2 Behavioural Learning and the Design of Incentives

2.1 Learning and Incentives

In a world of constant information, we are beset with new scientific research findings every day. Not all of this information becomes knowledge. Not all knowledge is used to make better decisions.

Only a small fraction of the available, and potentially useful, information becomes knowledge. Given that individuals have a limited cognitive load, information acquisition is cognitively costly. Hence, in judging new information, we must assess its value in terms of its relevance (is it salient to me and worth learning?) and accuracy (is it correct?). Yet, the relevance and accuracy of information may not be objectively or scientifically known. Indeed, Popper's approach to a science is that any scientific hypothesis always has the potential to be falsified (Popper, 1959). Hence, necessary judgments require some degree of confidence in the information sources. Understanding the way in which we process information is fundamental to understanding how new stimuli affect actions. For incentives (stimuli) to change action, they need to steer behaviour in the same way as information needs to be deemed to be relevant and accurate.

Nonetheless, understanding how information affects behaviours is far from trivial. Learning takes place both consciously and unconsciously. Information sources vary greatly in their personal salience to someone, and between people. The degree to which information sources fit with social norms or existing beliefs and identities matters to its impact on behaviour. Over many years smokers, for example, have been so bombarded with information about the health effects of smoking that some estimates suggest that smokers overestimate the risks of smoking (Viscusi, 1990). We learn from our own actions as well as deliberate inaction, in addition to observing others (social learning).

Introductory economic models assume that information is incorporated into decision-making through a process of information updating. If information does not reach its 'destination', then it is typically attributed to some form of information asymmetry, in which one party holds an information advantage. This is often seen in insurance and doctor patient decisions, for example, which give rise to a game of information advantage. The individual purchaser of health insurance will often know their risk profile and likelihood of a claim better than the insurer, and the insurer will, in theory, counter this with 'small print' in the contract limiting their exposure to future claims. Physicians or pharmacists paid on a fee-for-service basis might hold

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

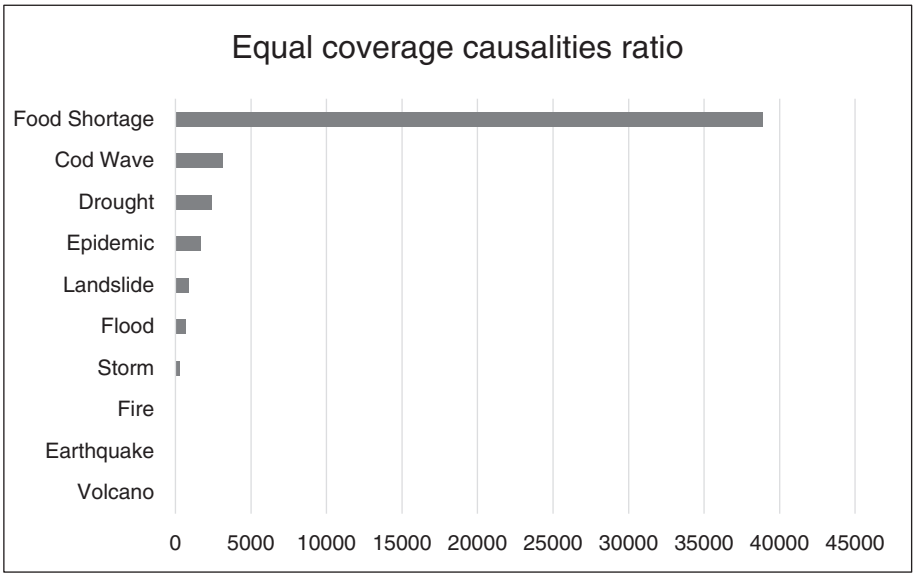

Figure 2.1 Inconsistent attention

a significant information advantage over the patient, leading to the risk of excessive health costs through 'supplier-induced demand'.

That said, not all information is equally salient. People often fail to update their beliefs even when the information reaches them. Cognitive limitations, emotions or inattention are often cited as underpinning reasons. A classic example that attention matters is provided by Eisensee and Stronberg (2007). They estimated that the number of deaths required in a range of events in order to garner attention equal to a single casualty from a volcano eruption (see Figure 2.1). Interestingly, an epidemic would require 1,696 deaths, a drought 2,395 deaths and a famine 38,920 deaths.

Unlike what we learn in an introductory economics textbook, a behavioural analysis accommodates the fact that people differ in how they use information to inform their behaviours. This includes systematically ignoring information, Hence, it is evidently important to understand better when such incentives might be ignored or will not be internalised into personal expectations and beliefs.

The process of information updating has to clear the substantial hurdle presented by motivated reasoning, which generates a selective 'belief update bias' in which people 'rationalise away unwanted evidence' (Kappes and Sharot, 2019). It even appears to be the case that when the solution to a problem is at odds with personal ideology, then this leads to denial of the problem itself, known as 'solution aversion' (Campbell, 2018).

To actually change behaviour, information needs to be distilled, internalised and synthesised, along with an underlying assessment of the messenger's intentions and goals. During the COVID-19 pandemic, it was observable that in the United States with a Democratic Party government calling for people to vaccinate, Republican voters were less likely to be vaccinated than Democratic voters, and in the United

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

Kingdom with a Conservative government, it was Labour voters less likely to be vaccinated (Klymak et al., 2021). This is what has been called the 'messenger effect' (Martin and Marks, 2019) in the belief updating bias.

Smoking is a classic example of this. In most circumstances, it would be hard for someone to claim nowadays that they do not know that smoking is a health risk. Even when smokers are fully aware of smoking risks (Murphy et al., 2014), smoking remains prevalent. Estimates for the United States suggested that around 480,000 people over the age of 35 were dying prematurely from smoking-related causes, and 122,000 in the United Kingdom (Peto et al., 2021). We will discuss this 'prevention failure' in Chapter 7. Clearly the availability of information alone does not guarantee learning and behavioural adaptation.

