

---

# LETTERS & NOTES ON REGULATION

---

## Heuristics and biases in regulatory decision making

George Yarrow

---

No. 3.3

July 2014

---

REGULATORY POLICY INSTITUTE

*Letters and notes on Regulation, Vol. 3*

*Published by the Regulatory Policy Institute  
300 Banbury Road, Oxford OX2 7ED, UK  
[www.rpieurope.org](http://www.rpieurope.org)*

First published July 2014

© Regulatory Policy Institute

George Yarrow<sup>1</sup>

## Heuristics and biases in regulatory decision making.

### 1. Introduction

In his thought provoking note *Applying behavioural economics at the Regulatory Conduct Authority*,<sup>2</sup> Stephen Littlechild has drawn attention to an important set of questions about the use of behavioural economics in regulation. The Regulatory Conduct Authority of the paper's title is an imaginary agency that made a brief, Brigadoon-like appearance on 1 April 2014. Its hypothetical purpose is to make use of behavioural economics in regulating other regulators.

As the title indicates the paper was stimulated by a document issued by the new UK Financial Conduct Authority (FCA), *Applying behavioural economics at the Financial Conduct Authority* (April 2013), which relied heavily upon the notion that a regulator can both identify and correct or mitigate 'biases' in consumer behaviour. The mythical RCA's analogue of the FCA document is therefore based on the proposition that an additional public authority can both identify and correct or mitigate 'biases' in the behaviour of other regulators.

Two distinct questions are raised by the paper. The first is whether behavioural economics is useful in developing a better understanding of consumer behaviour. The second is whether regulatory actions that are informed by behavioural insights and are aimed at directly modifying consumer behaviour are likely to be effective in serving policy goals, for example by correcting for observed 'biases' in consumer decision making. Answers to the questions are not necessarily the same. For example, a better understanding of economic behaviour does not logically lead to a conclusion that attempts to modify behaviour are warranted.

What follows is a short examination of these questions. The remarks build upon some of the arguments made in a longer, earlier paper.<sup>3</sup>

### 2. Heuristics and biases

It is useful to be clear about the meaning of words at the beginning of a discussion, particularly when the same words are used in both a technical literature and in more general public discourse, as is the case for behavioural economics. The online Oxford English Dictionary (OED) defines behavioural economics as *a method of economic analysis that applies psychological insights into human behaviour to explain economic decision-making*. This is a little clumsy since there is no single 'method', but it will suffice for current purposes.

---

<sup>1</sup> Chair, Regulatory Policy Institute.

<sup>2</sup> Available at <http://www.iaa.org.uk/blog/regulatory-conduct-authority-applies-behavioural-economics-to-regulators>

<sup>3</sup> "Dysfunctions in economic policymaking Part 1: Simple stories, complex systems and corrupted economics", *Essays in Regulation*, Regulatory Policy Institute, 2014.

The current upsurge in interest in the area stems from pioneering work by Kahneman and Tversky that came to be referred to as the *heuristics and biases* approach.<sup>4</sup> This labelling has not, I think, been helpful. The online OED defines heuristic (as an adjective) as *enabling a person to discover or learn something for themselves* and (as a noun) as *a heuristic process or method*. Put simply it is (as a noun) a method for discovering or learning about things.

There is nothing here that requires a heuristic to be a *simple* method, but the focus of Kahneman and Tversky's work was, as a matter of fact, on very simple, widely used heuristics<sup>5</sup> and the labelling has generally come to be taken to refer only to one end of the simplicity/complexity spectrum of methods for discovering or learning about things. The shift in meaning is unfortunate because it tends to lead to neglect of large parts of the wider range of heuristics, many of which are rather more elaborate and complex than those that have been the focus of recent interest in psychological research and its potential uses in public policy. These include many of the techniques used in non-behavioural economics which are themselves heuristic in nature (see further below).

This brings us to the problematic word *bias*. Kahneman and Tversky were clear enough in how they were using the word: to declare a bias obviously requires some benchmark against which behaviour and decisions are to be assessed and this was explicitly defined as rational choice theory, most perfectly formed in the work of von Neumann and Morgenstern.

