

What can Neuroscience Contribute to the Debate Over Nudging?

Gidon Felsen · Peter B. Reiner

Published online: 21 March 2015

© Springer Science+Business Media Dordrecht 2015

Abstract Strategies for improving individual decision making have attracted attention from a range of disciplines. Surprisingly, neuroscience has been largely absent from this conversation, despite the fact that it has recently begun illuminating the neural bases of how and why we make decisions, and is poised for further such advances. Here we address empirical and normative questions about “nudging” through the lens of neuroscience. We suggest that the neuroscience of decision making can provide a framework for understanding how nudges work, and how they can be improved. Towards this end, we first examine how nudges can be incorporated into a leading model of decision making supported by neurobiological data, and use the model to make predictions about the relative effectiveness of different classes of nudges. We then use the model to demonstrate how nudges can both infringe upon and promote autonomy. Finally, we explore the normative implications of the converging consensus from neuroscience and related fields that many everyday decisions are susceptible to covert external influences.

1 Introduction

Recent years have witnessed a vigorous discussion about strategies to improve individual decision making. The primary impetus for such strategies is the recognition that humans exhibit cognitive biases that can lead to suboptimal decisions (Kahneman 2011); by “nudging” in the “correct” direction, choice architects hope to improve decision outcomes (Thaler and Sunstein 2008). Interest in nudges has come primarily from the disciplines of psychology, economics, and philosophy, with substantial advances in public policy accruing

G. Felsen (✉)

Department of Physiology and Biophysics, University of Colorado School of Medicine, 12800 E. 19th Ave., Mail Stop 8307, Aurora, CO 80045, USA
e-mail: gidon.felsen@ucdenver.edu

G. Felsen

Center for Bioethics and Humanities, University of Colorado, 13080 E. 19th Ave., Mail Stop B137, Aurora, CO 80045, USA

G. Felsen · P. B. Reiner

National Core for Neuroethics, University of British Columbia, 2211 Westbrook Mall, Vancouver, BC V6T 2B5, Canada

(Sunstein 2014). Nonetheless, the perspective of neuroscience – which most directly examines *how* the brain performs its manifold functions – has largely been absent from the conversation, despite significant advances in our understanding of the neural bases of decision making (Levy and Glimcher 2012; Shadlen and Kiani 2013). In this paper, we suggest that neuroscience can contribute to both empirical and normative questions about nudging: Specifically, how to make nudges more effective, and how to understand whether they are consistent with our ethical principles.

In considering how to design nudges that lead to better decision outcomes, it is tempting to think that *what* choices are made under various conditions, and how cognitive biases influence these choices, are all that are important. These factors are largely the domain of psychology and behavioral economics. What value, then, can be added by understanding *how* decisions are made by the nervous system, as illuminated by neuroscience?

We address this question with an analogy to drug discovery. One approach to designing better drugs is to test a large set of randomly-selected compounds and observe whether any have the desired effects. This brute-force approach is, strictly speaking, sufficient: assuming the effects of the compounds are observable, the optimal drug can, in principle, be identified. However, understanding the biochemical mechanisms by which drugs act, instead of simply their overt effects, allows for rational drug design, not only increasing the efficiency of the process to identify the optimal compound but also, ideally, reducing the likelihood of side effects. In a similar way, psychology and behavioral economics are, strictly speaking, sufficient for determining which nudges most effectively promote the desired choices. Neuroscience, however, may allow for the “rational design” of nudges by providing insight into the mechanisms by which decisions are made, providing a biologically-grounded practical starting point for forming testable hypotheses about which nudges may work best. Determining the efficacy of nudges would still require field testing, but understanding the underlying neuroscience gives the choice architect both insights into, and new tools to manipulate, the decision-making process. Towards this end, in Section 2 we examine how nudges work in light of data and theories from neuroscience.