2.1.1 Understanding Learning

A central claim of behavioral economics is that we need a much more nuanced appreciation of how learning occurs. Indeed, learning has both a cognitive and a tangible cost to individuals in terms of effort, time and resources. Given the scarcity of time, money and cognitive resources, it seems only reasonable to learn what one can expect to bring some visible returns, or to rely on less costly unconscious learning. Therefore, the effectiveness of incentives will largely depend on whether information activates individual motivations and engages with the way people actually learn. What we call bounded learning models assume that people often prefer to use a shortcut in such decisions, which we usually refer to in behavioural economics as 'heuristics' or 'rules of thumb'. Such shortcuts are incredibly useful in minimising the costs of processing all the information and decisions faced on a daily basis. We might, for example, be predisposed to reject new evidence showing benefits from genetically modified foods simply because, as a rule of thumb, such food seems unnatural, and perhaps the opposite for organic produce or 'local' food. We might be hard to convince that a vaccine is safe if it has been trialled quickly for a pandemic response, regardless of the scale, robustness or intensity of those trials.

We learn from own internal sources. These may be past experiences, memories or stories that have been told to us. They influence what we like and dislike and influence our *hedonic forecasting* capacity, that is, whether a specific experience will improve our well-being in the future. Behavioural incentives are cast into this whirlpool of influences, and learning effects are central to whether they are noticed and effective. This applies regardless of whether the incentives involve money, social stimuli or gentle nudges within the choice architecture of decisions. Making learning easier where it most matters is particularly important. The easier it is, the less attractive it will be used to employ one of the usual bias-inducing heuristics. Besides, easing learning, incentives can be 'primed' somehow to make them more salient to the recipient, as can social norms, and narratives give cognitive structure to actions (instead of cognitive dissonance).

Learning is affected by our priors, the views of the world we have before processing new information. Our priors include what we already know (or think we

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

know), our attitude towards absorbing new information, levels of trust, as well as the credibility we attach to an information source. These priors affect the learning process alongside any entirely external stimuli shaped by others. Indeed, information sources are an external stimuli. Accordingly, we face external stimuli that might be broadly targeted across an entire population or group of people, as is normal for many public health campaigns. Analogously, we also encounter external stimuli that have been individualised for us, perhaps filtered to us by a trusted 'gatekeeper' or 'agent' such as a physician or nurse. Of course, we also learn much from others, without particularly targeting of anyone (social learning). We may wish to emulate people we admire, or groups we belong to or the expected behavior of those causes we identify with, just as much as we may wish to differ from certain people or groups. The avoidance of disapproval by our peers (Berger et al, 1977) is a powerful motivator, and a spur to learning how we might avoid social mistakes.

Crucially, many of these mechanisms for learning can involve unconscious action. They can occur through adaptation to an environment, and be the result of evolutionary mechanisms. In environments where there is prevailing uncertainty, people tend to form beliefs and expectations about the likelihood of future events by relying on assumptions about the world that reflect their own cognitive biases. 'Projection bias' or 'self-forecasting bias' can explain the overestimation of how much our future selves will share the same beliefs we hold today. Similarly, whilst exercising, we will likely overestimate our future physical activity, or when feeling particularly hungry, we overestimate how much we can eat when ordering food. To the best of our knowledge, no one who ever said 'I'm so hungry I could eat a horse' has eaten a 500 kg horse.

Old learning theories in psychology such as the Theory of Planned Behaviour (Ajzen, 1985) predict that people form their beliefs, from which they form behavioural intentions, which in turn determine behavioural actions. Individuals attempt to be consistent between their planned behaviour and their actual behaviour, but nonetheless exhibit some level of 'cognitive dissonance' a contradiction occurs between beliefs and actions (Akerlof and Dickens, 1982). Our capacity for self-control varies, thus affecting our resolve to deliver on earlier plans, and self-doubt can creep in regarding our self-efficiency to execute our plans (Bandura, 1999). Nevertheless, the core message for behavioural policy from the Theory of Planned Behaviour is that merely *having* a prior plan makes a particular action more likely is an important one. We also know, however, that subconscious factors play a role, alongside conscious plans. In a changing world, many circumstances are unlikely to be foreseen and planned for in advance. This is particularly true when it comes to health.

2.1.2 The Limits of Demand for Health Models

Economic theories of health behaviour have relied on the assumption that individuals generically demand 'health'. 'Health' is depicted as a household-produced good resulting from engaging in 'healthy behaviours', which are formed under the ideal assumption of perfect information. In such a framework, health produces two types of utility: direct utility arising from the *consumption* of health; as well as indirect utility from

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

investing in health, producing more productive time, which in turn makes it possible to produce more health and other commodities. The consumption benefit is the 'psychic' rate of return, whilst the investment benefit is the marginal monetary return on health investment. This basic model assumes that people have full information and can foresee (and thereby plan for) the future consequences of their choices. However, this ideal model is, of course, open to significant criticism for its descriptive validity, even if it is believed to have normative validity in depicting how health decisions *should* occur. Muurinen (1982), for example, describes it as unrealistic because it fails to account for information failures that undermine judgements about future health states.

Within the demand for health models' behaviour is predicted to change as a result of alterations of the budget constraint, perhaps through changes to taxes or income transfers. Similarity, behavior is sensitive to changes in the availability of health information. Although some studies reveal a causal positive relationship between education and health behaviour (Cutler and Lleras-Muney, 2010), others show just mixed results (Currie and Moretti, 2003; Clark and Royer, 2013). Education appears to exert two independent effects. It influences both knowledge and ability. Enhanced ability makes it easier to interpret health knowledge. However, knowledge alone does not give rise to better decisions. There is, for example, a body of evidence examining whether medical doctors make better health decisions. It seems that being a doctor and or a close relative of a doctor (for which we assume significant health knowledge) actually reduces adherence to medicine prescriptions (Finkelstein et al., 2021). Similarly, physicians are only slightly more likely to use high value care than non-physicians (Frakes et al., 2021).

2.2 Learning and Information

2.2.1 Sources of Information

Behaviours are the result of preferences, which have themselves been shaped by beliefs before leading to action. The sources of information that successfully produce preference-shaping beliefs are, as mentioned above, either generated internally or produced externally. Information can be *internally generated* or inferred from personal experience, reflection or generalisation (from both induction and deduction). This information source depends on our individual capacity to create information (reflective personality). In addition, information can be *externally produced*, and learning takes place by observing others' 'cues' (social learning) which leads us to update our own knowledge stock. We may, for example, develop our views of mask-wearing during a viral pandemic by seeing that most people like us are (or are not) wearing one, thus engaging in social learning. Alternatively, we may directly and consciously update our knowledge of the subject by listening to viral disease experts. Such external information may have been standardised for a whole population, perhaps over the national media, or personalised by a clinician we know and talk with.

