

Big data is changing how consumer markets work

From the way prices are set to the way customers are acquired, consumer markets are entering a new and unfamiliar phase.

How should the state interact with consumer markets? When, how, and on what basis, should the government intervene? These are some of the oldest questions in modern economics and they are becoming relevant again. Big data and new analytical methods are changing important aspects of the market economy, from the way prices are set to the way customers are acquired and retained. In time, this will require new approaches to economic regulation and policymaking.

Big data and price discrimination

To build this argument, let's start by looking at prices. In the front-running sectors of the world's mature economies, big data now allows firms to personalise their prices, asking different customers to pay different amounts for the same product or service.

In its purest form, this marks a break from how most prices were set in the mass consumer markets of the 20th century. In these decades, we grew used to the idea that a price is a feature of a product itself, not of the person buying the product. From the perspective of consumers, prices appeared to reflect intuitive factors: the cost of a good's manufacturing process or the scarcity of its raw materials.

Price discrimination, the theory that prices can also reflect the amount that each individual consumer is willing to pay, was an idea discussed in the university classroom but rarely applied in the company boardroom.¹ It made sense in theory but in practice its effects were limited by two practical constraints.

¹ For a discussion of the theory of price discrimination see, for example, Armstrong, M. (2006) *Price Discrimination*, University College London. See also p.74 onwards, Philipps, R. L. (2005), *Pricing and Revenue Optimisation*, Stanford University Press.

First, consumer-facing firms simply did not *know* the ‘reserve price’ of each of their customers—the maximum price that each customer would be willing to pay. Of course, firms could safely assume that some of their customers would be willing to pay more than others, but they could not practically identify who was who.

Second, even *had* companies known the reserve price of each of their customers, they would have found it impractical and uneconomic to act on this information. Imagine changing your price tags, or reprinting your menus, for each new customer that walked in the door. Even if a company could have rigged up such a system, it would not have been economically viable.

As a rule, therefore, price discrimination has been confined to marginal instances: prices might vary across regions, for example (you might pay more for a tin of beans in Belgravia than in Bognor), or in sectors where haggling survives (a savvy negotiator might pay less for a second hand car), or when tied to targeted discounts (a Young Person’s Railcard gives access to cheaper train tickets). Aside from examples like this, however, in mainstream consumer markets the rule has been: one product, one price.

A new era of big data—and new pricing possibilities

In the last 10 years, big data has melted away both of these constraints. Tech-savvy companies can now use data analytics and experimental pricing methods to find the reserve price of specific subgroups of customers. Having done this, they can then act on that information, identifying when a specific customer visits their website and, by using cookies left on that customer’s laptop or phone, present them with a personalised price.

A caveat is important here: these are early days and opinions differ on the uptake of price discrimination so far. Studies suggest that purer forms of price discrimination are not yet widespread but that differentiated pricing is certainly happening, with online retailers found charging different customers different prices for the same

item.² The media has rooted out newsworthy examples, such as an insurance company charging more to customers with hotmail addresses.³ And there is evidence of search discrimination, in which customers are shown different results despite searching for the same term, to steer them to a more profitable item.⁴

It's also clear that price discrimination is becoming dramatically more profitable.⁵ Big datasets, like a customer's browsing history, are *far* better predictors of a person's price sensitivity than the simplistic demographic data on which companies used to rely. This is spawning an entire industry to advise companies on how to optimise their prices, covering everything from petrol to apps. In their own words, these consultancies are there to help companies 'align' their prices to a "customers' willingness to pay".⁶

For a simple illustration of the sophistication of modern pricing strategies, just look at how prices now move over time. The chart below shows the price of the board game Monopoly on Amazon over time. Prices no longer rise or fall in annual or quarterly increments. They move daily, even hourly, like stocks, as an algorithm optimises profit and customer insight over time.