Can such insight into how nudges work inform the fiercely debated normative question of whether the use of nudges is ethical (Blumenthal-Barby 2013; Bovens 2009; Grüne-Yanoff 2012; Mitchell 2004; Sunstein and Thaler 2003; White 2013)? One position is that neuroscience can have nothing to say about the ethics of nudging, just as it may not say anything about any normative question: With respect to brain function, neuroscience describes what “is,” which may not be relevant to what we “ought” to do (Berker 2009; Kamm 2009). This position has some validity: For example, we may learn that certain information is processed via mechanism A and not mechanism B. If A and B are arbitrary mechanisms there is no a priori reason to think that the data are normatively significant. However, suppose that a particular capacity C depends on information being processed by mechanism B. Now, our finding that information is not processed by mechanism B *does* have implications—specifically, that it would be misplaced to assign much value to capacity C. In other words, neuroscience may demonstrate that some of our assumptions about the capacities of the brain are incorrect, prompting us to reconsider how they are valued (Greene 2003, 2014; Singer 2005). In Section 3, we apply this idea to the question of whether the effect of nudges on autonomy is normatively significant.

We conclude by discussing the prospects for future neuroscientific research, and interactions between neuroscience and other disciplines, to contribute to these empirical and normative questions.

2 Empirical Questions About Nudging

Choice architects wish to optimize nudges in order to maximize their effectiveness. The standard empirical approach to this problem is to compare the effectiveness of different forms of nudges in particular contexts (Downs et al. 2009; Sinclair et al. 2014; Sunstein 2011). We suggest that there is merit in considering a complementary approach that explores the question of *how* nudges exert their effects. Such a mechanistic perspective would naturally consider findings from neuroscience, a field that has increasingly provided insight into how decisions are made.

Ideally, these findings would derive from humans making exactly the sorts of decisions that choice architects wish to nudge. In the real world, however, experimental approaches in human subjects are limited: current means of recording (e.g., functional neuroimaging) and manipulating (e.g., transcranial stimulation) normal human brain activity offer insufficient spatiotemporal resolution. On the other hand, studies of decision making in animal models – including non-human primates (Hikosaka et al. 2014), rodents (Carandini and Churchland 2013) and even fruit flies (DasGupta et al. 2014) – offer deep insight into the neural mechanisms of decisions at the level of neurons and synapses. Moreover, modern technology allows not only direct recording of neuronal activity during decision making, but, once a model is developed, exceedingly subtle manipulation of the circuitry can be used to rigorously test hypotheses. The decisions that can be studied in animal models are often much simpler than those that choice architects may have in their sights. However, by exploiting the known similarities in structure and function across species, data from animals making simple decisions can inform strategies for influencing human decisions, in much the same way that animals can serve as simplified, yet indispensable, models for studying human disease (Jucker 2010). Our main focus in this paper is therefore on the data typically obtained from animal models and, in particular, the theories of decision making that they support.

One set of models in particular offers insights that may be useful to the choice architect. These “diffusion-to-bound” models originated in the psychological literature in order to explain observable behavioral data in subjects performing a range of tasks requiring a perceptual decision to be reported with a motor output (e.g., a button press or an eye movement; Ratcliff and Rouder 1998; Smith and Ratcliff 2004). The models have found substantial support from studies in which neuronal activity has been recorded or manipulated in animal subjects performing simple decision-making tasks (Carpenter and Williams 1995; Gold and Shadlen 2007). Diffusion-to-bound models provide a valuable framework for thinking about how nudges may affect decisions.

Decisions are influenced by multiple factors, including current sensory evidence about important features of the environment, stored memory about past experience, and the subjective value of each option. The general idea of diffusion-to-bound models is that these influences are integrated into a unified “decision variable” (DV) that exists in a space bounded by the available choices (Fig. 1a). The DV “diffuses” stochastically within this space, closer or further from each bound depending on the strengths of the

various influences, until reaching one of the bounds, at which point the corresponding choice is made.¹ This model captures the probabilistic nature of decision making, in the sense that factors influencing the DV may make one choice more likely but do not entirely determine the choice. For example, imagine you are in a cafeteria, deciding between a salad and a brownie. The aroma of the brownie and the short-term pleasure you will experience upon eating it – with which you are familiar from previous experience – drive the DV towards the “brownie bound.” However, your long-term goal of maintaining overall health tugs the DV back towards the “salad bound.” The relative strength of these competing influences ultimately results in the DV reaching one of the bounds, at which point the choice is made.