The effectiveness of a piece of information to affect beliefs and preferences will likely vary over time. At some phases in life, people are more prone to the formation

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

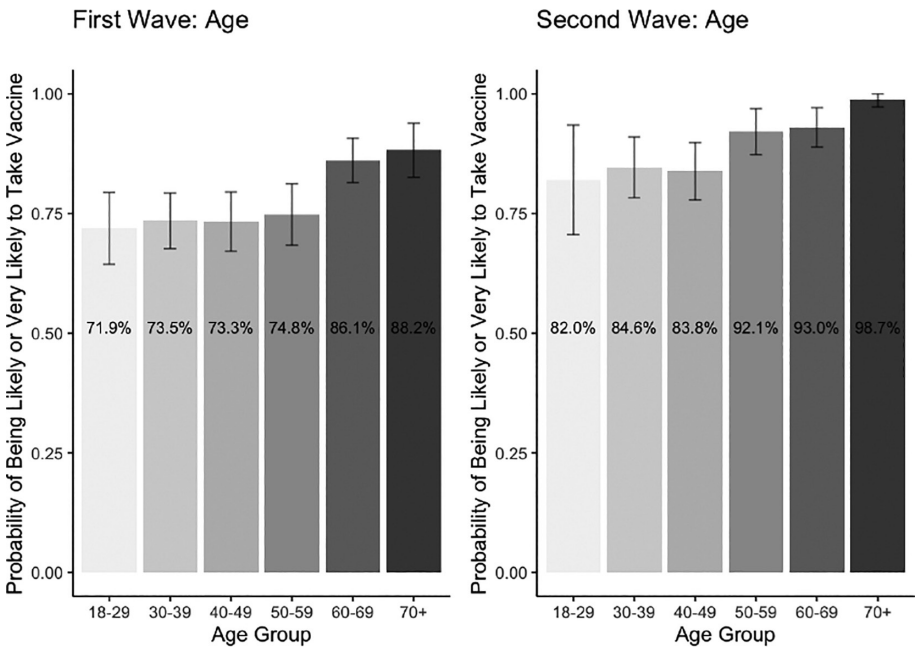

Figure 2.2 Experience effects
Source: YouGov 2020–21 several surveys.

of core beliefs; our 'impressionable years', or early adulthood (Krosnic and Alwin, 1989). For the most part, information can exert a limited effect on behaviour. Indeed, experimental data shows that different types of information about vaccine risk either does not change behaviour or shows a negligible effect (Coraece and Garber, 2014). Sometimes new information is not processed because it comes up against our core beliefs, giving rise to knowledge resistance, which we will turn to later in this chapter.

We learn from first-hand experience in a different way to other sources. Experience from exposure to events and outcomes 're-wires' our brains (Malmendier, 2021). People who have been personally exposed to a specific experience, perhaps a health or economic shock, will then respond differently to these events than those who have not, even if they are fully informed (Akerlof, 1983). Personal experience modifies beliefs and attitudes differently to the provision of information that such an event could occur. However, these effects are very domain specific. One example of how learning affects attitudes is shown in an Oxford University survey study conducted by YouGov during the first and second waves of the COVID-19 pandemic in the United Kingdom. Across most age groups, pro-vaccine sentiment increased between the pandemic waves, especially amongst those who are in the 50–59 age group. As personal experience increased with the virus, the intention to be vaccinated grew too. That said, other experiential factors would, of course, been involved in this shift, not least the time people had spent in 'lockdown' (Figure 2.2).

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

2.2.2 Rational Learning

One way to conceptualise how individuals learn is to assume that individuals form *rational expectations* – that they are the 'idealised' fully rational agents with perfect information (Erev and Roth, 2014). Under rational learning, individuals have complete willpower and no present bias (no conflict between the preferences of current and future selves). These fully economic beings will optimise their utility based on their existing knowledge of probabilities and pay-offs as represented in games, involving the range of potential scenarios. In this case, these rational individuals will be making their choices based on the maximization of their expected utility (EU), typically conceptualised as the weighted average of the utilities of the different states of the word (U_a, U_b), in which the weight is either the objective or subjective probabilities (p, 1−p) as follows:

$$EU = pU_a + (1-p)U_b$$

Nonetheless, existing evidence suggest that expected utility maximisation is more the exception rather than the norm, hence its unlikely that it can serve as a descriptive model of human behaviour (Schoemaker, 1982).

2.3 Search Costs

One of the limitations of naive learning models lies in that learning is far from costless. Active learning carries both tangible and cognitive costs, as discussed previously. Markets have frictions, which explain why individuals stay with the same insurance contract, bank or doctor for a long time. Indeed, not only is information costly, but often it is asymmetrically distributed amongst certain individuals. This creates a situation in which some (experts) have much better information available to them than others. It also takes time to gather and interpret new information. This explains why we have to be strategic in our learning behaviour. We selectively pay attention to different kinds of information sources based on the expected costs and benefits of such types of information, and we are also subject to 'regret aversion' in gathering information in areas for which we have developed a fear of missing out (Thaler, 1985). This is a behavioural regularity loved by advertising agencies. Advertising has the effect of reducing search costs, but very selectively.

2.4 Projection Bias and Rare Events

One of the claimed routes to bounded learning is that, as a shortcut, we will simply choose the option that appeared to have worked in the past and project it to the future. This becomes a 'projection bias' when no account is taken of changed circumstances and the past is just projected into the future. We overestimate the extent to which our future will look much like the present. Lowenstein et al. (2003) show that projection bias explains why people over-consume and under-save early in life, with little

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

engagement in preventative behaviours. But we do not just fail to anticipate needs later in life. We also attach too little weight to the probability of rare negative events occurring, using past experiences to make decisions today. Generally speaking, there is a tendency to attach excessive weight to the future likelihood of rare positive events, making people likely to gamble in the hope of an exceptionally positive outcome. We will come back to this subject in our later discussion of Prospect Theory.

2.4.1 Bayesian Learning

One way to conceptualise information updating is to assume that most decisions are based on re-evaluating pre-existing information. That is in adjusting a prior belief, a prior, influenced by cultural background and social environment. Through this adjustment process, each piece of new information incrementally decreases or increases the estimated belief (or probability) that a hypothesis is correct (e.g., cancer risk). Figure 2.3 illustrates an example of how the distribution of prior beliefs adjusts after some new information is revealed if individuals were Bayesian updaters.