The existence of 'biases' in decision making in this sense – behaviour and decisions that do not conform to those implied by rational choice theory – was well known before the emergence of the *heuristics and biases* stream of work,<sup>6</sup> although the descriptive language was different: a divergence between actual choices and the choices implied by the theory was described as just that, a divergence between theory and observation. Kahneman and Tversky's major contribution was to identify some of the simple heuristics in widespread use, in an attempt to account for such discrepancies.

The most usual interpretation of discrepancies between theoretical predictions and observed realities is that the model/theory is deficient. Thus, if consumers do not behave in ways implied by rational choice theory, the obvious inference is that such theory offers an inadequate account of behaviour. In such circumstances, development of alternative methods of understanding consumer behaviour is therefore obviously merited.

On the other hand, describing the discrepancies as 'biases' carries the ordinary-language implication that it is the actual decision maker, not the theorist, who has somehow gone wrong. Further, by wrong is meant something more than just making a mistake, such as a consumer

---

<sup>4</sup> See for example T. Gilovich, D. Griffin and D. Kahneman, *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press, 2002.

<sup>5</sup> The most prominent early examples are what they referred to as the availability, representativeness, and anchoring heuristics.

<sup>6</sup> The relevant stream of research dates back at least to Allais's work (1952/3), which appeared within ten years of von Neumann and Morgenstern's *Theory of Games and Economic Behavior*.

failing to choose an option that would, in a singular set of circumstances, have led to a better outcome. The word bias tends to imply a *systematic* tendency to make inferior choices.<sup>7</sup>

A claim of bias is, therefore, a very strong claim and it should properly require proportionately strong substantiation. Among other things it requires strong justification of the benchmark against which actual behaviour is to be assessed. The relevant question to ask is *What in the relevant circumstances would constitute unbiased behaviour?*, and it is typically not an easy question to answer. Decisions/choices are potentially affected by many factors, including factors that are idiosyncratic to the individual decision maker or context. Substantiation of an ‘unbiased’ benchmark requires that all materially relevant factors are taken into account.

Consider, for example, a consumer’s own choice of heuristic method. Among the factors affecting this choice are the costs of alternative heuristics for discovering and learning about products/services that might be purchased, including the costs of time taken up in gathering and assessing information and the usage costs of the scarce resource that we call cognitive effort. For a given transaction and taking account of the differing costs of effort associated with different heuristics, it can reasonably be asked: *how is it practically possible to identify an unbiased heuristic?*

In current regulatory discourse the problem of substantiating the ‘unbiased’ benchmark tends to be ignored or sidestepped. Instead, what is used is some hazy and subjective notion of how public policymakers think that consumers *should* behave, not infrequently distilled from some “academic scribbler of a few years’ back”.<sup>8</sup> The effect is to impute a wholly unmerited normative status to the scribblings.

Fortunately, there is a simple solution to these problems that does not in any way impede the effective use of psychological insights in economic analysis: stop using the word *bias* in policymaking save in those circumstances in which a convincing normative benchmark is easily identifiable or, if things are less clear cut, where the actual conduct is such as to deviate systematically from each and all of a number of plausible benchmarks.

## 2. The two questions

Having cleared the decks, let me return to the two questions posed. In relation to the first, my own view, apparently widely shared on the basis of reported reactions to the RCA paper, is that incorporating insights from psychology (the behavioural approach) *is* useful in helping understand consumer behaviour and also in understanding regulatory behaviour, although in the latter case social and organisational factors are likely to be at work alongside individual psychology and they may be the more important sources of influence on decision making. Put

---

<sup>7</sup> It be noted in passing that it is wrong in general to suppose that a bias is in itself a Bad Thing. In statistics, for example, an unbiased estimator is not necessarily the ‘best’ estimator. Thus, an estimator with a modest bias and a small variance may be preferred to an unbiased estimator with a much larger variance, because, for example, it might be particularly important in the relevant circumstances to avoid large errors.

<sup>8</sup> Keynes, *The General Theory*.

another way the behavioural approach in itself a potentially valuable, albeit rather general, heuristic or component of a more general heuristic.