We briefly describe the experimental evidence that supports one of the most commonly used diffusion-to-bound models of decision making for which there exists an exceedingly strong body of empirical data (Shadlen and Kiani 2013). In this paradigm, a non-human primate is presented with a visual stimulus consisting of moving dots (Newsome et al. 1989). Some percentage of the dots move in the same direction (“coherently”), and the primate must decide whether this coherent group of dots is moving to the left or the right and indicate its decision by executing a learned motor output; usually looking in the direction of coherent motion. Thus, in this context, the choice bounds represent the choices “look left” and “look right.” Correct choices are reinforced with a juice reward. The experimenter has control over all of the parameters that are integrated by the DV in this task: the percentage of coherently-moving dots, the probability that the direction of coherent motion on any given trial will be left or right, and how much reward will be provided for a correct choice.

The power of the model is revealed by experiments in which the activity of single neurons in specific brain regions are monitored while the animal performs the task. Using this approach, it has been shown that the rate at which individual neurons in the lateral intraparietal region of cortex fire action potentials – the currency for computing and communicating information in the nervous system – reflects the value of the DV. The firing rate during each trial increases faster on high-coherence than low-coherence trials, correlates with the primate’s left vs. right choice, and reaches the same level at the time the choice is made—exactly as would be predicted by the model (Mazurek et al. 2003; Roitman and Shadlen 2002). Other studies have shown that stimulating neurons in the same region predictably biases the decision and changes how quickly it is made (Hanks et al. 2006), demonstrating a causal relationship between neuronal activity and the DV.

This model is entirely consistent with the idea that nudges make a particular choice more likely without restricting the range of possible choices (Thaler and Sunstein 2008). In this context, a nudge can be considered as a shift in the starting position of the DV towards the bound corresponding to the preferred choice (i.e., the “preferred bound;” Fig. 1b).² All other influences on the DV being equal, the nudge makes the

¹ For simplicity of description and display we consider only two possible options and assume that one must be chosen; key features of the model can be extended to explain decisions among multiple alternatives (Bogacz et al. 2007).

² A nudge is one example of an external influence on decisions, other forms of which (e.g., recent trial history) have been shown to result in a similar shift in the starting position of the DV (Bode et al. 2012). Equivalently, any such influence (including nudges) can be thought to decrease the distance to the preferred bound or increase the distance to the “non-preferred bound.” As discussed below, nudges can also be modeled as a change in the drift rate of the DV towards the preferred bound.

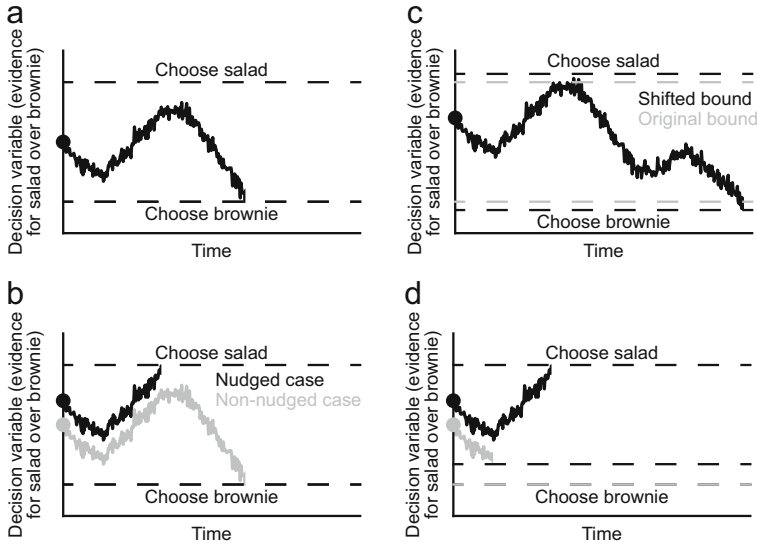


Fig. 1 Nudges considered within the diffusion-to-bound model of decision making. **a.** From its starting position (circle) the decision variable (DV) integrates, across time, the influences for one choice over the other. When the DV reaches one of the choice bounds (dashed lines), the process stops and the corresponding option (here, the brownie) is chosen. **b.** Nudges can be conceptualized as a shift in the starting position of the DV towards the “preferred choice” (here, the salad), making this choice more likely to be selected. **c.** Nudges do not preclude either choice: Shifting the choice bounds further from the starting position of the DV minimizes the effect of the nudge. **d.** Some nudges can be seen as debiasing. An internally-generated bias towards the brownie can be represented as a shift in the “choose brownie” bound closer to the starting position of the DV (shifted bound, black; original bound, gray). The DV in the non-nudged case (gray) would reach the brownie bound, but in the nudged case (black) would reach the salad bound