The outcome of the assessment for each piece of new information depends on how someone weights this new information against prior knowledge. For instance, risk perceptions π are formed by weighting prior risk perceptions ρ_0 and new risk information ρ_1. Hence, risk perceptions result from the relative weight people give to prior beliefs (μ) and to new information (ϑ):

$$\pi = \frac{\mu\rho_0}{\mu+\vartheta} + \frac{\vartheta\rho_1}{\mu+\vartheta}$$

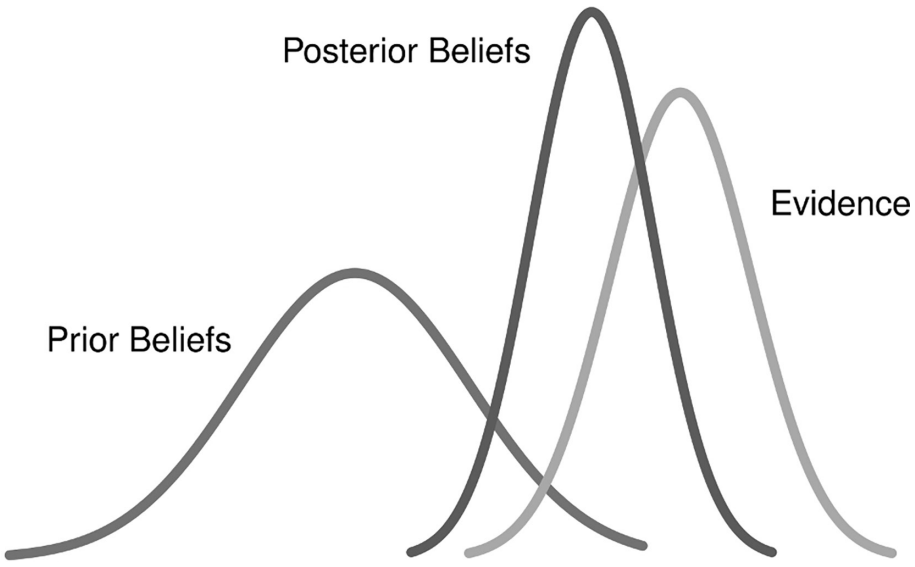

Figure 2.3 Belief update with evidence
Source: www.analyticsvidhya.com/blog/2016/06/bayesian-statistics-beginners-simple-english/

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

Bayesian learning might still work on a qualitative basis where individuals revise their judgements according to narratives to which they are exposed. Judgements may also be influenced by the particular salience for them, of some types of information. That is why 'priming' some types of information can have a central influence on belief formation (telephone calls are more effective than emails in gaining attention). There appear to be significant differences in the uptake of health information dependent upon both the information source and the receptor. Smith (2011), for example, found that elderly adults, white people, and those with relatively high incomes are more likely to trust their health professionals' information than other groups. Young people, who have grown up with internet use, are more likely to use and trust health information on the web than others. Similarly, highly educated people are more likely to consult health professionals than those with lower education levels, creating challenges of inequity in health system use even where access is free of charge (Dixon et al., 2003). Trust can have considerable importance within a health crisis dependent on individual behaviours, as seen during the COVID-19 pandemic. In such a situation, trust in government, in medicine and health agencies and in the life sciences industry all have a bearing on important aspects of behaviour such as compliance with pandemic restrictions (Costa-Font and Vilaplana, 2023).

2.4.2 Bounded Learning

When time and attention are scarce resources, people will adopt easy learning techniques to support decisions. These shortcuts in the learning process produce bounded learning, which creates risks of harmful bias. Bounded learning can also be the result of a failure to understand information due to limited technical ability or the relevance of information, perhaps due to limited objectivity or imagination.

Le Grand and New (2015: 83) describe four reasoning failures that produce 'bounded rationality':

- Limited technical ability
- Limited imagination/experience
- Limited willpower
- Limited objectivity

The same argument could be made in relation to bounded learning, and the relationship between these four traits and those discussed in this section are easy to spot. These leads people to economise on learning from new information, or even ignore information that seems unimportant or challenging.

2.4.3 Emotional Learning

Experiences that are associated with strong emotions help us remember them better, although we may also tend to write off bad experiences as 'bad luck' and refuse to learn from them. Emotions matter when learning new information (Loewenstein, 2000), as learning is heavily influenced by the role of 'visceral factors'.

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

These include regret, anger and disappointment, and refer to factors help some types and pieces of information really grab our attention. From this assessment, it is possible to simplify our decision-making: Utility (U) is dependent on $U(C,S)$, where C refers to consumption and S visceral states, so $dU(C,S)/dS>0$. A practical example would be that food tastes better when we are hungry, and warmth when we are cold.

The consideration of visceral factors is especially important when it comes to human evaluation of risks. Lowenstein has described 'risk as feelings' (Lowenstein et al., 2001). Emotions act positively or negatively. They reinforce beliefs, such as fear of a needle or fear of side effects from medication. Badger et al. (2007) argue that people who have not experienced an event, such as the craving felt by an addict, are unlikely to accurately predict their motivational force, or how they would react. They also show that, when their addiction is satiated, heroin addicts themselves underestimate the power of their craving. This is another example of 'projection bias'. One way to adjust to this bounded learning about the anticipation of events is to introduce prompts, or forms of reinforcement, that increase the salience of past experience and its likelihood of repetition.

2.4.4 Adaptive Learning

When decisions are repeated over time, such as consuming a certain food, playing a sport or taking a test, learning results from how we adjust decisions over time. We respond to the repeated experience. One form of adaptive learning is what can be labelled as 'analogic decision making'. This is when people mimic their planned reaction to analogous behaviours. Finding analogic decision frameworks help us simplify complex decisions. This form of learning was especially relevant, for example, within the context of learning about COVID-19 during that pandemic. At the start of the pandemic, the comparison to seasonal influenza was common. Narratives quickly shifted to language and understanding that it was, in fact, a very different kind of virus. Those who wanted to bring attention to severity used the emotive term 'war' (Martínez García, 2021) in relation to tackling the virus. This also has an analogic quality in connecting COVID-19 response to war.

