

6

Informed Consent, Autonomy, and Beliefs

Over the course of the next three chapters, I shall explicate the implications that my rationalist account of autonomy has for informed consent. Informed consent requirements are ubiquitous in health care, and they are regarded as a cornerstone of ethical medical practice. It is also often treated as a truism that these requirements are to be justified by the principle of respect for autonomy. However, whilst this view is still widely accepted, it has recently been brought into question, with critics suggesting that informed consent requirements are neither necessary nor sufficient for safeguarding individual autonomy in the biomedical sphere.

In the first part of this chapter, I shall suggest that this objection is misplaced, though I shall claim that it suggests that we should revise our understanding of what informed consent requires. In doing so, I shall extend my previous discussion of the relationship between beliefs and autonomous decision-making. This will provide the foundation for the final part of the chapter in which I shall outline a rationalist test of materiality. I suggest that this ought to undergird an adequate standard of information disclosure for informed consent, drawing on the recent hybrid account evidenced in the 2015 Montgomery judgement in the UK.

1. The Structure, Definition, and Limits of Informed Consent

As Gerald Dworkin points out, the doctrine of informed consent is a ‘creature of law’;¹ it has been developed in various legal domains in which one party sanctions another to perform ‘... some course of action to which the consented to party would otherwise have no moral right’.² The fundamental idea that we aim to capture when we claim that ‘A morally ought to obtain B’s informed consent to A’s doing x to B’, is that the moral permissibility of A’s doing x to B, is at least partly dependent on the following conditions being met:

- (i) B must be sufficiently informed with regards to the relevant facts concerning x to understand what x is (and what consequences are likely to occur as a result of x).

¹ Dworkin, *The Theory and Practice of Autonomy*, 101.

² Kleinig, ‘The Nature of Consent’, 8. See also Miller and Wertheimer, *The Ethics of Consent*, for examples of the domains in which informed consent can be invoked.

- (ii) On the basis of this information, *B herself* makes the decision to allow *A* to do *x*.³

One reason that the moral permissibility of *A*'s doing *x* to *B* can be at least partly dependent on whether *B* has provided informed consent to *x*, is that *B* may bear a right against *A* performing an interference of the sort involved in *x*. For example, *B* could have a right to bodily integrity that affords her protection against *A* performing an injection on her. However, many rights that incorporate these sorts of claims of protection also incorporate a second-order power to waive the claim in question.⁴ To authorize a medical treatment in accordance with the requirements of informed consent is to waive the rights of protection that might otherwise preclude the moral permissibility of the intervention. In providing informed consent to an injection, one waives one's claim right against the bodily interference that the injection involves. However, when *A* does *x* without *B*'s informed consent, *B*'s extant claim rights have not been waived, and are still 'in play'. In such cases, if *A* does *x*, *A* will have infringed (and perhaps violated) *B*'s right, and failed in her own duty to refrain from doing *x* to *B* (in the absence of *B*'s waiving that claim). We may also note that in some contexts, *A* may have a *positive* obligation to facilitate *B*'s ability to make an autonomous decision about whether to consent to *x*.

For reasons I shall further explain in my discussion of the value of autonomy in Chapter 9, we should not understand the requirements of informed consent to generate or impose a positive obligation to *provide x* to *B*. Consent can involve the waiver of moral protections that would otherwise render a medical intervention impermissible, but there may yet be other moral reasons that can outweigh the autonomy-based reasons we might have to provide *x* to *B*. This is clearest in cases where doing *x* to *B* will have harmful implications for others, or when *x* cannot be provided within a just framework of resource allocation.

However, harm to others is arguably not necessary for the moral or legal impermissibility of doing certain things to *B* to which she has consented. Indeed, in the legal context, even when a person provides valid consent, anything that causes that person actual bodily harm constitutes a criminal offence '... unless it can be shown that the act falls into one of the exceptional circumstances in which consent can provide a defence'.⁵ Naturally, 'proper medical treatment'⁶ is one such kind of act; however, consent may not provide a legal defence (and perhaps not a moral defence) of certain kinds of harmful action, such as acts of violence within a sado-masochistic context,⁷ killing, or medical procedures that fall outside the boundaries of 'proper medical treatment'.⁸ In adopting this view, the law appears to accept that we either (i) hold certain claim rights that do not incorporate the corresponding power to waive

³ There is a considerable debate about whether consent involves only a psychological state, or whether it also requires a behavioural expression. I lack the space to engage with this debate here, but note that medical law on informed consent and mental capacity implicitly reflects a behavioural understanding of consent. For discussion, see Hurd, 'The Moral Magic of Consent'; Kleinig, 'The Nature of Consent'.