preferred choice more likely. However, the nudge does not preclude the possibility of the DV reaching the non-preferred bound. Indeed, the bounds themselves can be shifted by top-down cortical control (Mulder et al. 2012), allowing the decision maker to mitigate the “advantage” conferred on the preferred choice by the nudge [Fig. 1c; cf. “policy setting” (Shadlen and Roskies 2012).]

For example, imagine you are back in the cafeteria and now a choice architect has conveniently placed the salad on the way to the cashier and the brownies in the far corner. The cost of the extra effort required to get to the brownies decreases their subjective value. This change is reflected in a shift in the starting position of the DV away from the “choose brownie” bound. All other factors being equal, the DV is now more likely to reach the “choose salad” bound before it reaches the “choose brownie” bound. Yet, by requiring that the DV reach a larger value before making a choice (via top-down control), the brain can still overcome the nudge and select the brownie. However, if you are engaged in a conversation with your co-worker while making this decision, your brain may be less likely to shift the bounds sufficiently to oppose the nudge, and therefore may be more likely to choose the salad.

This framework accounts for the effect of influences on decisions in a manner that is agnostic with respect to their source. In the same way that it explains the effects of external influences such as nudges, it can explain the effects of internal influences (Bode et al. 2014) such as the cognitive biases that nudges are designed to counteract. For example, steep discounting of future rewards may decrease the distance to the

bound corresponding to the option associated with higher near-term value (i.e., the brownie). Such delay discounting can thus be thought of as a source of bias. Considered in this light, the diffusion-to-bound framework lends credence to the idea that in some cases nudges can be seen to “debias,” rather than bias, decisions (Jolls and Sunstein 2006; Fig. 1d).

The framework also allows us to make testable predictions about the relative effectiveness of nudges. First, nudges should be more effective at influencing decisions that are made with less top-down, conscious control. Without such conscious reflection, choice bounds cannot be moved to counteract the nudge. This suggests that nudges that engage processes that are less under conscious control should be more effective than those that engage processes more under conscious control. Findings from some studies are consistent with this prediction (Wisdom et al. 2010), but to our knowledge this prediction has not been directly tested.

Second, although we have focused our discussion on the effects of nudges that shift the starting position of the DV, nudges that act upon dynamic sensory evidence may be more effective than those that act upon static factors (e.g., the relative subjective value of the options). Dynamic sensory evidence affects the *rate* at which the DV drifts toward the preferred bound (represented by the average slope of the DV path), while static factors decrease the distance from the starting position to the bound (represented by an intercept). Therefore, over time, the drift rate could have a larger effect on the DV than would a change in its starting position.³ Returning to the cafeteria, mitigating the aroma of the brownie – a sensory cue that can strongly influence the DV over time – could, over sufficient time, have a stronger effect than placing the brownies in a modestly inconvenient location.

There are of course some limitations of applying this framework – or any model – to nudges. Most obviously, while the framework captures some of the psychological and neurobiological processes underlying simple sensory-motor decisions, it is unknown whether these are sufficiently similar to the more complex decisions that the choice architect might want to nudge (Roskies 2010). We also note that we have not discussed several other neuroscience-inspired models, nor a wealth of data on the neural bases of (internal) cognitive and (external) contextual biases in human studies (De Martino et al. 2006; Deppe et al. 2005; Gilaie-Dotan et al. 2014; Murawski et al. 2012; Tom et al. 2007; Yu et al. 2010) that may also be relevant to nudges. Nevertheless, we hope to have demonstrated how the rational design of effective nudges can be informed by neuroscience.

