

[The new dot com bubble is here: it's called online advertising](#)

Jesse Frederik, Maurits Martijn, *The Correspondent*, November 6, 2019

In 2018 \$273bn was spent on digital ads globally. We delve into the world of clicks, banners and keywords to find out if any of it is real. What do we really know about the effectiveness of digital advertising?




Illustration: Jenna Arts for *The Correspondent*

Sometime in June 2003, Mel Karmazin, the president of Viacom, one of the largest media conglomerates in the world, walked into the Google offices in Mountain View, California. Google was a hip, young tech company that made money – actual money! – off the internet. Karmazin was there to find out how.

Larry Page and Eric Schmidt, Google's founder and its CEO respectively, were already seated in the conference room when co-founder Sergey Brin came in, out of breath. He was wearing shorts. And roller skates.

The Google guys told Karmazin that the search engine's earnings came from selling advertisements. Companies could buy paid links to websites that would appear at the top of users' search results. And Google worked as a middleman, connecting websites with ad space to advertisers eager to get their banners seen.

Schmidt continued: "Our business is highly measurable. We know that if you spend X dollars on ads, you'll get Y dollars in revenues." At Google, Schmidt maintained, you pay only for what works.

Karmazin was horrified. He was an old fashioned advertising man, and where he came from, a Super Bowl ad cost three million dollars. Why? Because that's how much it cost. What does it yield? Who knows.

"I'm selling \$25bn of advertising a year," Karmazin said. "Why would I want anyone to know what works and what doesn't?"

Leaning on the table, hands folded, he gazed at his hosts and told them: "You're fucking with the magic."

Is Google hacking your brain?

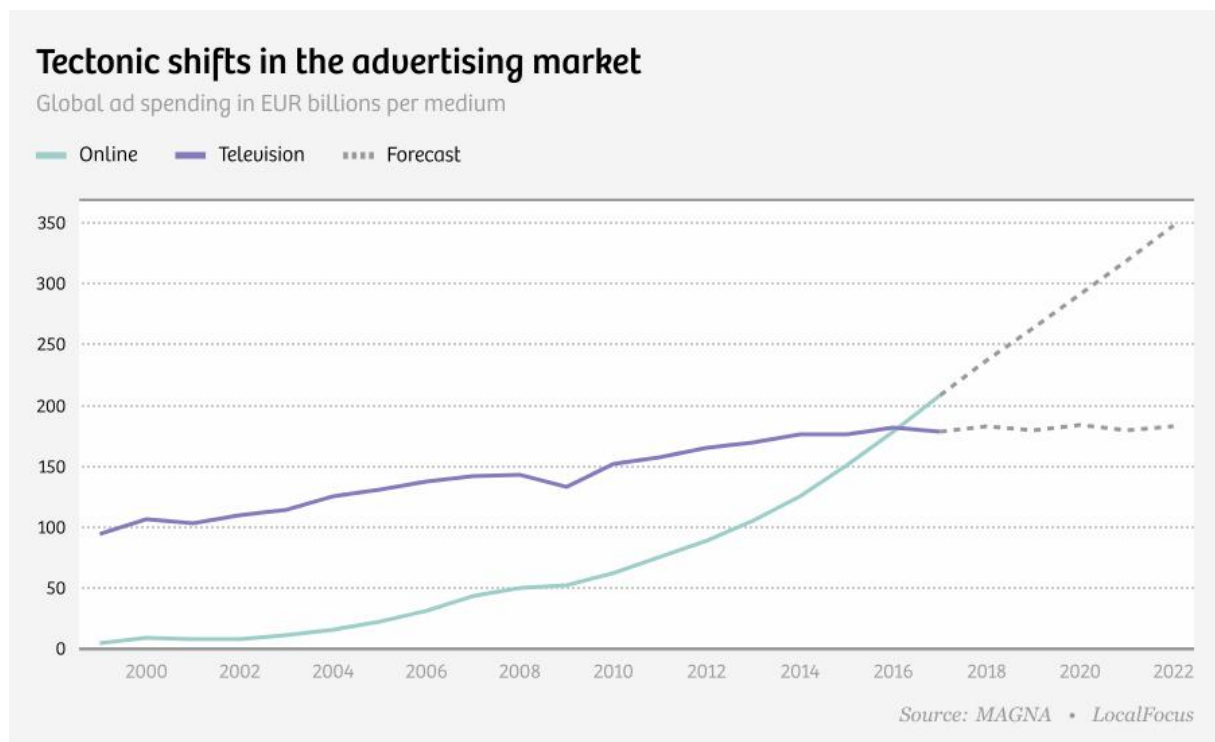
For more than a century, advertising was an art, not a science. Hard data didn't exist. An advertising guru of the Don Draper type proclaimed: "What you call love was invented by guys like me to sell nylons" – and advertisers could only hope it was true. You put your commercials on the air, you put your brand in the paper, and you started praying. Would anyone see the ad? Would anyone act on it? Nobody knew.