2.4.5 Social Learning

Learning affects identities and the roles we choose, as social animals, to play within society. We learn from others by observing, imitating and modelling. Even if you have never experienced the death of a loved one form COVID-19, you would probably be able to figure out what are the consequences yourselves. Non-verbal cues are important in learning about the environments in which we live. Mental models are formed by worldviews, narratives, categories, concepts and identities. We display co-ordination utility in our clothing choices (e.g., fashion) as we try to 'fit in'. We are 'enculturated actors' (Hoff and Stiglitz, 2016) of repeated games, coordinating preferences that result from maximising a complex utility function:

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

$$Max\ U\ (X_i, X^{t-1}{}_i, X^{t-2}{}_i, \ldots, X^{t-1}{}_{-i}, X^{t-2}{}_{-i}, \ldots, I_i, I_{-i})$$

In this function, U is our utility X consumption in the period $t-1$ by an individual i, and by other individuals I, subject to a common social environment represented by an individual's identity I_i, and that of others I_{-i}. Hence, our actions are determined by reciprocal determinism, in which our own behaviour is influenced by the social environment which has shaped past behaviours (X_i). Our actions and those of others' shape who we are. We emulate others we admire ('we fit in'), based on what we believe others will do (X_{-i}). We choose the clubs to join (I_i) as do others (I_{-i}). Belonging to a club e.g., vegeterian or teetotaler might give rise to behaviour along the line of the values of the club.

2.5 Biases in Learning

2.5.1 Framing and Prospect Theory

Daniel Kahneman and Amos Tversky (1979) proposed 'Prospect Theory' as a critique of expected utility based on their observation that decisions involving risk and uncertainty do not comply with its predictions, whether in treating gains and losses equally or in showing consistent risk attitudes. They argued that people evaluate risky choices not on absolute outcomes above or below zero, but on the 'prospect' of gains and losses against their psychological reference point (commonly the status quo) as depicted in Figure 2.4. Attitudes to risk switch around the reference point, with strong aversion to small losses but a risk-seeking attitude to potential gains. They also argue that risk attitudes are not consistent but vary according to the scale of prospective

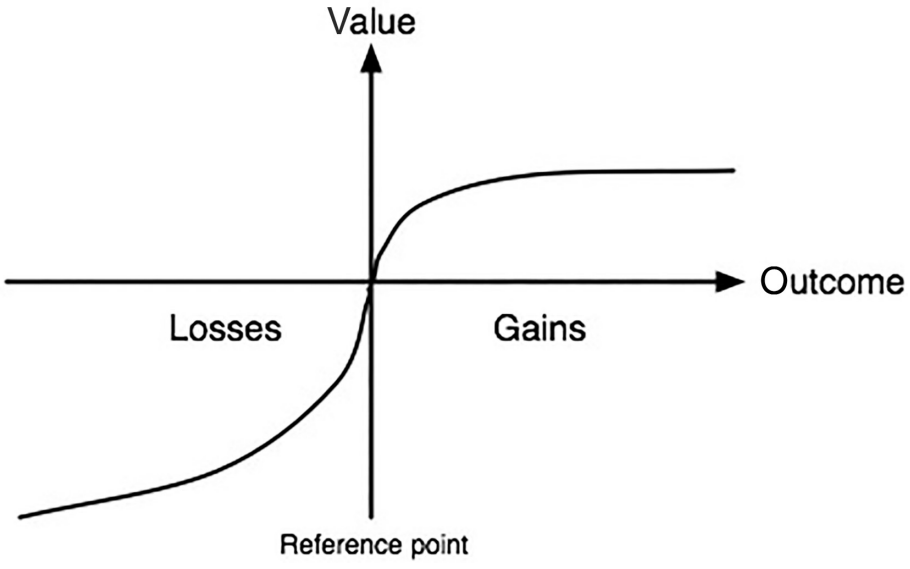

Figure 2.4 The Prospect Theory value function

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

gains or losses. In a graphic depiction of Prospect Theory, this heavy psychological value (pain) attached to prospective losses near the reference point is reflected in a steep initial curve below the reference point as the pain of loss is strongly felt. The curve for gains near the reference point is much gentler.

Most people prefer the certainty of winning £500 over a 50% chance of winning £1,000 or nothing, demonstrating risk aversion because the expected utility of both is the same. But most people would also prefer to gamble on a 50% chance of losing £1,000 as an alternative to the certainty of losing £500, showing risk-seeking behaviour. Hence, the way a decision is presented influences choices (irrespective of the informational content). Tversky and Kahneman offered the following example to demonstrate this point, using a policy decision over possible disease interventions. The first question is presented in a gain frame. Program A is the most popular choice. This shows a preference for certainty in a gain over the possibility of much higher gains but also the risk of no gain at all (Figure 2.5).

In their second scenario they apply the language of a loss frame. Under this framing of the same problem Program D proves to be the most popular option. With a loss frame of people 'dying' instead of being 'saved' by the intervention people opt for the gamble in order to avert a certain loss of 400 lives.

This has been a very brief introduction to Prospect Theory, but its influence on behavioural economics since 1979 becomes obvious through the remainder of the book. The award of the Nobel Prize in Economics to Kahneman in 2002 (Tversky sadly died in 1996 at the age of 59) seemed an important milestone in the acknowledgement of the huge influence and development of behavioural economics as a recognised discipline.

2.5.2 More Information Might Lead to Risk Overestimation

Other biases that frame decisions involving risk and uncertainty include the fact that we tend to overestimate risks that are highly publicised. Think back to the earlier example of the variable attention to deaths from volcano eruptions. Kahneman calls this the 'availability' heuristic, which we will discuss shortly. This may be part of the reason why people worry about some rare diseases more than some common diseases, or some cancers (Viscusi, 1990). They also tend generally to overestimate the probability of many minor risks and underestimate major risks (particularly around activities that involve some pleasure). However, not all information is equally salient, as 'dread' lends psychological weight to some health fears, triggering visceral responses. Think, for example, about the relative fear of 'cancer' and of diabetes.

2.5.3 More Information Is Not Better: The 'Paradox of Choice'

Introductory economics assumes that more choice and more information are always preferable to a limited supply of both. However, evidence suggests a different picture; that more choice options fail to improve well-being due to 'choice overload'. However, whether more information is preferable depends on the type of choice, more

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

Figure 2.5 Probability weighting
Source: Tversky, A. and D. Kahneman. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481): 453–458.

specifically, on functional or utilitarian matters (Dhar and Wertenbroch, 2000). In situations in which it is undesirable to spend time on decisions, there are modest returns to having a lot of choice. Whatever option does the job is good enough in this case. When it comes to some hedonic choices, however, the heightened benefit motivates the investment of more time and effort for the best decision.