⁴ Wenar, 'The Nature of Rights'.

⁵ Herring, *Medical Law and Ethics*, 150.

⁶ For definition of this, see *ibid*.

⁷ Athanassoulis, 'The Role of Consent in Sado-Masochistic Practices'.

⁸ Herring, *Medical Law and Ethics*, 204–5.

the claim in question, or (ii) that we have overriding reasons in non-ideal contexts to avoid false positive assessments of decision-making capacity to waive claims against particularly significant harms.⁹ However, as I shall suggest in Chapters 7 and 9, there are philosophical grounds for objecting to this understanding of the claim rights we enjoy.

Accordingly, *B* may not have a powerful positive right to medical intervention *x* in many cases, but she may hold a number of negative claim rights that generate a powerful duty for *A* to refrain from doing *x* to *B*, in the absence of *B*'s consent.¹⁰ Since rights are typically understood to provide trumping¹¹ or exclusionary moral reasons,¹² it will typically take far stronger moral reasons (corresponding to competing rights) to justify overriding the negative obligation to refrain from performing non-consensual treatment on an individual, than it takes to justify overriding the moral reasons we have to provide treatment to which they have consented.¹³ That said, even this negative obligation to refrain from non-consensual treatment can plausibly be overridden in certain contexts. For instance, in public health, there may be cases in which this negative obligation can be overridden by the competing rights of others.

I shall consider the strength of these obligations in greater detail in Chapter 9. Here though, I want to turn to a second aspect of informed consent. The moral and legal requirement to obtain *A*'s consent does not apply if *A* lacks the capacities that are necessary for providing valid consent.¹⁴ To give a clear example, there cannot be any meaningful requirement to obtain informed consent to a medical intervention from a patient who is in a coma. Instead, the decision about whether to perform the intervention must instead be guided by alternative approaches, drawing on advance directives and proxy decision-makers (where applicable), and/or an assessment of what is in the patient's best interests.¹⁵ Accordingly, our understanding of the conditions of decision-making capacity will have a considerable bearing on the scope of the requirement to obtain informed consent and to respect individual treatment decisions. I shall consider this aspect of informed consent in the next chapter.

With these limitations in mind, let us reconsider the constitutive conditions of informed consent. The conditions outlined at the beginning of this section can be understood to broadly map onto the two senses of voluntariness that are identified in the Aristotelian distinction outlined in the introductory chapter.

⁹ Foster, *Choosing Life, Choosing Death*, 89.

¹⁰ We may also note that an individual's positive rights are typically understood to be weaker than her negative rights. See Foot, *Virtues and Vices*.

¹¹ Dworkin, 'Rights as Trumps'.

¹² Raz, *Practical Reason and Norms*, 35–48.

¹³ This partly explains Foster's observation that '[a]utonomy in the medico-legal arena is (rightly) much more concerned with preventing unwanted violations than in guaranteeing a right to positive benefits'. Foster, *Choosing Life, Choosing Death*, 28.

¹⁴ This is compatible with the claim that there is nonetheless a requirement to first take measures to enable the individual to attain decision-making capacity.

¹⁵ In the legal context in England and Wales, these approaches are outlined in Mental Capacity Act 2005. For philosophical and legal discussions of these approaches, see Buchanan and Brock, *Deciding for Others*; Hope, Savulescu, and Hendrick, *Medical Ethics and Law*, 82; Herring, *Medical Law and Ethics*, 171–87.