The price of the board game Monopoly on Amazon

October 2013 - April 2018, £, per item

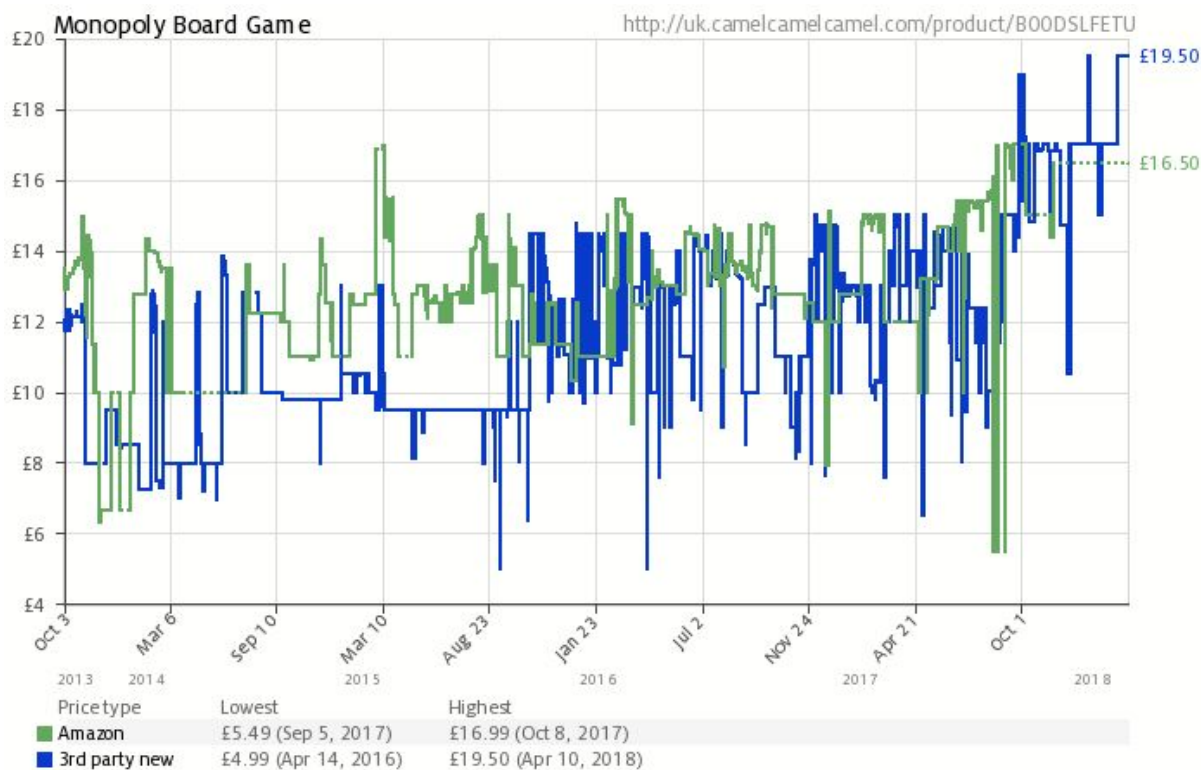
² Executive Office of the President of the United States, Council of Economic Advisors (2015), *Big data and differential pricing*, Washington, D.C. See also Hannak et al (2014), *Measuring Price Discrimination and Steering on E-commerce Websites*, Northeastern University. For a specific examples, see Valentino-DeVries, J., Singer-Vine, J. and Soltani, A. (24 December 2012), *Websites Vary Prices, Deals Based on Users' Information*, The Wall Street Journal

³ Hodge, K. and Leo, B. (23 January 2018), *Admiral bikes insurance costs for drivers using Hotmail email addresses*, The Sun (<https://bit.ly/2nBoEZ>, last accessed 10 April 2018)

⁴ Mattioli, D. (23 August 2012) *On Orbitz, Mac Users Steered to Pricier Hotels*, The Wall Street Journal. See also Hannak et al (ibid). See also Mikians, J. et al (2012), *Detecting price and search discrimination on the Internet*, Universitat Politècnica de Catalunya

⁵ Shiller, B. R. (2014) *First-Degree Price Discrimination Using Big Data*, Brandeis University

⁶ See, for example, Vendavo (<https://bit.ly/2Hdf1vW>, last accessed 10 April 2018). The company says its Price Optimisation Manager "enables you to grow revenue and profit by providing your sales team optimized target prices that are aligned with your customers' willingness to pay." See also A2ISystems and Sweet Pricing (<https://bit.ly/2qlEOaN> and <https://bit.ly/2Jz7GW7> respectively, accessed 10 April 2018)



Source: CamelCamelCamel (<https://bit.ly/1puxX8M>, last accessed 10 April 2018)

The winners and losers from price discrimination

There are indications, then, that in the tech-savvy sectors of our economy, companies are experimenting with setting prices in new and more sophisticated ways. As big data becomes ever more ubiquitous, how could this play out?