3 Normative Questions About Nudging

Despite the “libertarian” attempt to preserve the full range of options, one of the most prominent arguments against nudging is that it disrespects autonomy (Blumenthal-Barby 2013; Bovens 2009; Grüne-Yanoff 2012; White 2013). Given the respect for

³ This prediction has been tested in the framework of the neurobiological experiments described above: Stimulating neurons in the medial temporal lobe, which represent stimulus motion and therefore increases the drift rate of the DV, has an even larger effect on choice than stimulating lateral intraparietal neurons, which adds an offset to the DV (Ditterich et al. 2003; Hanks et al. 2006).

autonomy prevalent in modern Western society, this allegation is serious. We suggest that there are several ways in which neuroscience can offer insight into whether the effects of nudges on autonomy makes them ethically questionable.

According to a consensus conception of autonomy, factors influencing the autonomy of a decision include whether it is rational, free from undue external influences, and consistent with the agent's higher-order desires – i.e., those corresponding to fundamental goals, as opposed to lower-order desires corresponding to physiological needs (Dworkin 1988; Felsen and Reiner 2011; Frankfurt 1971; Hill 1989). The neuroscience-supported diffusion-to-bound model described above provides a framework within which to examine how nudges may affect each of these factors, and therefore autonomy. Note that we make no attempt to claim that a particular decision made under the influence of a particular nudge is categorically “autonomous” or “not autonomous;” rather, we consider autonomy as a graded phenomenon, and we seek to determine whether a particular nudge increases or decreases autonomy.

Imagine a covert nudge that affects a decision in which the two options are otherwise equally likely to be selected – i.e., the bounds are equidistant from the non-nudged starting position of the DV (as shown in Fig. 1b). Returning once again to the cafeteria, placing the brownies in a less convenient location than the salad is likely to exert such a covert influence.⁴ This nudge would violate the condition that the decision be free from undue influence, and thus decrease autonomy. This is captured by the model: Even a small external influence, such as a “gentle” nudge, may have a significant effect on a decision between equally-valued options (Bode et al. 2014), which would be difficult to counteract. Indeed, covert priming has been shown to influence free choices between equally-valued options, and can be accounted for by the diffusion-to-bound model (Mattler and Palmer 2012; Schlaghecken and Eimer 2004). Note, however, that if the influence were instead overt – e.g., displaying nutritional information for each item in the cafeteria – the agent could counteract it by shifting one or both choice bounds further from the starting point of the DV via top-down control (as in Fig. 1c).

Another case involves a nudge deployed to counteract an internally-generated bias (Fig. 1d), arguably the most common situation for which nudges were designed. If by counteracting a bias towards one choice (e.g., the brownie) the nudge promotes the choice (e.g., the salad) that is in line with higher-order desires (e.g., to maintain overall health), it may be seen to increase the autonomy of the decision. We suggest that this is true even for covert nudges, since the agent would likely endorse the influence if she were aware of it, since it promotes the attainment of her higher-order desires.

Note that even if two agents share the same higher-order desire – e.g., maintaining overall health – the same nudge could increase autonomy for one (as described above) but decrease it for the other. Imagine, following White (2013), that for this second agent, this cafeteria's brownies evoke pleasant memories of past brownies shared with his late grandfather, promoting overall well-being and health. Here, a nudge towards the salad decreases the likelihood that the agent will make the choice most consistent with his higher-order desires. The dilemma for choice architects is that they have only limited information about each agent's choice bounds and their relative value to their personal life plan, and therefore nudges invariably decrease autonomy for some even if

⁴ While a nudge may not be exclusively covert or overt, it may still exert its influence more covertly than overtly, or vice versa.

they succeed in respecting autonomy for most (White 2013). Proponents of nudging suggest that this issue may, in the future, be addressable via “personalized” choice architecture (Sunstein 2014), but at present it remains a practical concern.