In the early 1990s, the internet sounded the death knell for that era of advertising. Today, we no longer live in the age of Mad Men, but of Math Men.

Looking for customers, clicks, conversions? Google and Facebook know where to find them. With unprecedented precision, these data giants will get the right message delivered to the right people at the right time. Unassuming internet users are lured into online shops, undecided voters are informed about the evils of US presidential candidate Elizabeth Warren, and cars zip by on the screens of potential buyers – a test drive is only a click away.

But is any of it real? What do we really know about the effectiveness of digital advertising? Are advertising platforms any good at manipulating us?

You'd be forgiven for thinking the answer to that last question is: yes, extremely good. After all, the market is huge. The amount of money spent on internet ads goes up each year. In 2018, more than \$273bn dollars was spent on digital ads globally, according to research firm eMarketer. Most of those ads were purchased from two companies: Google (\$116bn in 2018) and Facebook (\$54.5bn in 2018).



Newspapers are teeming with treatises about these tech giants' saturnine activities. An essay by best-selling author Yuval Noah Harari on "the end of free will" exemplifies the genre: according to the Israeli thinker, it's only a matter of time before big data systems "understand humans much better than we understand ourselves."

In a highly acclaimed new book, Harvard professor Shoshana Zuboff predicts a "seventh extinction wave", where human beings lose "the will to will". Cunning marketers can predict and manipulate our behaviour. Facebook knows your soul. Google is hacking your brain.

‘The best minds of our generation’

I too used to believe that these tech giants were all-knowing entities. But while writing this story, I have come to realise that this belief is as wrong as it is popular.

A former Facebook engineer once said (and he’s been quoted a thousand times over): “The best minds of my generation are thinking about how to make people click on ads.” I spoke to some of those best minds: economists employed and formerly employed by the most powerful companies in Silicon Valley: Yahoo!, Google, Microsoft, eBay, Facebook, Netflix, Pandora and Amazon.

This story is about organisations and why they are so hard to change. And it’s about us, and how easy we are to manipulate.

They weren’t always easy to get hold of. I’d send an email in late October, and perhaps they would have an hour in January, at some ungodly hour in Europe. And then there was the language barrier. These guys did not speak English, but fluent Economese. My Economese is not all bad, but two hours of Hausman tests, incremental bidding, and exogenous variation is too much, even for an enthusiast such as myself.

Still, mixed in with the jargon, there were enough anecdotes to make your head spin. They would utter one line in plain English that I’d end up mulling over the rest of the day.

The story that emerged from these conversations is about much more than just online advertising. It’s about a market of a quarter of a trillion dollars governed by irrationality. It’s about knowables, about how even the biggest data sets don’t always provide insight. It’s about organisations and why they are so hard to change. And it’s about us, and how easy we are to manipulate.

One thing was gradually becoming clear: these guys are fucking with the magic. And nobody knows it. Or as Garrett Johnson, who used to work for Yahoo!, told me: “I don’t have anybody pounding down my door telling me I’m fucking with their magic, because ... well, they don’t even know who I am.”

Steve Tadelis was the most accessible of the bunch. He instantly replied to my email: “I would be delighted to talk.”