We should not worry about a dramatic scenario in which prices fragment and consumers routinely pay vastly different amounts for the same goods and services. This scenario—so-called ‘pure price discrimination’—is constrained by other factors. In particular, inter-firm competition limits the extent to which companies can raise prices for particular groups of customers. Some economists argue that this effect will dominate: some firms might experiment with tailored prices, but competition will pull markets back to one uniform low price.⁷

⁷ Choe, C. (2016) *Pricing with Cookies: Behavior-Based Price Discrimination and Spatial Competition*, Monash University

Even so, a middle scenario, short of perfect price discrimination but noticeably different to today, does seem plausible. For one thing, competition, particularly in regulated markets, is imperfect. For another thing, even successful competition does not insulate consumers entirely; recall that the very idea behind price discrimination is that companies will be better able to identify, and price differentiate, subgroups of price *insensitive* customers. This strategy works even if the *overall* market is competitive for the median customer.

These new capabilities, then, could herald a meaningful change to our lives as consumers in important markets. In the past, a small subset of price sensitive customers kept prices low for everyone else. Consumer-facing firms could not identify this subgroup and so, hidden among the overall customer population, they acted like undercover cops, keeping firms on their best behaviour.

Now, as big data becomes increasingly powerful, this dynamic will weaken, and prices will better reflect our *individual* willingness to pay.

What will happen as a result? All else equal, people who are insensitive to price changes—the rich, the time-poor, or those with little choice in the moment—will pay more. Meanwhile, people who are highly *sensitive* to price changes—the careful shoppers, the plan-aheaders, the savvy app-users—will pay less.

More concerning, prices will also better reflect the *cost* of serving individual customers. It will be easier for firms to identify, and charge more to, high cost customers—people who might not pay back their debts, or who have costlier needs. This development bites hardest in the insurance sector because it weakens the market's function as a risk-pooling mechanism. It also means prices could vary by protected characteristics: age, gender, or even ethnicity. And that explains growing concerns that discrimination could re-enter consumer markets via the backdoor of the algorithm.⁸

⁸ See, for example, BBC News (23 January 2018) *Admiral and M&S insurance firms deny 'racism' claims by The Sun* (<https://bbc.in/2IGPdWk>, last accessed 10 April 2018)

Big data *also* helps companies to shape consumer behaviour

These are interesting issues in their own right. But to appreciate the full implications of big data for consumers markets, we must turn to a second trend: the growing power of behavioural insights.

If big data helps companies to optimise their prices, it *also* helps them to optimise the design of products and services, maximising the chance that consumers will part with their money.

The idea of nudging consumers is nothing new. Retailers are old-masters of the cleverly-structured decision. Think, for example, of the sweets and gum by the till, to encourage impulse purchases, or the 3-for-2 offer that means you leave the store with three items when you only went in for one.

What *is* new is the unprecedented sophistication that big data allows. Large, consumer-facing companies like Amazon are now data companies as much as they are retailers or FMCG-manufacturers. These companies are running, in effect, massive, live behavioural experiments, in which they can tweak products, services, prices, or terms and conditions, and monitor the effect on consumer decisions.

This growing capability helps to explain the feeling that our lives as consumers are now strongly flavoured by frustration and regret. Consider, for example, that free trial you signed up for in a matter of seconds, but that has now become a monthly subscription that is a hassle to leave. Or consider the slow and clumsy switching process that leaves you sticking, grudgingly, to your energy deal. Or the insurance policy that you auto-renewed by mistake, paying a hugely increased premium.

These aspects of profit-maximising customer service are no less carefully designed than the more positive parts of a customer experience. And they now flavour our lives as consumers just as strongly as those sweet moments, thankfully also still common, of charming customer service or flawless product design.

