

How the net traps us all in our own little bubbles

Adapted from The Filter Bubble by Eli Pariser.

An invisible revolution has taken place in the way we use the net, but the increasing personalisation of information by search engines such as Google threatens to limit our access to information and enclose us in a self-reinforcing world view, writes Eli Pariser in an extract from The Filter Bubble

The Observer, Sunday 12 June 2011

A slide from Eli Pariser's TED Talk presentation which discusses how major internet players are tailoring information to individuals.

Few people noticed the post that appeared on Google's corporate blog on 4 December 2009. It didn't beg attention – no sweeping pronouncements, no Silicon Valley hype, just a few paragraphs sandwiched between a round-up of top search terms and an update on Google's finance software. Not everyone missed it. Search-engine blogger Danny Sullivan pores over the items on Google's blog, looking for clues about where the monolith is headed next, and to him, the post was a big deal. In fact, he wrote later that day, it was "the biggest change that has ever happened in search engines". For Danny, the headline said it all: "Personalised search for everyone".

Starting that morning, Google would use 57 signals – everything from where you were logging in from to what browser you were using to what you had searched for before – to make guesses about who you were and what kinds of sites you'd like. Even if you were logged out, it would customise its results, showing you the pages it predicted you were most likely to click on.

Most of us assume that when we google a term, we all see the same results – the ones that the company's famous Page Rank algorithm suggests are the most authoritative based on other pages' links. But since December 2009, this is no longer true. Now you get the result that Google's algorithm suggests is best for you in particular – and someone else may see something entirely different. In other words, there is no standard Google any more.

It's not hard to see this difference in action. In the spring of 2010, while the remains of the Deepwater Horizon oil rig were spewing oil into the Gulf of Mexico, I asked two friends to search for the term "BP". They're prettysimilar – educated white left-leaning women who live in the north-east. But the results they saw were quite different. One saw investment information about BP. The other saw news. For one, the first page of results contained links about the oil spill; for the other, there was nothing about it except for a promotional ad from BP. Even the number of results returned differed – 180 million for one friend and 139 million for the other. If the results were that different for these two progressive east-coast women, imagine how different they would be for my friends and, say, an elderly Republican in Texas (or, for that matter, a businessman in Japan).

With Google personalised for everyone, the query "stem cells" might produce diametrically opposed results for scientists who support stem-cell research and activists who oppose it. "Proof of climate change" might turn up different results for an environmental activist and an oil-company executive. A huge majority of us assume search engines are unbiased. But that may be just because they're increasingly biased to share our own views. More and more, your computer monitor is a kind of one-way mirror, reflecting your own interests while algorithmic observers watch what you click. Google's announcement marked the turning point of an important but nearly invisible revolution in how we consume information. You could say that on 4 December 2009 the era of personalisation began. With little notice or fanfare, the digital world is fundamentally changing.

What was once an anonymous medium where anyone could be anyone – where, in the words of the famous New Yorker cartoon, nobody knows you're a dog – is now a tool for soliciting and analysing our personal data.

According to one Wall Street Journal study, the top 50 internet sites, from CNN to Yahoo to MSN, install an average of 64 data-laden cookies and personal tracking beacons each. Search for a word like "depression" on Dictionary.com, and the site installs up to 223 tracking cookies and beacons on your computer so that other websites can target you with antidepressants. Open a page listing signs that your spouse may be cheating, and prepare to be haunted with DNA paternity-test ads. The new internet doesn't just know you're a dog: it knows your breed and wants to sell you a bowl of premium dog food. The race to know as much as possible about you has become the central battle of the era for internet giants like Google, Facebook, Apple and Microsoft. As Chris Palmer of the Electronic Frontier Foundation explained to me: "You're getting a free service, and the cost is information about you. And Google and Facebook translate that pretty directly into money." While Gmail and Facebook may be helpful, free tools, they are also extremely effective and voracious extraction engines into which we pour the most intimate details of our lives. Your smooth new iPhone knows exactly where you go, whom you call, what you read; with its built-in microphone, gyroscope and GPS, it can tell whether you're walking or in a car or at a party.

