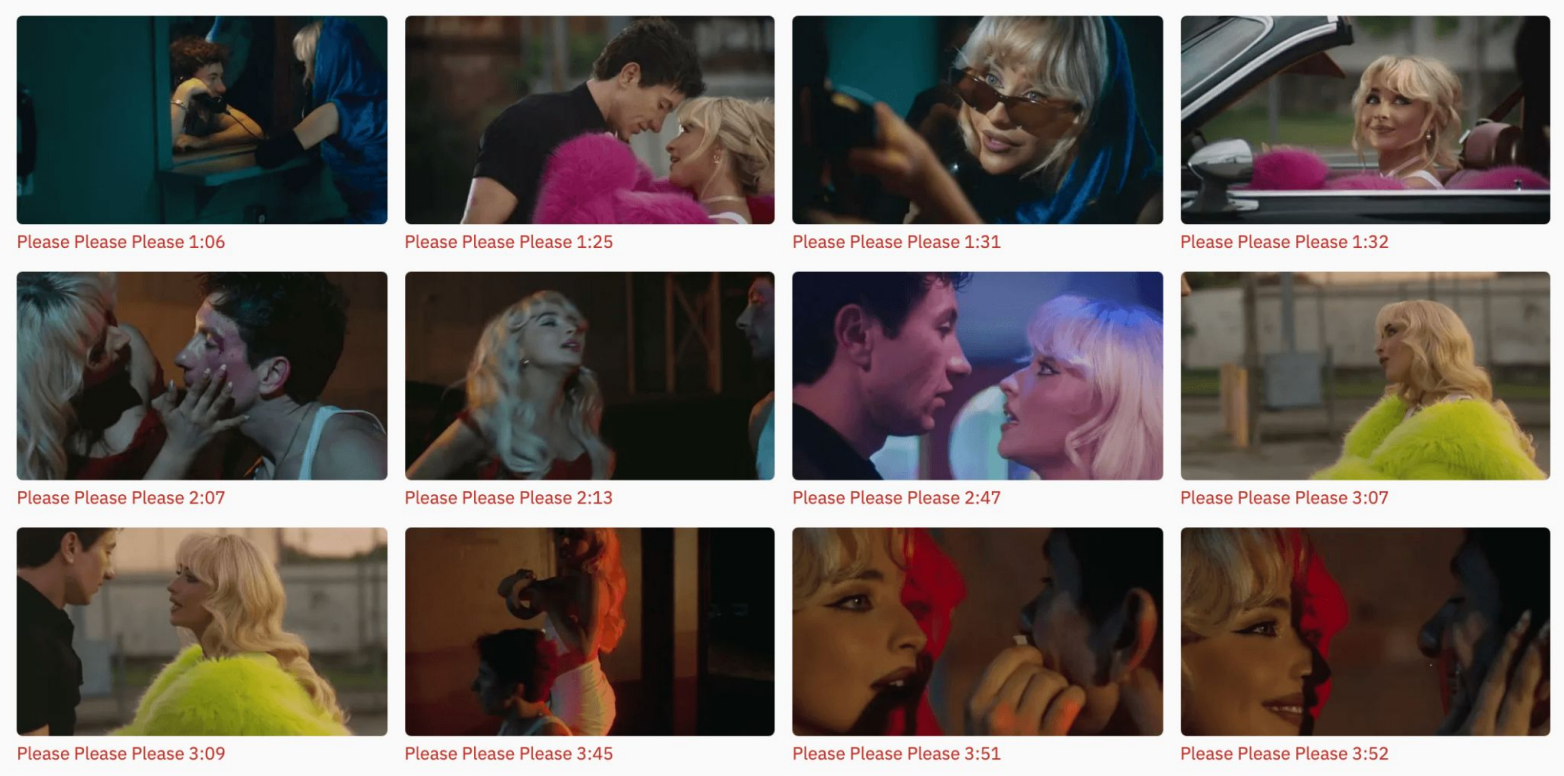

Suddenly it’s even more obvious what we’re missing in the other scenes -- perhaps the real intensity measurement is not eye gaze, but instead lip proximity.

One for the next study.

3. Technically ex-boyfriend now; but man that does not stop the chemistry.

THE LOOP

CONTRIBUTORS

❀❀❀

Editorial Content and Design by
SLIME MOLD TIME MOLD

Half-Fortnighter | Ben Grossman
Water Wizard | Paul Kreiner
Juan Ponce de León Impersonator | Valentino di SMTM
Potions Professor | Daystar Eld
Book Critic | Odessa Paige Turner
👉 | Karthik Bala
🤙 | wadapan
souperstar | Sahiti
Bungee Jumping Instructor | Hugo First
Rat Health Inspector | Amy
Dr Diet Coke | Harjas Sindhu
Pop Gazer | Daniel Hails
Intern | Igor Beaver

❀❀❀

Volume II, Issue 4, November 2026

Visit looploop.blog for more

KWARDS POOL BACKWARDS POOL BACKWARDS POOL BACK
KWARDS POOL BACKWARDS POOL BACKWARDS POOL BACK

BACK COVER

looploop.blog