As already indicated, however, it cannot be concluded from this that better understanding of consumer psychology will by and of itself increase the effectiveness of regulatory policy. Regulatory decisions will be a function of the behavioural patterns to be found in regulatory organisations as well as in the population of consumers and, to take the most obvious example, policy outcomes may deteriorate if the effect of reliance on behavioural insights about the conduct of consumers is to increase the interventions of an agency that is itself behaviourally inclined toward excessive intervention. Arguably this is just another example of the general principle of ‘second best’: there is no reason to believe that actions that would improve matters in circumstances in which all other economic conditions were optimised will necessarily improve matters in the presence of other ‘distortions’ (such as might be caused by organisational behaviour).

Market conduct will, then, be *co-determined* by the interacting behavioural patterns of both regulatory agencies and consumers and in examining this process there is, I think, one thing that should immediately strike an economically informed spectator, the disparity in market power between a regulatory agency and an individual consumer. By this I mean that the choice of one particular heuristic or method of approach to decision making will tend to have significant, market-wide effects in the case of the regulator, but it will tend not to have anything like such wide-ranging effects in the case of an individual consumer.

This points toward a similarity between the question “*when might it be beneficial to make outcomes more sensitive to the heuristics of a regulator than to the heuristics of consumers?*” (by subjecting consumer behaviour to greater regulatory influence, for example by way of ‘nudges’ or other, more stringent, interventions) and the question “*when might it be beneficial to subject market outcomes to greater monopolistic influences?*”. The latter question is one that has been subject to a great deal of scrutiny in traditional economics, opening up a rich source of potential insights.

### *Differing heuristics*

To draw on some of these insights consider again the notion of a heuristic – a method for discovering or learning about things – and the question: *how are heuristics developed and used in decision making?* More specifically, consider first the sub-set of heuristics that might broadly labelled as economic models. These are explicitly developed in the context of learning about economic behaviour and its effects, and have the characteristics that they are: simplified representations of complex realities; created or adopted by their users; capable of endless variation; compete with one another; subject to selection processes of various kinds (i.e. some methods survive and gain supporters and users, others don’t). There is also considerable variety in the ways in which models are formulated, used and judged/evaluated.

Everyday heuristics of the type studied by Kahneman and Tversky are not dissimilar, the chief differences being largely to do with complexity and formalism. Thus even when a generic type such as the *availability heuristic* is identified it will be found to exhibit great variety in use, if for no other reason than that the referencing information that is readily available to one consumer will be different from that available to another: every individual has a unique information set. Everyday heuristics can also be expected to adapt and change over time as new contexts are encountered and as individual information sets change. The same decision maker may have recourse to different heuristics in different contexts (just as an economics teacher may reach for different models or model-variants when addressing A-level students and when addressing post-graduate students). To cut a long story short, everyday heuristics, like their more complex cousins, are subject to discovery, adaptation and selection processes.

The relevant processes typically take place within social settings such as markets or other types of economic institution. In the study of such settings one of the central themes of economics from its inception has been that, compared with competitive alternatives, monopolised processes tend generally to be inferior (and often substantially inferior) in terms of promoting discovery, adaptation and selection. This applies as much to the development and use of heuristics as it does to other aspects of economic life. This is not to say that competition is *always* preferable to monopoly: there are obviously circumstances in which it isn't. However, they are dominated by the much larger number of economic contexts in which monopolistic processes perform less well than competitive alternatives.

The point to which all this leads is that, since regulatory processes lie at the monopoly end of the spectrum of ways of determining economic outcomes, when regulators seek to take responsibility for consumer decisions or to engage in behaviour modification via 'nudges' applied across markets *they are typically engaged in a process of monopolisation*. More specifically, the effect is that market outcomes become more subject to the influence of a single economic agent, the regulator.<sup>9</sup>

The various pitfalls associated with monopoly therefore indicate that, as a minimum, wary scepticism is advisable when contemplating the extension of regulatory activity into areas where competition is active.

#### **4. Bias in regulatory heuristics: the example of 'market failure'**

Although consumer 'biases' have been the centre of attention in policy debate, a number of potential regulatory 'biases' have in fact been examined in recent papers published by the Regulatory Policy Institute.<sup>10</sup> This short paper is not the place for general extensions of that material, but there is one form of bias that is both closely related to points raised above and

---

<sup>9</sup> Unless the interventions are offset by reductions in other interventions.