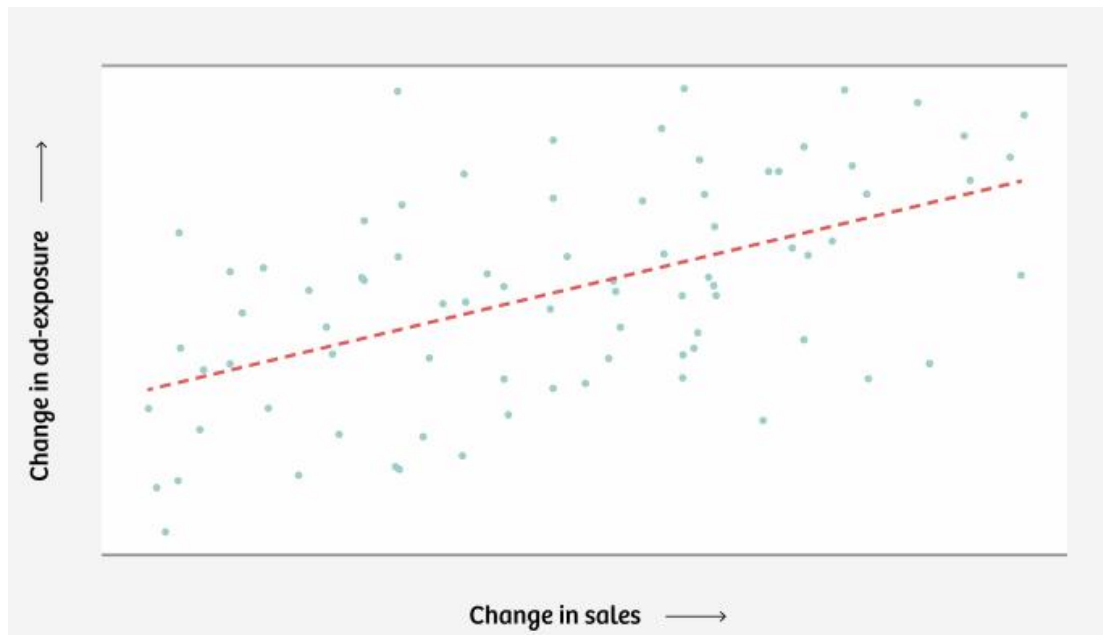
Tadelis told me about his work for eBay, a broad grin lining his face.

It all started with a surrealistic phone call to a data consultant. Tadelis was a professor of economics at the University of California, Berkeley when he went and spent a year at eBay in August 2011.

During one of his first conversations with eBay’s marketing team, they invited him to sit down with their consultants. The consultants could tell him how profitable each of eBay’s ad campaigns had been. And since Tadelis was an economist, maybe he’d like to quiz them about their methods.

“Proprietary transformation functions,” one of the consultants had said on the phone when Tadelis reached out. They used proprietary transformation functions, had 25 years of experience, and a long list of prominent clients.

When Tadelis pressed them he realised that “proprietary transformation functions” was only a clever disguise for your garden-variety statistics. You take the weekly expenditure on ads, combine it with the weekly sales, and voila! Fold the mixture into a scatter plot and see what happens. Easy as that!



"This is garbage," Tadelis thought.

Correlation, as any Statistics 101 class will inform you, is not causation. What do these impressive numbers mean if the people who see your ad are the exact same people who were going to use eBay anyway? eBay is no small fry. Surely lots of people looking for shoes end up on the online auction site all by themselves, whether they see an ad or not?

Picture this. Luigi's Pizzeria hires three teenagers to hand out coupons to passersby. After a few weeks of flyering, one of the three turns out to be a marketing genius. Customers keep showing up with coupons distributed by this particular kid. The other two can't make any sense of it: how does he do it? When they ask him, he explains: "I stand in the waiting area of the pizzeria."

It's plain to see that junior's no marketing whiz. Pizzerias do not attract more customers by giving coupons to people already planning to order a quattro stagioni five minutes from now.

I can get my head around a cynical advertising world, but a naive one?

Economists refer to this as a "selection effect." It is crucial for advertisers to distinguish such a selection effect (people see your ad, but were already going to click, buy, register, or download) from the advertising effect (people see your ad, and that's why they start clicking, buying, registering, downloading). Tadelis asked how exactly the consultants made this distinction.

"We use Lagrange multipliers," one of them said. And for a second, Tadelis was astounded. What? Lagrange multipliers? But Lagrange multipliers don't have anything to do with ... "Then it hit me," Tadelis recalled. "This guy is trying to out-jargon me!"