Price discrimination and behavioural nudges: A dangerous combination

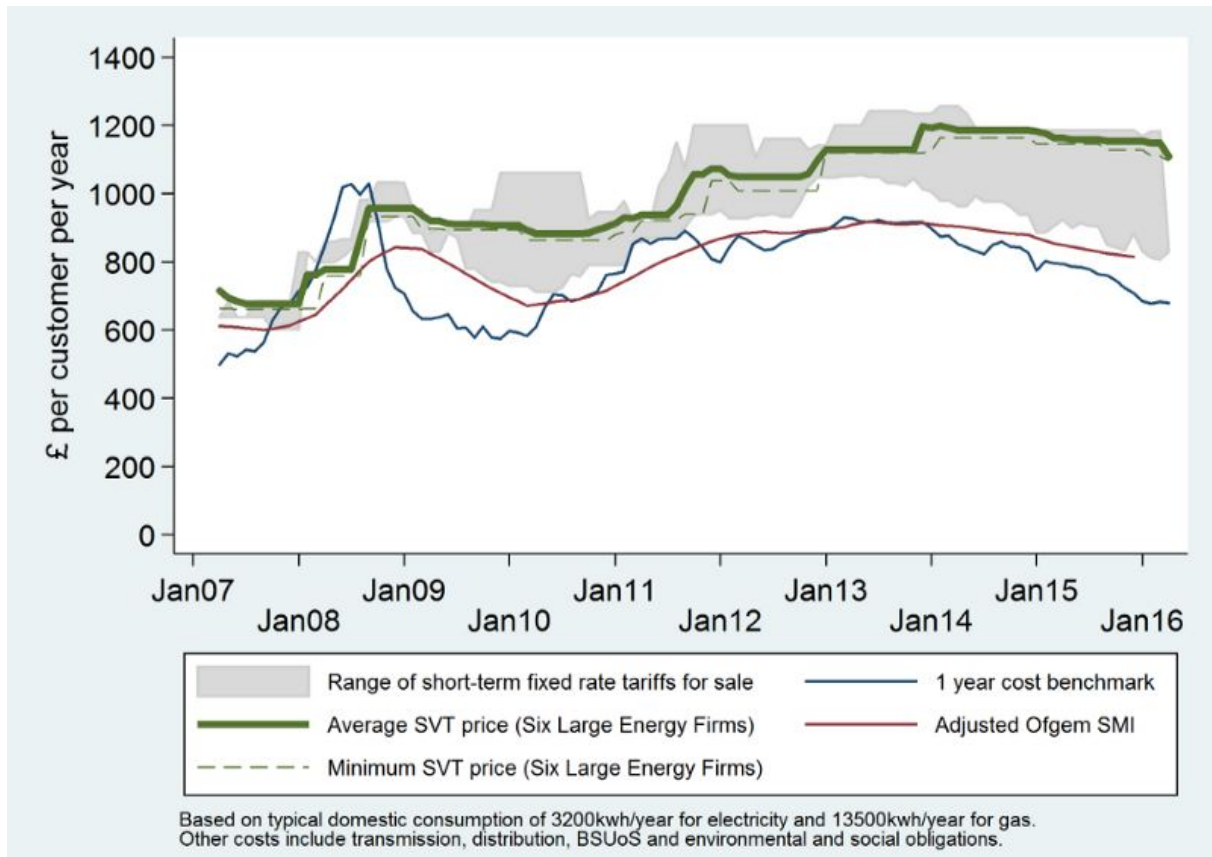
Taken individually, these two growing capabilities—the optimisation of prices and the optimisation of product- or service-design—are powerful. When they're combined, however, they become both more significant and more pernicious.

A good example is the penalty consumers now pay for inertia. In sectors from energy to telecoms to banking, companies now pursue a harmful business model that could be called 'bait and squeeze'. This means they compete fiercely with an attractive acquisition price, tempting people to sign a contract; then they design the customer experience to maximise the chance that people will auto-renew their contract by mistake; and *then* they hike the price for that disengaged group, charging them more for exactly the same service.

This might sound like a cynical idea, but you don't need to assume bad motives to believe that these dynamics take place; the only motive at play is the maximisation of profit. And, as the chart below shows, these behaviours are feeding through into the numbers. The chart focuses on Britain's energy market where, in recent years, the gap between competitive acquisition tariffs (the bottom of the grey area) and default Standard Variable Tariffs (the green line) grew to such an unsustainable level that the UK government felt it had to intervene to cap prices.