<sup>10</sup> See: C. Decker and G. Yarrow, "On the discovery and assessment of economic evidence in competition law", *Studies in Regulation*, Regulatory Policy Institute, 2011, which focuses on confirmation bias; R. Finn and S. Less, "Capture of independent sectoral regulators", *Letters and notes on regulation*, Regulatory Policy Institute, June 2013, which addresses sources of social/political biases; and G. Yarrow, *op cit.* (footnote 3).

central to the current role of behavioural economics in UK public policymaking, particularly in financial services markets. It is associated with what can be called the *market failure heuristic*, which appears still to be widely used in economic regulation. This heuristic is based on comparisons of observed market behaviours and outcomes with the hypothetical behaviours and outcomes implied by an abstract economic model/theory that *inter alia* rests on an assumption of zero transactions costs, including zero costs of acquiring, processing, transmitting and making use of information.

No sane economist believes that the assumptions of the relevant model/theory reflect reality, but the model/theory itself continues to be taught on the basis (often not explained to students) that it can, depending upon circumstances, sometimes be a useful heuristic. Unfortunately (and similar to what has happened in relation to rational choice theory) in addition to being used as a contingent heuristic, to be adopted when it is potentially illuminating and to be set aside when it is not, the model/theory has come to be used as a *normative* benchmark in the assessment of market performance. The market failure heuristic labels divergences between market realities and the theory's implications as *market* 'failures', implying that there is (or, in more sophisticated uses, might be) something wrong with the market: market performance *should* be better (hence the normative aspect). The result is a bias toward over-diagnosis of 'failure'.

Perhaps the most egregious example of the effect is to be found in the view that asymmetric information is a market failure. Whilst it is possible to find contexts in which regulatory measures to reduce information asymmetries can improve market performance, the fact of the matter is that, for the most part, asymmetric information is highly efficient: it is an aspect of the division of labour that has contributed enormously to economic progress over the centuries.

In summary then, just as in the case of use of the rational choice theory benchmark when assessing consumer behaviour, the use of an irrelevant, normative benchmark is implicated in the establishment and persistence of regulatory biases.

There is something of an irony here in that greater use of behavioural economics in public policymaking potentially offers a way back out of the dead end to which the market failure heuristic has led. It is notable, for example, that the concept of market failure had been particularly influential in UK financial services regulation in the past, specifically in relation to issues involving asymmetric information, and that the newly established FCA has, in emphasising behavioural approaches, subsequently been at the forefront of UK regulation in reducing reliance on the market failure heuristic. A transition toward greater reliance on behavioural evidence can also be interpreted to a return to a more classical tradition in economics: Smith's *Theory of Moral Sentiments* is packed full of psychological insights and the most distinctive component parts of Keynes's *General Theory* (the consumption function, animal spirits in investment behaviour, the speculative demand for money, sticky nominal wages, etc.) have behavioural (rather than deductively rational) foundations.

Unfortunately, in taking a step away from reliance on a defective heuristic (market failure) based on an irrelevant normative standard, financial services regulation appears to be at risk of

leaping to reliance on another irrelevant, normative standard (rational choice theory), as manifest in the frequent use of the word ‘bias’. The words frying pan and fire come to mind.

This leads to a second irony: in the first of its incarnations in the modern period, led by Herbert Simon<sup>11</sup>, behaviourism was associated with a rejection of the notion of ‘optimisation’, the idea that a best outcome could be computed, and with an emphasis on the centrality of ‘search behaviour’ in the face of limits on knowledge arising from the limited information processing or cognitive capacity of any one individual or organisation (i.e. bounded rationality). This view-of-the-world is not far distant from the emphasis placed on processes of discovery, adaptation and selection in dynamic economics (in contrast to the emphasis of economic statics on notions of optimisation and equilibrium). In its second incarnation, however, behaviourism appears unwittingly to have taken a step backwards by giving a ghostly afterlife to concepts used in static theorising about a frictionless reality, which have been resurrected in their new manifestations as normative benchmarks against which consumer behaviour is to be judged, and usually found wanting.