"I resisted the temptation to say: 'I'm sorry, you're fucked, I actually teach this stuff.'" Instead, Tadelis decided to continue the conversation in Economese.

"Lagrange multipliers, that's fascinating," he replied. "So now I know you have a constrained optimisation model, and as we all know the Lagrange multipliers are the shadow values of the constraints in the objective function. We all know this, right?"

The line went silent.

"So what is your objective function, and what are your constraints?"

...

"Steve, are you on a cell phone? Because you're breaking up and I can't hear you."

A not-so-brilliant advertising campaign

Two weeks later, Tadelis met the marketing consultants in the flesh. The advisers had put together a slick presentation demonstrating how eBay was raking in piles of cash with its brilliant ad campaigns. Tadelis recalled: "I looked around the room, and all I saw were people nodding their heads."

Brand keyword advertising, the presentation informed him, was eBay's most successful advertising method. Somebody googles "eBay" and for a fee, Google places a link to eBay at the top of the search results. Lots of people, apparently, click on this paid link. So many people, according to the consultants, that the auction website earns at least \$12.28 for every dollar it spends on brand keyword advertising – a hefty profit!

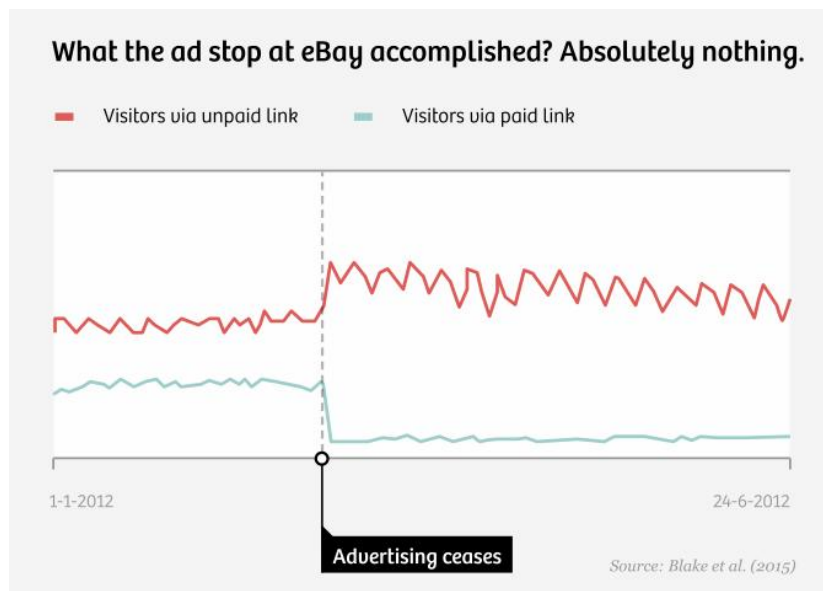
Tadelis didn't buy it. "I thought it was fantastic, and I don't mean extraordinarily good or attractive. I mean imaginative, fanciful, remote from reality." His rationale? People really do click on the paid-link to eBay.com an awful lot. But if that link weren't there, presumably they would click on the link just below it: the free link to eBay.com. The data consultants were basing their profit calculations on clicks they would be getting anyway.

Tadelis suggested an experiment: stop advertising for a while, and let's see whether brand keyword advertising really works. The consultants grumbled.

When, a few weeks later, Tadelis contacted the consultants about a follow-up meeting, he was told the follow-up had come and gone. He hadn't been invited.

A few months after the awkward presentation, though, Tadelis got the chance to conduct his experiment after all. There was a clash going on between the marketing department at eBay and the MSN network (Bing and Yahoo!). Ebay wanted to negotiate lower prices, and to get leverage decided to stop ads for the keyword 'eBay'.

Tadelis got right down to business. Together with his team, he carefully analysed the effects of the ad stop. Three months later, the results were clear: all the traffic that had previously come from paid links was now coming in through ordinary links. Tadelis had been right all along. Annually, eBay was burning a good \$20m on ads targeting the keyword 'eBay'.



When Tadelis presented his findings to the company, eBay's financial department finally woke up.

The economist was given a free hand: he was permitted to halt all of eBay's ads on Google for three months throughout a third of the United States. Not just those for the brand's own name, but also those targeted to match simple keywords like "shoes", "shirts" and "glassware".