Another way in which neuroscience can inform the question of whether the effect of nudges on autonomy makes them less ethical is to examine the neural bases of the conditions necessary for autonomy themselves. In doing so, Felsen and Reiner (2011) noted that the neuroscientific evidence about how everyday decisions are made is inconsistent with the idea that they are free from undue external influence; in fact, external influences are the norm rather than the exception (Schüür and Haggard 2011). These influences are often incorporated into our decision processes covertly; that is, once they become a factor in a particular decision, the source of the influence is discarded. For example, functional imaging experiments have suggested that options framed as losses elicit different choices than do identical options framed as gains because the former elicit covert emotional responses (De Martino et al. 2006). Such effects are ubiquitous (Kahneman 2011), likely reflecting an evolutionarily advantageous feature of brain function, and call into question whether goal-directed decisions can be fully free of external influences (Bode et al. 2014; Custers and Aarts 2010). Therefore, neuroscience suggests that the degree to which our everyday decisions are autonomous – according to the consensus conception employed here – is limited.

Does this analysis have normative implications? As described in Section 1, if it is true that autonomy depends on the capacity to make decisions free from covert external influences, and the evidence from neuroscience supports only a diminished form of this capacity, then it is worth considering whether the high value with which we regard autonomy may be misplaced. While it of course does not follow from this line of reasoning that nudges are ethical, this analysis counters the common argument that nudges are unethical *because* they violate autonomy. We discuss other ways in which neuroscience can inform the ethics of nudges below.

4 Conclusions and Future Directions

We have described how our current understanding of the neural bases of decision making can inform both empirical and normative questions about nudging. Indeed, insights into the underlying neuroscience provide a useful framework for considering the relative effectiveness of nudges (Section 2), as well as for illuminating the normative implications of nudges on autonomy (Section 3).

The insights at which we have arrived here are likely to be supplanted by future studies, as the neuroscience of decision making continues to advance, driven by technological and methodological innovation (Insel et al. 2013). For example, there exists a consequentialist argument suggesting that agents will become habituated to rely upon nudges rather than reflect upon their choices (White 2013; Zhang et al. 2011) and therefore fail to make optimal decisions in the absence of nudges, ultimately leading to poorer overall decision outcomes. The neural mechanisms that underpin habit formation and even sober reflection are areas of intense neuroscientific study (Badre 2008; Graybiel 2008) that may allow for the design of nudges less susceptible to this problem. These and other advances in neuroscience may have a substantive impact upon both the design of nudges and their normative implications. Given the potential for new

neuroscientific data to inform policies intended to improve decision making, we call for frequent and vigorous communication between the disciplines of neuroscience and public policy. While the field of behavioral law and economics has integrated behavioral psychology into legal and economic theories (Camerer et al. 2003; Jolls et al. 1998), incorporating the related neuroscientific data, with an eye towards their application to public policy, is becoming increasingly important.

Acknowledgments We thank the editors and two anonymous reviewers for their constructive suggestions. This work was supported by the Greenwall Foundation's Faculty Scholars Program in Bioethics (G. F.).