The price of energy to UK consumers

£, per customer per year



Source: Competition and Markets Authority. Based on typical domestic consumption of 3,200 kwh/year for electricity and 13,500 kwh/year for gas. Other costs include transmission, distribution, BSUoS and environmental and social obligations.


In the UK, this problem has been most salient in the energy market, but it goes much further than this. In Britain's broadband market, prices jump 43% on average at the end of a fixed-term deal—more than in energy.⁹ In banking, 1.2 million mortgagors on standard variable rates pay hundreds of pounds more than they did under their fixed-term deal.¹⁰ Again, this is not to imply sinister motives; in these industries, bait and squeeze has simply become the equilibrium pricing strategy, making it hard if not impossible for any one company to price in a different way.

In insurance, these pricing strategies are now nothing less than the norm. Take the picture below as an anecdotal example of a practice that will be familiar. It shows an insurance renewal letter I received a few weeks ago, proposing a 50%

⁹ Merola, R. and Greenhalgh, L. (2017) *Exploring the loyalty penalty in the broadband market*, Citizens Advice

¹⁰ Merola, R. et al. (2017) *Exploring the loyalty penalty in the mortgage market*, Citizens Advice

year-on-year price increase. No-one would choose to pay this increase. And of course this is the point; the company in question does not expect anyone to *choose* to renew at this price; the business model rests on people auto-renewing by mistake. And it works: around 12.9 million UK households auto-renew their home insurance after 1 year, with vulnerable people most likely to fall into the trap.



Implications for economic policy-makers

How should economic policymakers respond to these developments? They need to adapt their approach in at least three ways.

First, these dynamics complicate some of the fundamental premises on which regulatory economics is based.

Consider the idea of ‘regret’. It came up frequently in the examples above: the free trial you *regret* not cancelling, the exorbitant insurance you *regret* auto-renewing, the extortionate out-of-bundle mobile phone charges that make you *regret* signing up to your ‘£20 a month’ mobile phone contract.

What these experiences have in common is the irksome feeling that you did not really ‘choose’ the outcome at all, despite having been technically free to do so. These outcomes therefore complicate the orthodox defence of free consumer markets—because they complicate the idea of consumer choice itself.

When a consumer market works well, with low levels of regret, consumer choice is a powerful concept. It lets the economist argue that each trade *necessarily* adds value for the consumer. After all, if the person didn't value the product, and value it more than the price they paid, they would not have made the trade.

When behavioural tricks come into play, this argument goes weak at the knees. Does the fact that you're paying for a monthly subscription, after a free trial period ended, mean you *want* that subscription? What if you have not switched your energy supplier because they make the process a hassle? Are these purchases generating a surplus for you as a consumer? Not necessarily. If the 'choices' consumers are making, in ostensibly free and competitive markets, are not necessarily raising their utility, that undermines a longstanding premise of regulatory decision-making.

So these developments should prompt a healthy reappraisal of the more orthodox principles that sit behind consumer policy. As Martin Wolf put it in a recent op-ed on technology monopolies: "policymakers must get an intellectual grip on what is happening".¹¹

Second, these developments also raise practical challenges for economic regulation. How should you govern behavioural nudges in otherwise competitive markets? And how should you respond when competition *itself* is the locus of these tricks? When companies are competing not over 'who has the best product?' but over 'who can most effectively nudge consumers into spending money they later regret?'

These are fundamental questions. And insurance auto-renewal is again a good example of how they play out in practice. In recent years, the Financial Conduct Authority identified that insurance companies were nudging their customers to auto-renew by mistake while hitting these customers with massive year-on-year

¹¹ Wolf, M. (14 November 2017) *Taming the masters of the tech universe*, The Financial Times

price increases. Auto-renewal letters had come to sound like adverts for a spa weekend: ‘Sit back and relax—you don’t need to do anything’.

To stop this deception, the FCA now prescribes the design of auto-renewal letters, requiring firms to show customers the price they paid last year alongside their renewal price (hence my letter above).¹² Now, a new battle has started: some companies are gaming the system by showing each price on a different page of the letter, or by using different calculation methods so that the two figures are hard to compare.