The marketing department anticipated a disaster: sales, they thought, were certain to drop at least 5%.

Week 1: All quiet.

Week 2: Still quiet.

Week 3: Zip, zero, zilch.

The experiment continued for another eight weeks. What was the effect of pulling the ads? Almost none. For every dollar eBay spent on search advertising, they lost roughly 63 cents, according to Tadelis's calculations.

The experiment ended up showing that, for years, eBay had been spending millions of dollars on fruitless online advertising excess, and that the joke had been entirely on the company.

To the marketing department everything had been going brilliantly. The high-paid consultants had believed that the campaigns that incurred the biggest losses were the most profitable: they saw brand keyword advertising not as a \$20m expense, but a \$245.6m return.

For Tadelis, it was an eye-opener. "I kind of had the belief that most economists have: businesses are advertising, so it must be good. Because otherwise why would they do it?" He added: "But after my experience at eBay that's all out of the window."

I felt just as surprised as Tadelis had been. A cynical advertising world is something I can get my head around, but a naive one? But the more I talked to these economists, the more I realised that eBay was not alone in making this mistake.

I spoke to Randall Lewis, who used to work for Yahoo!, Google and Netflix, and is currently head of research for the ad platform Nanigans. For the type, Lewis was something of an unparalleled virtuoso. (Johnson, who worked with Lewis at Google, later confessed to me: "One of the strengths that I bring to the table is that I'm good at translating Randall's genius into something the rest of us can understand.")

The benchmarks that advertising companies use are fundamentally misleading.

Lewis had been part of a small group of economists at Yahoo! who had done lots of experiments with advertising. This meant that he had more than eight years of experience disappointing advertisers. "It's always awkward," he told me. "They think everything is rosy. But when you're running these experiments ... well, when things look too good to be true, they usually are."

Lewis explained that the whole thing is pear-shaped because for the most part, the industry is in the sway of the same wrong-headed statistical line of thought. The online marketing world has the same strategy as Luigi's Pizzeria and the flyer-handling teens. That's where eBay had gone wrong with search advertising. But the same thing happens with banner advertising, Instagram videos and Facebook ads.

The benchmarks that advertising companies use – intended to measure the number of clicks, sales and downloads that occur after an ad is viewed – are fundamentally misleading. None of these benchmarks distinguish between the selection effect (clicks, purchases and downloads that are happening anyway) and the advertising effect (clicks, purchases and downloads that would not have happened without ads).

It gets worse: the brightest minds of this generation are creating algorithms which only increase the effects of selection.

Consider the following: if Amazon buys clicks from Facebook and Google, the advertising platforms' algorithms will seek out Amazon clickers. And who is most likely to click on Amazon? Presumably Amazon's regular customers. In that case the algorithms are generating clicks, but not necessarily extra clicks.

Advertising platforms are not the only ones susceptible to this flawed way of thinking. Advertisers make the same error. They're targeting personalised ads at an audience that is already very likely to buy their product. You watch a Renault commercial, and then your screen is taken over by Twingos. You put a dress in an online shopping basket, then it's stalking you across the internet. You liked World of Warcraft, and now your timeline is full of Larp events ("OrcFest 2019: bring your battle axe!").

But who knows, maybe you really would have bought that dress anyway, maybe you've had your eye on a Twingo for months, and perhaps you've just ordered a battle axe.

The majority of advertising companies feed their complex algorithms silos full of data even though the practice never delivers the desired result.

I had never really thought about this. Algorithmic targeting may be technologically ingenious, but if you're targeting the wrong thing then it's of no use to advertisers. Most advertising platforms can't tell clients whether their algorithms are just putting fully-automated teenagers in the waiting area (increasing the selection effect) or whether they're bringing in people who wouldn't have come in otherwise (increasing the advertising effect).

"We are setting ourselves up for failure," Lewis explained, "because we are optimising for the wrong thing."

We currently assume that advertising companies always benefit from more data. And certainly, live-gaming is more likely to appeal to gamers. But the majority of advertising companies feed their complex algorithms silos full of data even though the practice never delivers the desired result. In the worst case, all that invasion of privacy can even lead to targeting the wrong group of people.