This 'paradox of choice' (Schwartz, 2004) is supported in analysis of the system of open enrolment for US health insurance. When households must choose their own

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

Medicare Part D (pharmaceuticals) insurance plans, the evidence suggests that people struggle to pick the plan that is most suitable for them (Abaluck and Gruber, 2011). The more Medicare prescription drug plans available to them, the more the quality of decision-making may be diminished (Hanoch et al., 2009). More recently, guideposts and signals have been introduced to try to address this problem and to help people navigate health insurance plan choices. In the US health insurance marketplaces, metal categories (platinum, gold, silver, bronze) are used to denote plan features, which include the monthly premiums and the level of 'deductible' (portion of claims) paid by the insure in a year, known in other systems as the 'excess' that the client must pay (CMMS, 2015). In any attempt to improve the navigability of complex choice sets, there are also important decisions that could be taken on the basic choice architecture. Which plan should be listed first? Should there be a default option, and if so which would this be? Once the paradox of choice is acknowledged, then there are decisions too for any would-be 'choice architect'.

2.5.4 The 'Availability Heuristic'

People tend to give considerable weight to information that is easily available to them – things that are easy to recall, perhaps having a high media profile, or which conjure up a memorable mental image, or on which they have some prior knowledge or recent experience. Most individuals are more likely to insure for flooding after a flooding experience, but only for a while until the availability of the experience wanes. They might select a health insurance with a small deductible after being hit with a painful bill when needing care, but again the effect will probably diminish, and the lure of a lower premium becomes irresistible. Similarly, governments are much more likely to invest in pandemic preparedness after being hit by a pandemic, but after a while memories of the pandemic fade, and people revert to their old priorities. Sarah Lichtenstein and her colleagues used experiments to reveal tendencies to under- and over-estimate the frequency of lethal events (Lichtenstein et al., 1978). They found, for example, that people thought death by homicide was almost as frequent as death by stroke, when in fact stroke was killing at least 10 times as many people as were murdered. This was perhaps less surprising considering their finding that homicides were given 5,000 newspaper inches to 130 inches for strokes in a 6-month period. They found that this relationship was much the same between deaths by floods and by asthma.

The tendency to overestimate misbehaving amongst peers can spread harm. Students who believe their colleagues drink more than they do are unlikely to reduce their own drinking habits, and feel pressure to conform to a mistaken behavioural norm. Such judgements can spread purely by walking past the student's union and seeing drinking throughout the day. The visual image overrides an assessment of what proportion of the total student population are in the bar. As is often said: 'A picture is worth a thousand words'. This is certainly true when it comes to probability judgements. It is often the case that we follow traditions not because we prefer them to other action, but because we think that other will appreciate it, or like or, or even praised

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

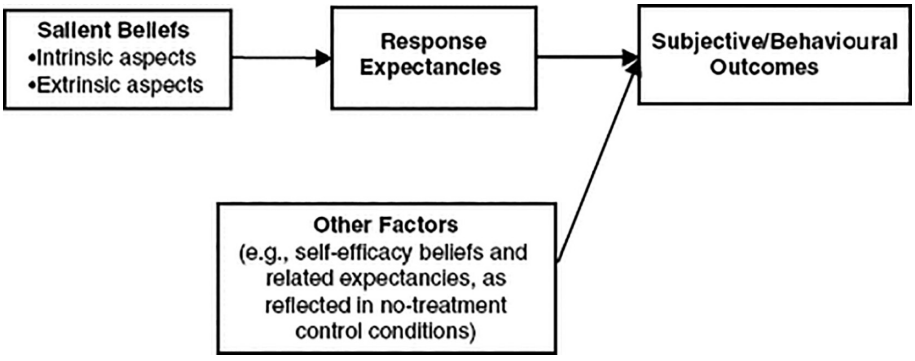

Figure 2.6 Framework for placebo effects
Source: Shiv et al. (2005)

us. Hence, some adolescents who do not wish to start smoking might hold positive views about drugs because they think others hold them too, whether this is true or not. This is often called 'pluralistic ignorance', which gets corrected when shocks happen such as the death of a celebrity from overdose. Such events lead some to express their previously suppressed 'true views', which in turn incentivises others to follow. Another example can be found in how the #MeToo and recent racial inequality and anti-discrimination movements quickly grew as suppressed outrage was released.

2.5.5 Anchoring and Placebo Effects

In evaluating new health information, which carries some uncertainty, we tend to make judgments relative to an available 'anchor'. These allow us to relate to information in a way that brings some sense to a novel situation. We can readily anchor on something known to us. For instance, when deciding whether a virus such as COVID-19 is a high, moderate or low risk, people tend to compare COVID-19 to the risks of more familiar viruses, perhaps seasonal flu, SARS or HIV. Judgements over new information will have different effects on behaviour dependent on the anchor used and the veracity of this anchor as a legitimate guide.

Anchoring may partly explain the 'placebo effect', which sees placebos delivering surprisingly positive effects, whether relieving pain or treating physical and mental illness. Shiv et al. (2005) suggest that we use price as a 'signal' in selecting amongst medicines, and that the placebo effect is stronger for more expensive drugs. When a person receives a treatment, their salient beliefs about the substance activate anticipations of behavioural consequences from such treatments, which in turn lead to the behavioural outcomes or placebo effects (Figure 2.6).

Similarly, it seems that this can also operate in negative ways, known as the 'nocebo effect'. In this situation, people who expect adverse effects from treatments suffer a worsening of their symptoms from treatment or side effects when receiving the placebo in a clinical trial (Planès et al., 2016).