While Google has (so far) promised to keep your personal data to itself, other popular websites and apps make no

such guarantees. Behind the pages you visit, a massive new market for information about what you do online is growing, driven by low-profile but highly profitable personal data companies like BlueKai and Acxiom. Acxiom alone has accumulated an average of 1,500 pieces of data on each person on its database – which includes 96% of Americans – along with data about everything from their credit scores to whether they've bought medication for incontinence. And any website – not just the Googles and Facebooks of the world – can now participate in the fun. In the view of the "behaviour market" vendors, every "click signal" you create is a commodity, and every move of your mouse can be auctioned off within microseconds to the highest commercial bidder.

As a business strategy, the internet giants' formula is simple: the more personally relevant their information offerings are, the more ads they can sell, and the more likely you are to buy the products they're offering. And the formula works. Amazon sells billions of dollars in merchandise by predicting what each customer is interested in and putting it in the front of the virtual store. Up to 60% of US film download and DVD-by-mail site Netflix's rentals come from the guesses it can make about each customer's preferences.

In the next three to five years, says Facebook chief operating officer Sheryl Sandberg, the idea of a website that isn't customised to a particular user will seem quaint. Yahoo vice president Tapan Bhat agrees: "The future of the web is about personalisation... now the web is about 'me'. It's about weaving the web together in a way that is smart and personalised for the user." Google CEO Eric Schmidt enthuses that the "product I've always wanted to build" is Google code that will "guess what I'm trying to type". Google Instant, which guesses what you're searching for as you type, and was rolled out in the autumn of 2010, is just the start – Schmidt believes that what customers want is for Google to "tell them what they should be doing next".

It would be one thing if all this customisation was just about targeted advertising. But personalisation isn't just shaping what we buy. For a quickly rising percentage of us, personalised news feeds like Facebook are becoming a primary news source – 36% of Americans aged under 30 get their news through social-networking sites. And Facebook's popularity is skyrocketing worldwide, with nearly a million more people joining each day. As founder Mark Zuckerberg likes to brag, Facebook may be the biggest source of news in the world (at least for some definitions of "news"). And personalisation is shaping how information flows far beyond Facebook, as websites from Yahoo News to the New York Times-funded startup News.me cater their headlines to our particular interests and desires. It's influencing what videos we watch on YouTube and what blog posts we see. It's affecting whose emails we get, which potential mates we run into on OkCupid, and which restaurants are recommended to us on Yelp – which means that personalisation could easily have a hand not only in who goes on a date with whom but in where they go and what they talk about. The algorithms that orchestrate our ads are starting to orchestrate our lives.

The basic code at the heart of the new internet is pretty simple. The new generation of internet filters looks at the things you seem to like – the actual things you've done, or the things people like you like – and tries to extrapolate. They are prediction engines, constantly creating and refining a theory of who you are and what you'll do and want next. Together, these engines create a unique universe of information for each of us – what I've come to call a filter bubble – which fundamentally alters the way we encounter ideas and information. Of course, to some extent we've always consumed media that appealed to our interests and avocations and ignored much of the rest. But the filter bubble introduces three dynamics we've never dealt with before.

First, you're alone in it. A cable channel that caters to a narrow interest (say, golf) has other viewers with whom you share a frame of reference. But you're the only person in your bubble. In an age when shared information is the bedrock of shared experience, the filter bubble is a centrifugal force, pulling us apart.

Second, the filter bubble is invisible. Most viewers of conservative or liberal news sources know that they're going to a station curated to serve a particular political viewpoint. But Google's agenda is opaque. Google doesn't tell you who it thinks you are or why it's showing you the results you're seeing. You don't know if its assumptions about you are right or wrong – and you might not even know it's making assumptions about you in the first place. My friend who got more investment-oriented information about BP still has no idea why that was the case – she's not a stockbroker. Because you haven't chosen the criteria by which sites filter information in and out, it's easy to imagine that the information that comes through a filter bubble is unbiased, objective, true. But it's not. In fact, from within the bubble, it's nearly impossible to see how biased it is.

Finally, you don't choose to enter the bubble. When you turn on Fox News or read The New Statesman, you're making a decision about what kind of filter to use to make sense of the world. It's an active process, and like putting on a pair of tinted glasses, you can guess how the editors' leanings shape your perception. You don't make the same kind of choice with personalised filters. They come to you – and because they drive up profits for the websites that use them, they'll become harder and harder to avoid. Personalisation is based on a bargain. In exchange for the service of filtering, you hand large companies an enormous amount of data about your daily life – much of which you might not trust friends with. These companies are getting better at drawing on this data to

make decisions every day. But the trust we place in them to handle it with care is not always warranted, and when decisions are made on the basis of this data that affect you negatively, they're usually not revealed.