This insight is conspicuously absent from the debate about online privacy. At the moment, we don't even know whether all this privacy violation works as advertised.

Run an experiment!

Luckily there is a way to measure the unadulterated effect of ads: do an experiment. Divide the target group into two random cohorts in advance: one group sees the ad, the other does not. Designing the experiment thus excludes the effects of selection.

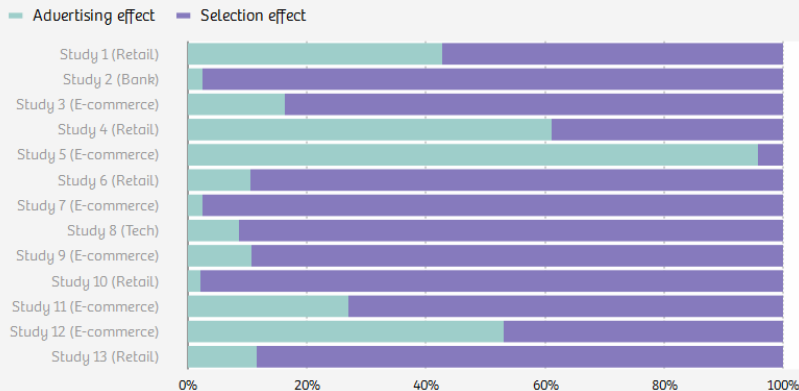
Economists at Facebook conducted 15 experiments that showed the enormous impact of selection effects. A large retailer launched a Facebook campaign. Initially it was assumed that the retailer's ad would only have to be shown 1,490 times before one person actually bought something. But the experiment revealed that many of those people would have shopped there anyway; only one in 14,300 found the webshop because of the ad. In other words, the selection effects were almost 10 times stronger than the advertising effect alone!

And this was no exception. Selection effects substantially outweighed advertising effects in most of these Facebook experiments. At its strongest, the selection bias was even 50 (!) times more influential¹.

¹ Now any good data analyst would say: that's all fine and well, but there's no need to experiment when you can just correct. Let's return to the flyer guys in the pizzeria waiting area for a minute. It's possible to create a model in which the distance from the pizzeria is also accounted for. This way, flyers handed out inside the pizzeria are rewarded a lower 'score' than flyers given to people outside the pizzeria. In short, the influence of all manner of things that play into purchasing decisions can be corrected for, leaving you with only that of the ad itself. That means you have to know all the factors that could possibly shape the purchasing decision (such as age, gender, post code, preferences, online activity, and so on and so forth) and you need good data on these factors. What Zettelmeyer et al. demonstrated at Facebook was that even with the very best data – the researchers knew basically everything that anyone's ever done on Facebook – he could not correct for the selection bias. "The methods we study still yielded biased estimates of the causal effects of advertising in a majority of cases." A delicate economic palate is required to appreciate this point, but it's worth savouring. Observational methods correct for any non-random factors, which leaves just a mote of random ones to work with. Machine learning algorithms are intended, though, to exclude the random: who ultimately sees an ad is increasingly determined. So, as the algorithms become more effective, it becomes harder to assess their effectiveness!

Facebook ads are far less effective than we thought

When people see an ad, what percentage of the total effect on sales can be attributed to the selection effect, and how much is due to the advertising effect?



Source: [Gordon et al. \(2018\)](#) • LocalFocus

In seven of the 15 Facebook experiments, advertising effects without selection effects were so small as to be statistically indistinguishable from zero.

Now we arrive at perhaps the most fundamental question: what, in the end, is there really to know in advertising? Can advertisers ever know exactly what their ad brings in?


Google CEO Eric Schmidt told his TV colleague Mel Karmazin that when it comes to online advertising, that question was easy to answer. Lewis went on to work for Schmidt, but research he conducted for Yahoo! in 2011 puts the lie to that claim. The title of his paper: On the near impossibility of measuring the returns to advertising.

Disappointment had been the study's driving force. At Yahoo!, Lewis had run 25 gigantic ad-experiments. And still, he was left with a lot of uncertainty about the actual effects of advertising.

"People thought that after a one-million-person experiment, we could walk away, and know exactly how advertising works," Johnson recalled. He added, "if you've got a million you should be able to count angels dancing on pins."