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

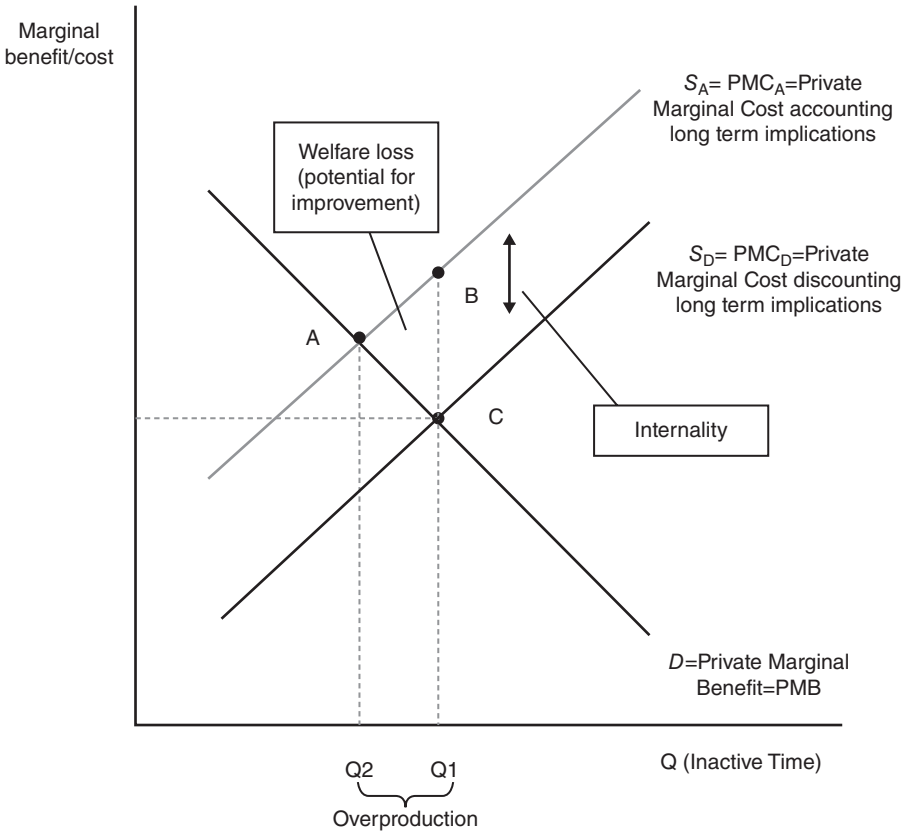

Figure 2.7 Calculating an internality – limited consideration of the long-term costs of inactivity
Source: Correa-Burrows (2014)

2.5.6 We Learn from Today: 'Present Bias'

Individuals often overweight costs and benefits incurred today relative to the costs and benefits incurred in the future. This leads individuals to forgo healthy behaviours in a way that is inconsistent across time. Starting and continuing to exercise or eating healthily will often be felt as a cost in time lost and pleasure foregone today, whatever the potential benefits tomorrow. Present bias can generate 'internalities', as we harm our own future well-being by overweighting today's enjoyment over tomorrow's returns.

Figure 2.7 illustrates the economic consequences when we do not fully account for the long-term costs, we will bear due to physical inactivity, which thus leads us to under-invest time in exercise. When we undervalue the long-term effects of physical inactivity our actual activity is Q_1. Our costs (summarized by curve PMC_D) exceed our benefits (summarised by curve PMB). In this example, the welfare loss is equal to the area ABC. Instead, if we were to fully internalise the long-term consequences of

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

physical inactivity, then we would increase our physical activity to the welfare maximising level, Q_2. At this level, costs accounting for long term effects match gains. Hence, as illustrated, the internality incurred amounts to the additional costs as a result of the long-term consequences from our self-indulgent choices today self that rousts from making a choice now.

2.6 Narratives

Narratives are the simple *stories* that help us relate to issues. They offer simple explanations of events. Narratives are found, for example, in media headlines; whether a negotiation is described as a 'showdown' or simply 'talks' affects our beliefs and expectation of the event. Narratives mix information and emotion. Their importance lies in that they hugely affect the credibility of information. We are attuned to receive and remember stories better than disjointed facts and figures.

Narratives can produce social norms. A single story may prompt someone to avoid eating meat, use a particular hospital or to stop taking a medicine. They invoke the availability heuristic by making a piece of information salient and easy to recall.

Narratives can, of course, be responsible for some information distortion; they feature in some mental health conditions. For example, 'dual narratives' can be part of schizophrenia. Narratives matter because their underlying ideas influence our identities, and the creation of mental models that shape and limit the way we learn. By naming a condition as an 'epidemic' we imply the need for a policy reaction. This is why agencies use such labels. Consider how the naming of a tropical or winter storm enhances the attention we pay to it and levels of fear.

Whilst some narratives are much more effective than others, it would be incorrect that there is always some 'one-size-fits-all' narrative that is universally impactful. Just as with facts and figures, narratives compete for our attention and vary in their success to cut through the noise into our attention. Humans react to individual narratives differently, meaning that they vary in their effect as behavioural incentives. Appreciating the personal role of narratives and how they change can be fundamental to achieving successful behaviour change.

Narratives with impact emerge from many sources. Media stories play an important part in forming and propagating them. Cultural sources, particularly movies, can help narratives establish and shift narratives. Think how smoking was once an integral part of building the character of a movie role, and how this has changed over time. The same was true of its use in advertising imagery, until this was banned. These uses support representation biases. If I wish to show a glamorous or rebellious character, then I either should or should not be seen smoking in order to best fit the stereotype, according to how the dominant narrative of the current time and place is working. Of course, many narratives are consciously generated by opinion leaders in order to support a desired behaviour. Religious texts have always used memorable stories to better convey a behavioural norm.

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

2.6.1 Fake News

Some narratives often have little or no basis in fact. 'Just do it', 'Because you're worth it', 'Adds life', convey emotion not fact. Each narrative, if successful, is intended to benefit its own interest group, whether that is political, religious, a brand or something else. Akerlof and Shiller (2015) describe a *'phishing equilibrium' within* which competitors manipulate narratives with a degree of limited honesty in order to promote a product. Pharmaceutical companies often use vague but very positive lifestyle imagery to advertise their products, rather than focus on specific characteristics. The narrative of alcohol as something to consume for relaxation, or after an intense effort, is prevalent in certain professional groups, and frequently portrayed in movies and popular TV series. The influence on lifestyle routines can be significant, and is often underestimated.

2.6.2 Representative Heuristic

How do we form expectations when the situations we face are unusual to us or outcomes ambiguous, when we cannot rely on past experience and lessons learned? Kahneman and Tversky (1973) argue that we tend grasp at any '*similarity of circumstance',* rather than invest in reflection on the probabilities of gains and losses.

In health, this means that clinicians may fail to diagnose symptoms of heart attack in a young woman, because a young woman is not represented by the stereotype of the typical heart attack patient. If an unusual situation is not representative of a stereotype, then, regardless of probabilities, decisions may be biased as our minds are guided by the representative heuristic.

Similarities of circumstance can also play a part in processes of change. The experience of encountering multiple similar circumstances can generate a tipping point between holding one subjective view and another. Tipping points give rise to the situation in which a 'one-off' event tips the balance in favour of a new point of view, attitude or norm. For example, when several individuals within a person's social environment receive a particular clinical treatment option, this can trigger a perception of it being 'normal'. This new perception then increases the likelihood of that person making the same choice in favour of an option previously seen as novel. The same effect is seen in judgments over different forms of cancer. Hearing a mention of 'breast cancer' perhaps we can bring to mind the individual experience of a personal contact, again affecting decisions taken. Professionals are not immune to these effects, and these may explain some of the variability in clinical practice we will discuss later in Chapter 8.

