

The Yes Column

I handed my creative judgement to a machine. It showed me how ideas really get rejected.

Boğaçhan Ülker

MSc Advertising & Creativity, TU Dublin · September 2026

For people who make, judge or buy creative work.

The challenge

Advertising runs on a belief so common that nobody states it: creative judgement is instinct. A good creative director can feel that an idea is wrong long before they can say why, and the feeling is the expertise. “It’s not it. I can’t tell you why yet” is said in every agency, every week. I have heard it in an internship, a newsroom and a classroom, and said it myself.

I want to question that belief. Not the part about seniors having better judgement; they usually do. The part where the judgement cannot be put into words. My suspicion is this: *a standard you cannot write down is not judgement. It is habit.* Judgement can be explained, argued with and handed to someone else. Habit can only be repeated.

Paul Arden built a whole book on the instruction to think the opposite of whatever you think. So I handed my judgement to something that cannot guess.

The experiment

I took two live briefs from this course and, instead of working them myself, built a small agency out of AI agents to do it for me. The first was Saint Vincent de Paul Ireland, *The Ripple Effect*: one small act of local help changes a household out of all proportion. The second was Currys around the 2026 World Cup: move the retailer from “sells televisions” to “makes your home the place people choose to watch”.

Each build had the roles a real team would have, including a Creative Director whose job was to review the work, score it, kill it or approve it, and write down the reasons.

The part of the design that matters most is what I did with myself. I set up the structure, and then I removed the approval points. I wanted to see how far the system's reasoning would go with nobody standing over it. Its decision log records it in one line:

```
The user handed all creative control to the agency and removed the approval points.
```

Two complete campaigns came out, each in roughly forty-five minutes. Currys landed on a platform called PEOPLE OVER. SVP landed on Breathing Room: financial pressure as a plane pressing down on a household, and one small act of support opening just enough space for them to make their own next move.

I read them two or three times. They had reasonable parts. Both were defensible. And I decided not to take either further, for a reason I could only put in the phrase this industry now uses a dozen times a day: too AI.

The results

I stayed with that phrase, because it was exactly the judgement I had set out to test. Forced to be specific, the objection came out as this: the Creative Director was too rational. Where I would expect a team to be testing different tensions and chasing the route that scares them a little, the system had moved, every time, to the most logical option.

That was my whole objection, and it was no more useful than "it's not it". So I went back to the archive and asked a simpler question. What had I actually told the reviewer a good idea was?

I had not told it. Not in my own words. For Currys I let the system write its own job description. For SVP I asked a different AI to draft the reviewer's brief, read it once, approved it and passed it on. I did this deliberately, for a clean experiment, and I was comfortable with it: I had set the preconditions. What I had not noticed was what they contained.

Reading the criteria two months later was the moment the experiment turned. Brief fit, distinctiveness, brand ownership, feasibility, "could a competitor sign this line". Every single one is a reason to say no. Not one line asks: does this surprise me? Does it frighten me a

little? Would I be jealous if another agency had made it? One rule I signed off on was pure suspicion: at most one approval in the first round; if you are inclined to approve two, you are reviewing too kindly. I had written doubt into the rules and never once written want.

The system did exactly what those rules told it to do. On Currys, the route with the highest emotional voltage scored five out of five on that criterion. PEOPLE OVER. scored four and won, and the reviewer explained why with perfect clarity: the higher-voltage route made the product secondary and risked implying that a purchase buys popularity. Both objections are correct. Both are reasons to say no. The reviewer even wrote the sentence that sounds most like judgement:

The score supports rather than makes the decision.

True. But what made the decision instead was, every time, the logical objection. Nothing in the system had permission to say: I know the risk, and I want it anyway. My own notes from the day of the run had named the system's danger as sterility. I read that in July as a fact about the software. It was a fact about the instructions.

The learning

The lazy conclusion is that the machine cannot do the magic part. The archive does not support it. The system did not fail to be creative; it failed to optimise for a criterion nobody gave it. When I removed the approval points, I did not remove a rubber stamp. I removed the only part of the process carrying the unwritten half of the standard: the half that says yes.

That is the finding, and it is not about artificial intelligence. Every creative department runs this experiment, slowly, every day. A senior hands a junior the ways to be wrong, keeps the ways to be right in their head, and wonders why the work comes back correct and dull. The junior has been given the no column and nothing else.

It is not that the yes column cannot be written. D&AD's *The Copy Book* is fifty copywriters explaining, in plain sentences, how they get to a line: a whole book of judgement written down. Rick Rubin, in *The Creative Act*, treats taste as attention, something you notice rather than own. Noticing can be described. We simply do not bother, because "I can't tell you why yet" is accepted as an answer.

So here is what I now do, and what I would ask you to try.

Write the yes column. For every brief, keep two lists. The first is the one every creative already keeps: the ways this idea could fail. The second is the one almost nobody keeps: what would make me want this idea even if it failed one of the first tests. On Curry's mine would have read: I want the route that makes the host slightly nervous. On SVP: I want the version where the help is almost invisible and the household is almost too proud. Either would have changed the outcome.

Give it to your team before the work starts. As sentences, not a mood board. They will still bring you defensible work. They will also bring you the other kind, because you have told them it is allowed.

Treat "I can't tell you why yet" as a debt, not a credential. Come back to it after the work has shipped. The second attempt is usually a real criterion.

The fear in this industry is that machines will do the creative work. The system I built could not out-create me. It could only out-execute the half of my standards I had bothered to write down. The risk is not the tool. It is unwritten judgement, in whichever hands happen to be executing it.

Write the yes column.