So what went wrong? If you want to measure something small, you have to go big. Let's say I want to know how many people have the rare disease cystic fibrosis. Cystic fibrosis affects one in 3,400 people (0.03%). But let's say I don't know that.

So I open the phone book and I call 10,000 people. Plus another 10,000. And another 10,000. Then another 10,000.



So you see, the results of my poll are all over the place. 10,000 is simply too small a sample to get reliable estimates. We'd better call a million people. And another million. And another million. Now we're getting somewhere.

Imagine, then, that I had wanted to know how many people had contracted the flu last year (one in 20). Ten thousand calls would have been enough to get reliable estimates. More people get the flu, so a flu study can have smaller test groups.

The point is, advertising is like cystic fibrosis, not the flu. And even that's extremely unfair to cystic fibrosis, since people buying things because they saw an ad is even rarer than cystic fibrosis.

- **Who has the flu? 10,000 calls**



Johnson sighs: "It's very hard to change behaviour by showing people pictures and movies that they don't want to look at."

To illustrate, consider Steve Tadelis's eBay research. Ebay lost 63 cents on every dollar they put into Google search advertising, but that's actually an imprecise estimate. If the experiment were to be replicated infinitely (and another ad stop, and another ad stop, and another ad stop...), in 95% of all ad stops the true loss would fall in the range of negative \$1.24 and negative \$0.03. This is what statisticians call the confidence interval. In advertising research, the confidence interval tends to be huge.

Ebay's performance was so shoddy that the only logical conclusion would have been: stop buying search ads! But if eBay's marketing had been just a tiny bit more effective – say they only lost 10 cents on every dollar they invested – then their experiment would have shown that the marketing department had delivered something between a 70 cent loss and a 50 cent profit.

Advertising does far less than most advertisers believe

What good is information like this? Such experiments tend to have an either-or conclusion: the campaign was either profitable or it wasn't. This may give you a sense of direction, but it cannot provide certainty. "Simply rejecting that a campaign was a total waste of money is not an ambitious goal," Randall Lewis wrote in his study. Still, in practice, that proved 'nearly impossible'.

Advertising rationally, the way it's described in economic textbooks, is unattainable. Then how do advertisers know what they ought to pay for ads?

"Yeah, basically they don't know," Lewis said in one of those throw-away clauses that kept running through my head for days after.

Keep that in mind the next time you read one of those calamity stories about Google, Facebook or Cambridge Analytica. If people were easier to manipulate with images and videos they don't really want to see, economists would have a much easier task. Realistically, advertising does something, but only a small something – and at any rate it does far less than most advertisers believe.

"What frustrates me is there's a bit of magical thinking here," Johnson says. "As if Cambridge Analytica has hacked our brains so that we're going to be like lemmings and jump off cliffs. As if we are powerless."

So we arrive at our final question: who wants to know the truth?

It's a question that has long fascinated the economist Justin Rao (who's worked for Yahoo!, Microsoft and others). Before he worked on advertising, he did field research with a cult that predicted the end of days on 21 May 2011. Rao awarded prizes to cult members. Those willing to accept their prize

after Judgement Day – when the world would be annihilated and the faithful would ascend to heaven – were promised more money. Their belief in the apocalypse proved uncompromising. Even an extra 500 dollars couldn't seduce the cultists.

"Beliefs formed on insufficient evidence seem tough to move," Rao wrote.

When Rao joined Microsoft, the eBay studies by Steve Tadelis had just been published and were all over the media: the Harvard Business Review, The Economist, The Atlantic and the BBC all covered the story. Marketing blogs couldn't stop talking about it.

According to Rao, "Probably even Steve's mother e-mailed him."

But did it matter? At Microsoft, Rao had a search engine at his disposal: Bing. Following the news about the millions of dollars eBay had wasted, brand keyword advertising only declined by 10%. The vast majority of businesses proved hell-bent on throwing away their money.

The fact that the eBay news did not even encourage advertisers to experiment more was perhaps the most striking.

Rao did observe the occasional ad stop at Bing. Rao was able to use ad stops like these, just as Tadelis had at eBay, to assess the effects on search traffic.

When these experiments showed that ads were utterly pointless, advertisers were not bothered in the slightest. They charged gaily ahead, buying ad after ad. Even when they knew, or could have known, that their ad campaigns were not very profitable, it had no impact on how they behaved.