References

- Badre, D. 2008. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences* 12(5): 193–200.
- Berker, S. 2009. The normative insignificance of neuroscience. *Philosophy and Public Affairs* 37(4): 293–329.
- Blumenthal-Barby, J.S. 2013. Choice architecture: improving choice while preserving liberty? In *Paternalism: theory and practice*, ed. C. Coons and M. Weber, 178–196. Cambridge: Cambridge University Press.
- Bode, S., C. Murawski, C.S. Soon, P. Bode, J. Stahl, and P.L. Smith. 2014. Demystifying “free will”: the role of contextual information and evidence accumulation for predictive brain activity. *Neuroscience & Biobehavioral Reviews* 47: 636–645.
- Bode, S., D.K. Sewell, S. Lilburn, J.D. Forte, P.L. Smith, and J. Stahl. 2012. Predicting perceptual decision biases from early brain activity. *The Journal of Neuroscience* 32(36): 12488–12498.
- Bogacz, R., M. Usher, J. Zhang, and J.L. McClelland. 2007. Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 362(1485): 1655–1670.
- Bovens, L. 2009. The ethics of nudge. In *Preference change: Approaches from Philosophy, Economics and Psychology*, ed. T. Grüne-Yanoff and S.O. Hansson, 207–219. Dordrecht: Springer.
- Camerer, C., S. Issacharoff, G. Loewenstein, T. O'Donoghue, and M. Rabin. 2003. Regulation for conservatives: behavioral economics and the case for “asymmetric paternalism.”. *University of Pennsylvania Law Review* 151(3): 1211–1254.
- Carandini, M., and A.K. Churchland. 2013. Probing perceptual decisions in rodents. *Nature Neuroscience* 16(7): 824–831.
- Carpenter, R.H.S., and M.L.L. Williams. 1995. Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377(6544): 59–62.
- Custers, R., and H. Aarts. 2010. The unconscious will: How the pursuit of goals operates outside of conscious awareness. *Science* 329(5987): 47–50.
- DasGupta, S., C.H. Ferreira, and G. Miesenböck. 2014. FoxP influences the speed and accuracy of a perceptual decision in *Drosophila*. *Science* 344(6186): 901–904.
- De Martino, B., D. Kumaran, B. Seymour, and R.J. Dolan. 2006. Frames, biases, and rational decision-making in the human brain. *Science* 313(5787): 684–687.
- Deppe, M., W. Schwindt, H. Kugel, H. Plassmann, and P. Kenning. 2005. Nonlinear responses within the medial prefrontal cortex reveal when specific implicit information influences economic decision making. *Journal of Neuroimaging* 15(2): 171–182.
- Ditterich, J., M.E. Mazurek, and M.N. Shadlen. 2003. Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience* 6(8): 891–898.
- Downs, J.S., G. Loewenstein, and J. Wisdom. 2009. Strategies for promoting healthier food choices. *American Economic Review* 99(2): 159–164.
- Dworkin, G. 1988. *The theory and practice of autonomy*. Cambridge: Cambridge University Press.
- Felsen, G., and P.B. Reiner. 2011. How the neuroscience of decision making informs our conception of autonomy. *American Journal of Bioethics Neuroscience* 2(3): 3–14.
- Frankfurt, H. 1971. Freedom of the will and the concept of a person. *The Journal of Philosophy* 68: 5–20.
- Gilaie-Dotan, S., A. Tymula, N. Cooper, J.W. Kable, P.W. Glimcher, and I. Levy. 2014. Neuroanatomy predicts individual risk attitudes. *The Journal of Neuroscience* 34(37): 12394–12401.
- Gold, J.I., and M.N. Shadlen. 2007. The neural basis of decision making. *Annual Review of Neuroscience* 30: 535–574.