This crystallises a difficult challenge facing regulators. If they stick to their traditional approach, they will find themselves in an arms race, setting increasingly detailed rules to stop companies playing increasingly sophisticated behavioural tricks, and continually having to amend these rules to stop each new trick. Entirely new approaches to regulation will therefore be needed.

Third, although it’s early days, we can start to sketch the elements of a new regulatory settlement - one that fits how markets work in a world of big data.

Given the uncertainty about where these trends will lead, the task is to sketch a new approach that fits the direction of travel. It now seems all but certain, for example, that big data will become increasingly powerful and ubiquitous. And it seems clear that the speed of computers and communications networks will soar and their cost plummet, that mobile devices will take over, and that retail will continue to shift online. This means we know roughly what tomorrow’s consumer markets will look like—they will look a lot like the technology sector today.

Whether markets work better or worse in this future is an open debate. But it’s clear they will work *differently* in important ways. Prices will likely move closer to the reference price of individual consumers, breaking down the cross-subsidies that are implicit in flat pricing. Behavioural nudges will become more sophisticated, making

¹² Financial Conduct Authority (8 March 2017), Transparency in insurance renewals, FCA

it harder to claim that a person's purchase of a product or service necessarily 'reveals their preference' for that product or service. Some industries will get locked into harmful equilibriums, in which there's more money to be made perfecting behavioural tricks than perfecting pleasing products, and this will direct effort and investment into activities that don't add value for consumers.

Policymakers must not overreact to these changes. Big data creates huge opportunities for consumers, as well risks, not least by disrupting incumbents or disintermediating large suppliers. And economic regulators must strike a careful balance: it is important to maintain the clarity and predictability of the rules that government economic regulation today, while also changing fast enough to preserve the legitimacy of the regulatory regime.

Within these constraints, however, three specific shifts are likely to be needed.

One, there will need to be a greater willingness to intervene, in targeted ways, based on clear rules, to limit extortionate prices for subgroups of consumers, particularly vulnerable people in markets for essential services.

This entails a greater acceptance that *behavioural* market failures sometimes justify supply-side interventions, even when a market seems competitive and when consumers, in theory, have freedom of choice.

Two, for the reasons outlined above, there will need to be a general rebalancing away from highly prescriptive regulatory rules toward broader principles-based or outcomes-based regulation, backed by tough enforcement. This means scrapping detailed rules and, instead, requiring companies not to violate general ideas like 'misleading customers' or 'exploiting vulnerable people'.

This avoids an arms race in which regulators prescribe company behaviour in ever more detail to stop ever cleverer behavioural nudges. It also shifts the responsibility of compliance onto industry, and off of the regulator. Instead of the regulator having to check that every company has ticked every compliance box, the

companies themselves have to weigh up whether they're confident they're acting within the broad principles set out by the regulator.

Three, a fuller and stronger set of *horizontal* institutions will be needed, to make sure these new technologies—big data and machine learning in particular—benefit consumers as much as producers. Without the government's support, it might prove easier for large firms to use these new capabilities than it is for consumers themselves, and that could squeeze consumer surplus.

These horizontal interventions could include, for example, rules to make data more open, so that price-comparison intermediaries have easy access to data on tariffs and pricing, as well as rules to put more data into consumers' hands. And it might require a role for government to encourage or to seed-fund innovations that would put the power of machine learning and big data in consumers' hands. Both approaches are likely to be particularly valuable for vulnerable consumers.

In closing, it is important to note that a full response to these trends requires more than a technocratic solution. The subtle changes we're seeing to consumer markets are just one aspect of growing unease about how the world's mature market economies function. This sits alongside other concerns, such as the material crisis in living standards and the changing nature of work.

Big issues like these demand and deserve a well-informed and public economic debate, not least to surface and test public opinion. What outcomes do we *expect* from well-functioning consumer markets? What are the *ethics* of practices like price discrimination? What role do we want the government to play and what role falls to independent regulators? These are political questions as much as they are theoretical or technocratic ones, so they deserve a healthy public debate.

Title: Big data is changing how consumer markets work

Author: James Plunkett