"Beliefs formed on insufficient evidence seem tough to move."

#FitGirl analytics

Steve Tadelis saw this first-hand too. The financial director of eBay asked Tadelis to look into the second item on the list of so-called success campaigns: affiliate marketing. An example of this type of advertising could be eBay paying some influencer #fitgirl to embed a link to a particular brand of yoga pants in an Instagram post.

The affiliate marketing boss was okay with Tadelis experimenting, but he did issue a caveat. "Let me tell you something Steve," he had said. "If we run this experiment, and the results look like what you showed us with search advertising, I'm not going to believe you."

"It was clear to me that he meant it," Tadelis recalled. "So I told him: 'Well, if this is about religion, I can't help you. I have nothing against religion, I just don't think it has a place in marketing analytics.'"

It might sound crazy, but companies are not equipped to assess whether their ad spending actually makes money. It is in the best interest of a firm like eBay to know whether its campaigns are profitable, but not so for eBay's marketing department.

Its own interest is in securing the largest possible budget, which is much easier if you can demonstrate that what you do actually works. Within the marketing department, TV, print and digital compete with each other to show who's more important, a dynamic that hardly promotes honest reporting.

The fact that management often has no idea how to interpret the numbers is not helpful either. The highest numbers win.

Randall Lewis told me about a meeting with the man responsible for evaluating Yahoo's marketing strategy. The man had apparently done everything Lewis had advised against – and worse. He graciously admitted that he either added or omitted data to his model if it led to the 'wrong' results. Lewis: "I was like: oh man. All of that is bad scientific practice, but it's actually great job preservation practice."

"Bad methodology makes everyone happy," said David Reiley, who used to head Yahoo's economics team and is now working for streaming service Pandora. "It will make the publisher happy. It will make the person who bought the media happy. It will make the boss of the person who bought the media happy. It will make the ad agency happy. Everybody can brag that they had a very successful campaign."

Marketers are often most successful at marketing their own marketing.

Is online advertising working? We simply don't know

Perhaps what's driving this phenomenon is something much more profound. Something that applies not just to advertising. "There is a fear that saying 'I don't know' amounts to an admission of incompetence," Tadelis observed. "But ignorance is not incompetence, curiosity is not incompetence."

We want certainty. We used to find it in the Don Drapers of the world, the ones with the best one-liners up their sleeves. Today we look for certainty from data analysts who are supposed to just show us the numbers.

Lewis admitted that it's not all bad. Decisions have to be made, somebody has to lay out a strategy, doubt must stop at some point. For that reason, companies hire overconfident people who act like they know what they cannot possibly know.

Lewis could never do the sort of work they do. "I would feel like it's a random coin toss for most decisions," he said. But somebody has to toss the coin. And a company full of Randalls only leads to analysis paralysis. Nothing happens.

Randall Lewis had left Google and was working for Netflix when he attended the Datalead Conference in Paris in November 2015.

His time at Yahoo! and Google had taught him how difficult it is to advertise better. But Lewis wasn't out to blind with science, he didn't want to burn it all down and turn his back. He wanted to make the near impossible just a little bit more possible. And let's be honest: advertising a bit better is actually quite a lot compared to stumbling about in the dark. It can prevent blunders of the eBay sort.

"Marketeers actually believe that their marketing works, even if it doesn't."

Lewis had come to Paris to present one of his improvements. At Google, he had built a platform that gives advertisers a cheap and simple way to experiment with banner ads. "I think this is a revolution in advertising," he said proudly. Buyers could finally optimise for the right thing.

About a quarter of the way through his presentation, an audience member stood up and asked: do advertising companies actually want to know this? Aren't they primarily interested in research that reassures?

"That's actually endemic to the entire industry," Lewis replied. He started on one of his brilliantly inaccessible Lewisian responses. "The moral hazard problem is a series of cognitive dissonance biases ..."

Halfway through his impenetrable answer, another audience member interruption. This time it came from Steve Tadelis. "What Randall is trying to say," the former eBay economist interjected, "is that marketeers actually believe that their marketing works, even if it doesn't. Just like we believe our research is important, even if it isn't."

Lewis laughed. "Thank you, Steve."