- Graybiel, A.M. 2008. Habits, rituals, and the evaluative brain. *Annual Review of Neuroscience* 31(1): 359–387.
- Greene, J. 2003. From neural “is” to moral “ought”: what are the moral implications of neuroscientific moral psychology? *Nature Reviews Neuroscience* 4(10): 846–850.
- Greene, J.D. 2014. Beyond point-and-shoot morality: Why cognitive (neuro)science matters for ethics. *Ethics* 124(4): 695–726.
- Grüne-Yanoff, T. 2012. Old wine in new casks: libertarian paternalism still violates liberal principles. *Social Choice and Welfare* 38(4): 635–645.
- Hanks, T.D., J. Ditterich, and M.N. Shadlen. 2006. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience* 9(5): 682–689.
- Hikosaka, O., H.F. Kim, M. Yasuda, and S. Yamamoto. 2014. Basal ganglia circuits for reward value-guided behavior. *Annual Review of Neuroscience* 37(1): 289–306.
- Hill, T. 1989. The Kantian conception of autonomy. In *The inner citadel: essays on individual autonomy*, ed. J.P. Christman, 91–108. New York: Oxford University Press.
- Insel, T.R., S.C. Landis, and F.S. Collins. 2013. The NIH BRAIN initiative. *Science* 340(6133): 687–688.
- Jolls, C., and C.R. Sunstein. 2006. Debiasing through law. *The Journal of Legal Studies* 35(1): 199–242.
- Jolls, C., C.R. Sunstein, and R. Thaler. 1998. A behavioral approach to law and economics. *Stanford Law Review* 50(5): 1471–1550.
- Jucker, M. 2010. The benefits and limitations of animal models for translational research in neurodegenerative diseases. *Nature Medicine* 16(11): 1210–1214.
- Kahneman, D. 2011. *Thinking, fast and slow*. New York: Farrar, Straus, and Giroux.
- Kamm, F.M. 2009. Neuroscience and moral reasoning: a note on recent research. *Philosophy and Public Affairs* 37(4): 330–345.
- Levy, D.J., and P.W. Glimcher. 2012. The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology* 22(6): 1027–1038.
- Mattler, U., and S. Palmer. 2012. Time course of free-choice priming effects explained by a simple accumulator model. *Cognition* 123(3): 347–360.
- Mazurek, M.E., J.D. Roitman, J. Ditterich, and M.N. Shadlen. 2003. A role for neural integrators in perceptual decision making. *Cerebral Cortex* 13(11): 1257–1269.
- Mitchell, G. 2004. Libertarian paternalism is an oxymoron. *Northwestern University Law Review* 99: 1245–1277.
- Mulder, M.J., E.-J. Wagenmakers, R. Ratcliff, W. Boekel, and B.U. Forstmann. 2012. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience* 32(7): 2335–2343.
- Murawski, C., P.G. Harris, S. Bode, D.J.F. Domínguez, and G.F. Egan. 2012. Led into temptation? rewarding brand logos bias the neural encoding of incidental economic decisions. *PLOS ONE* 7(3): e34155.
- Newsome, W.T., K.H. Britten, and J.A. Movshon. 1989. Neuronal correlates of a perceptual decision. *Nature* 341(6237): 52–54.
- Ratcliff, R., and J.N. Rouder. 1998. Modeling response times for two-choice decisions. *Psychological Science* 9(5): 347–356.
- Roitman, J.D., and M.N. Shadlen. 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience* 22(21): 9475–9489.
- Roskies, A.L. 2010. How does neuroscience affect our conception of volition? *Annual Review of Neuroscience* 33(1): 109–130.
- Schlaghecken, F., and M. Eimer. 2004. Masked prime stimuli can bias ‘free’ choices between response alternatives. *Psychonomic Bulletin & Review* 11(3): 463–468.
- Schüür, F., and P. Haggard. 2011. What are self-generated actions? *Consciousness and Cognition* 20(4): 1697–1704.
- Shadlen, M.N., and R. Kiani. 2013. Decision making as a window on cognition. *Neuron* 80(3): 791–806.
- Shadlen, M.N., and A.L. Roskies. 2012. The neurobiology of decision-making and responsibility: reconciling mechanism and mindedness. *Frontiers in Decision Neuroscience* 6: 56.
- Sinclair, S.E., M. Cooper, and E.D. Mansfield. 2014. The influence of menu labeling on calories selected or consumed: a systematic review and meta-analysis. *Journal of the Academy of Nutrition and Dietetics* 114(9): 1375–1388.
- Singer, P. 2005. Ethics and intuitions. *The Journal of Ethics* 9(3–4): 331–352.
- Smith, P.L., and R. Ratcliff. 2004. Psychology and neurobiology of simple decisions. *Trends in Neurosciences* 27(3): 161–168.
- Sunstein, C.R. 2011. Empirically informed regulation. *The University of Chicago Law Review* 78(4): 1349–1429.
- Sunstein, C.R. 2014. *Why nudge?: the politics of libertarian paternalism*. New Haven: Yale University Press.

- Sunstein, C.R., and R.H. Thaler. 2003. Libertarian paternalism is not an oxymoron. *University of Chicago Law Review* 70: 1159–1202.
- Thaler, R.H., and C.R. Sunstein. 2008. *Nudge: improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.
- Tom, S.M., C.R. Fox, C. Trepel, and R.A. Poldrack. 2007. The neural basis of loss aversion in decision-making under risk. *Science* 315(5811): 515–518.
- White, M.D. 2013. *The manipulation of choice: ethics and libertarian paternalism*. New York: Palgrave Macmillan.
- Wisdom, J., J.S. Downs, and G. Loewenstein. 2010. Promoting healthy choices: information versus convenience. *American Economic Journal: Applied Economics* 2(2): 164–178.
- Yu, R., D. Mobbs, B. Seymour, and A.J. Calder. 2010. Insula and striatum mediate the default bias. *The Journal of Neuroscience* 30(44): 14702–14707.
- Zhang, Y., J. Xu, Zixi Jiang, and S.C. Huang. 2011. Been there, done that: the impact of effort investment on goal value and consumer motivation. *Journal of Consumer Research* 38(1): 78–93.