## 5. Concluding thoughts

My initial reaction to the RCA paper was that its cutting edge derives mainly from two aspects: (a) the invitation to consider whether the original FCA document was based on a double-standard in which consumers are to be judged according to more demanding criteria than are regulators and (b) the nature of the institutional set-up that it implies: a new regulator to regulate the conduct of other regulators. In relation to the second, I was reminded of a passage in Stephen Hawking’s book *A Brief History of Time*:

*A well-known scientist (some say it was Bertrand Russell) once gave a public lecture on astronomy. He described how the earth orbits around the sun and how the sun, in turn, orbits around the center of a vast collection of stars called our galaxy. At the end of the lecture, a little old lady at the back of the room got up and said: "What you have told us is rubbish. The world is really a flat plate supported on the back of a giant tortoise." The scientist gave a superior smile before replying, "What is the tortoise standing on?" "You're very clever, young man, very clever," said the old lady. "But it's tortoises all the way down!"*

In the regulatory cosmos there appears to be a tendency towards tortoises all the way up, which has a feel of ‘too many tortoises’ about it, particularly since each tortoise enjoys significant market power/influence. The RCA should, therefore, clearly retain its Brigadoon characteristics and not be given a more substantive presence, but this conclusion obviously does little or nothing to contribute to the resolution of current issues surrounding the use of behavioural insights in regulation.

On the other hand, a regulatory approach to consumer behaviour that is based on greater awareness of some of the dysfunctions of regulatory bureaucracies and that is less inclined toward asymmetric judgmentalism does not necessarily preclude regulatory actions that are akin to nudging. The crucial questions seem to me to be not *whether* the use of behavioural

---

<sup>11</sup> See, for example, “A behavioral model of rational choice”, *Quarterly Journal of Economics*, 1955.

insights by a regulator is likely to be helpful in furthering the interests of consumers, but rather who decides what those interests are and how these matters are to be decided.

Nudges that are determined by a regulator may be described as benevolent or liberal paternalism because they leave the final choices about transactions to the consumer, but they are paternalism nonetheless and have the undesirable property that they represent exercises in monopoly power, undertaken (in the words of competition law) *independently* of the consumer and with all the weaknesses that such a decision process entails.

In contrast, the classic liberal answer to the who and how questions is that consumers should decide matters themselves, either in competitive conditions or, where some element of collective choice is involved, by means of a process that commands their approval and consent. Nudging can be consistent with this view, but only if it is implemented with the informed consent of the bulk of the consumers who are likely to be affected.

In effect, regulatory actions, including nudging strategies, represent changes in the market rule-book and have market-wide effects. The underlying questions therefore relate to participation in rule-making. Regulators have a distinct role to play in rule-making, but should they play that role independently of the views and preferences of consumers? I suggest not: it is the regulatory equivalent of abuse of a dominant position in competition law.

Developing participatory arrangements for market governance that involve consumers (and businesses) as well as regulators is a broad agenda, but there are modest, practical steps that can be taken along the way. For example, in circumstances where a regulator believes that there is a strong case for an intervention motivated by behavioural considerations, that case might be put to consumers and their consent sought: at the end stage of a consultation process, arguments both in favour and against a proposed market-wide intervention directed at consumer behaviour could be promulgated and consumers polled in some way or other. Such an approach might have the additional advantage of encouraging regulators to think hard and write in plain English, but its chief advantage is that it would introduce at least a limited form of competition into the regulatory arrangements: opposing views would be exposed to a competitive selection process in which consumers are sovereign.

Finally, in any area of analysis it is generally informative to look at the history of ideas, and the use of psychological insights in economic policymaking is no exception. Examining the work of Smith, Keynes and Simon, it is striking that their (considerable) reliance on psychological insights is used to help understand the behaviour of the *individual* units of an economic system (consumers, workers, investors, firms, etc.), but not in any direct sense to understand the workings of an economic system as a whole or of sub-systems such as markets.

This suggests that, for the purposes of understanding and regulating markets, behavioural insights are likely to be most productively used as inputs into assessments that are focused on the structures and processes (i.e. on the workings of economic *systems*) that have been the traditional concern of economics and political economy. If, however, contemporary interest in behaviourism diverts attention and effort away from analysis of economic systems, the effects on policymaking cannot be expected to be benign.