2.7 Trust

The credibility granted to different information sources may be influenced by who delivers that information, known sometimes as the 'messenger effect' (Martin and

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

Marks, 2019). A study of self-reported compliance with behaviours to limit viral transmission during the COVID-19 pandemic found that the level of personal trust in science had a significant effect on compliance, and much more so than trust in government (Bicchieri et al., 2021). Hence, the credibility of an information sources depends on both the *type* of information being used and the *messenger* used to convey it.

Trust is integral to whether we even consider new information as relevant, appropriate or helpful – its salience to us as individuals. Levels of trust in physicians and insurers are sensitive to the amount of contact a patient has previously had with both, and the extent to which they have choice in physician and insurer. People retrieve information from a variety of differing sources. There is not only a digital divide in access to information, but also a differential access to care and insurance providers, across ethnicity, language, location (particularly between rural and urban areas). These all affect the range of information sources to which we are each exposed, which in turn varies in its salience to us as individuals, whether it is noticed and whether it has any influence on behaviour.

2.7.1 Negativity Bias

In the process of learning from new information, we are affected by a 'negativity bias'. This is an evolutionary tendency to generally pay more attention to negative information, in an endeavour to avoid the risk of losses. It is, of course, an important part of our general loss aversion. Such a bias can explain why information that plays on fears of death during a viral pandemic seems to have longer lasting effects on behaviour than information based on the positive effects of new technology in prevention and treatment. Similarly, they can trigger 'food scares' in the presence of a potential epidemic of avian flu. Negative information looms longer.

2.7.2 Knowledge Resistance

In an ideal world, we would absorb only information with high levels of validity and would dismiss falsehoods or unproven claims. However, the association between knowledge and behaviour is far from straightforward. People exhibit 'knowledge resistance' rejecting information that is both valid and salient to their circumstances (Klintman, 2019). Cultural values, beliefs and identities generate incentives and disincentives for knowledge updating. This phenomenon has a clear evolutionary advantage, as people work hard to adapt to cultural norms to improve their chances of survival and reproduction. Knowledge resistance can strengthen group bonds and enhance collaboration within the group. Klintman (2019) argues that the more the beliefs of a group deviate from that of other groups, the stronger their group cohesion, and the more likely are 'knowledge tribes' to develop.

The problem lies when there is an overlap between knowledge and moral claims. Beliefs on contentious issues such as global warming, or abortion are divisive in

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

society, are mostly a reflection of people's moral values, especially at the extremes of the ideological debate. For insurance, progressives are more likely to neglect findings that challenge organic agriculture.

The importance of moral values lies in how people frame such problems. One typical way to explain knowledge resistance is that individuals make a choice about their 'desired conclusion' and work their arguments out backwards. From research on 'Solution Aversion' Campbell explains that:

> 'people are motivated to deny problems and the scientific evidence supporting the existence of the problems when they are averse to the solutions' (Campbell, 2018)

Knowledge resistance and solution aversion are universal. Humans have not evolved to be knowledge maximisers. It is arguable, similarly, that we have also not evolved to be trust seekers. There is a strong social incentive in the form of status that comes from group collaboration[1] rather than from challenging a group's 'received wisdom'. Anyone challenging group norms faces real risks in doing so, even if they personally dislike a behavioural norm, at least until a tipping point is reached when the gains from change become significant and salient and the norm collapses (Sunstein, 2019).

Knowledge resistance can be explained by loss aversion, not by a failure of intellect. Intelligence, defined as the ability to accomplish complex goals, produces more competent knowledge resisters (Klintman, 2019), which may be contrary to what many would expect. For knowledge resisters, the euphoria of winning an argument can override any potential satisfaction from the learning from argument and evidence. Knowledge resistance mitigates any anxiety around fitting in with group norms. It could be regarded as 'ignorance with a purpose'.

Caplan defines similar behaviour as 'rational irrationality' (Caplan, 2001a), arguing that in some cases not only it might be rational to remain ignorant, but even to adopt beliefs that are not grounded in evidence. Such hypothesis are used to explain voting behaviours (Caplan, 2001b) and the formation of beliefs around novel technologies (Costa-Font et al., 2006).

Klintman (2019) illustrates some other examples of knowledge resistance in the health domain. Indeed, prospective parents prefer often to wait until childbirth to know the gender of their baby rather than discover this in advance. The additional knowledge is knowingly and decisively resisted. The same approach to information is seen in resistance to knowing genetic predisposition for specific diseases. This is 'strategic ignorance'. It is an accepted, and seemingly acceptable, form of knowledge resistance in many cultures. If, however, someone has family history of a disease with a genetic component, then the calculus around genetic testing may change, due to the proximity (or 'availability') of a particular risk.

[1] Although it is common in public debates to accuse each other of being 'knowledge resistant', when deeply defined social interests' conflict with factual knowledge, we tend to choose satisfaction of our social interests.

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press

2.8 Conclusions

Learning is far from the simple and automatic process of information acquisition as described in the models of introductory economics. Information is not equally available to all and is processed in a different way according to the circumstances (including emotions). Accordingly, there is a considerable role for cultural priors, alongside trust and credibility in processing 'objective' information. This is especially the case for health information, which is affected by significant information learning costs.

Learning strategies are at the core of behavioural economics. They reflect the way we form our knowledge, attitudes and behaviours. Behavioural learning considers these social, psychological and developmental processes.

Learning processes are affected by several cognitive biases (framing, priming, social cues, present bias, availability biases). As we progress through this book, we will focus on understanding these biases better and their implications for health. This understanding enables the development of incentives that account for these biases whilst considering long-term ambitions for individual and social well-being. These incentives may harness social norms, facilitate changing norms and use rewards or penalties or other behavioural interventions to counter known biases in the learning process. Recognising and understanding bounded learning creates a wider range of opportunities for intervention within the choice architecture of our lives, to give information greater impact on behaviour.

2.9 Questions to Ponder

1. What examples of knowledge resistance have you encountered?
2. In which aspect of health might less choice be desirable?

https://doi.org/10.1017/9781009168113.003 Published online by Cambridge University Press