Ultimately, the filter bubble can affect your ability to choose how you want to live. To be the author of your life, professor Yochai Benkler argues, you have to be aware of a diverse array of options and lifestyles. When you enter a filter bubble, you're letting the companies that construct it choose which options you're aware of. You may think you're the captain of your own destiny, but personalisation can lead you down a road to a kind of informational determinism in which what you've clicked on in the past determines what you see next – a web history you're doomed to repeat. You can get stuck in a static, ever-narrowing version of yourself – an endless you-loop.

And there are broader consequences. In *Bowling Alone*, his book on the decline of civic life in America, Robert Putnam looked at the problem of the major decrease in "social capital" – the bonds of trust and allegiance that encourage people to do each other favours, work together to solve common problems, and collaborate. Putnam identified two kinds of social capital: there's the in-group-oriented "bonding" capital created when you attend a meeting of your college alumni, and then there's "bridging" capital, which is created at an event like a town meeting when people from lots of different backgrounds come together to meet each other. Bridging capital is potent: build more of it, and you're more likely to be able to find that next job or an investor for your small business, because it allows you to tap into lots of different networks for help.

Everybody expected the internet to be a huge source of bridging capital. Writing at the height of the dotcom bubble, Tom Friedman declared that the internet would "make us all next-door neighbours". In fact, this idea was the core of his thesis in *The Lexus and the Olive Tree*: "The internet is going to be like a huge vice that takes the globalisation system ... and keeps tightening and tightening that system around everyone, in ways that will only make the world smaller and smaller and faster and faster with each passing day." Friedman seemed to have in mind a kind of global village in which kids in Africa and executives in New York would build a community together. But that's not what's happening: our virtual neighbours look more and more like our real-world neighbours, and our real-world neighbours look more and more like us. We're getting a lot of bonding but very little bridging. And this is important because it's bridging that creates our sense of the "public" – the space where we address the problems that transcend our narrow self-interests.

We are predisposed to respond to a pretty narrow set of stimuli – if a piece of news is about sex, power, gossip, violence, celebrity or humour, we are likely to read it first. This is the content that most easily makes it into the filter bubble. It's easy to push "Like" and increase the visibility of a friend's post about finishing a marathon or an instructional article about how to make onion soup. It's harder to push the "Like" button on an article titled "Darfur sees bloodiest month in two years". In a personalised world, important but complex or unpleasant issues – the rising prison population, for example, or homelessness – are less likely to come to our attention at all.

As a consumer, it's hard to argue with blotting out the irrelevant and unlikable. But what is good for consumers is not necessarily good for citizens. What I seem to like may not be what I actually want, let alone what I need to know to be an informed member of my community or country. "It's a civic virtue to be exposed to things that appear to be outside your interest," technology journalist Clive Thompson told me. Cultural critic Lee Siegel puts it a different way: "Customers are always right, but people aren't."

The era of personalisation is here, and it's upending many of our predictions about what the internet would do. The creators of the internet envisioned something bigger and more important than a global system for sharing pictures of pets. The manifesto that helped launch the Electronic Frontier Foundation in the early 1990s championed a "civilisation of Mind in cyberspace" – a kind of worldwide metabrain. But personalised filters sever the synapses in that brain. Without knowing it, we may be giving ourselves a kind of global lobotomy instead.

Early internet enthusiasts like web creator Tim Berners-Lee hoped it would be a new platform for tackling global problems. I believe it still can be, but first we need to pull back the curtain – to understand the forces that are taking the internet in its current direction. We need to lay bare the bugs in the code – and the coders – that brought personalisation to us. If "code is law", as Creative Commons founder Larry Lessig declared, it's important to understand what the new lawmakers are trying to do. We need to understand what the programmers at Google and Facebook believe in. We need to understand the economic and social forces that are driving personalisation, some of which are inevitable and some of which are not. And we need to understand what all this means for our politics, our culture and our future.

source <http://www.guardian.co.uk/technology/2011/jun/12/google-personalisation-internet-data-filtering>