Although my account of decisional autonomy departs from the standard account of autonomy in bioethics, the Aristotelian distinction can also be broadly understood to implicitly frame the latter. As I explained in the introductory chapter, on the standard view, an agent is autonomous with respect to a particular act if it is carried out:

- (1) Intentionally,
- (2) With understanding,
- And
- (3) Without controlling influences that determine their action.¹⁶

Importantly for my purposes in this chapter, Faden and Beauchamp suggest that this account of autonomy can be used to undergird a theory of informed consent understood as a form of *autonomous authorization*.¹⁷ On this view, to give informed consent is to perform a specific kind of autonomous action, one that ‘... authorises a professional to initiate a medical plan for the patient’.¹⁸ A corollary of this is that in order to give informed consent, patients must have certain abilities (such as those required for understanding material information) that are *causally* necessary for meeting the above constitutive conditions of informed consent.¹⁹ To have the abilities in question is to have decision-making capacity, or competence (I shall distinguish the two in the next chapter).

If a patient agrees to a medical procedure without sufficient understanding, unintentionally, or as a result of controlling influence, then their consent may be described as invalid. There is some debate as to whether the concept of invalid consent is morally meaningful. For instance, writing about coerced consent, John Kleinig claims:

... invalid consent no more counts as consent than an invalid vote counts as a vote. It has form but no substance. It is, I believe, more accurate to say that although *A* gave his assent, this did not amount to consent.²⁰

I believe that Kleinig is correct to say this of coerced consent; the assent of the victim of coercion lacks moral substance because their assent reflects the coercer’s authority over them, rather than their own autonomous authorization. However, it would be a mistake to assume that *all* kinds of invalid consent lack any moral substance. In some cases, mere assent can make some moral difference, even if it does not amount to the

¹⁶ Faden and Beauchamp, *A History and Theory of Informed Consent*, 238.

¹⁷ Beauchamp and Childress’ *Principles of Biomedical Ethics* defends a similar view. I consider Faden and Beauchamp’s *A History and Theory of Informed Consent* in this chapter rather than Beauchamp and Childress’ view for two reasons. First, Faden and Beauchamp’s work is solely on the nature of informed consent, and so represents a more focused discussion of the concept. Second, the views on informed consent that Beauchamp and Childress have espoused in *The Principles of Biomedical Ethics* have undergone significant revisions over the numerous editions of the book. However, as Walker acknowledges, Faden and Beauchamp’s account is very similar to the view that is apparent in editions of *The Principles of Biomedical Ethics* that followed it. Walker, ‘Respect for Rational Autonomy’, fn. 3.

¹⁸ Faden and Beauchamp, *A History and Theory of Informed Consent*, 278.

¹⁹ Some accounts include competence itself as a condition of informed consent; however, I prefer to avoid this conflation of constitutive and causal conditions.

²⁰ Kleinig, ‘The Nature of Consent’, 15.

moral significance denoted by the full-blown authorization of valid consent. To illustrate, there is a stronger moral justification for treating a patient who lacks capacity if they assent to the treatment, than if they dissent to it.²¹ One reason for this is that a treatment to which a patient has dissented is likely to be far more distressing for the patient. Furthermore, a process of assent can be important because it allows a patient who lacks capacity to have at least some sort of an input into a decision that affects them, an input that is commensurate with the capacities that they do have.²²

Faden and Beauchamp distinguish informed consent as autonomous authorization from a second institutional sense of informed consent, which pertains to the rules and policies that actually govern informed consent in institutional contexts. What qualifies as an informed consent in this second sense, may or may not amount to the sort of act of autonomous authorization that the first sense means to identify.²³ For instance, there may be cases in which minors may have the capacity to autonomously authorize a medical procedure (and thus provide valid consent in accordance with the first sense), and yet this decision will not be legally recognized as a token of valid consent (in the second sense).²⁴

The institutional sense can incorporate a wide range of conceptions of what 'informed consent' might require. For the purposes of my discussion in the second half of this chapter, it is important to be clear about two different institutional senses of 'informed' consent, which are required to provide sufficient protection against different kinds of torts in medical law, namely battery and negligence.²⁵ First, a physician can be liable to be charged with battery if she touches a patient *without* their valid consent. As I have discussed in the first part of this book, the voluntariness of a patient's decision can be undermined in a manner that may serve to invalidate their consent by controlling influences such as coercion, deception, and manipulation. This much is straightforward; the more complex question is what degree of understanding a patient must have in order to provide valid consent in this institutional sense. Following the case of *Chatterton v Gerson* (1981), the answer to this question in England and Wales is that the patient must be 'informed in broad terms of the nature of the procedure'.²⁶ As Maclean points out, it is relatively easy for physicians

²¹ This is recognized in the law in various ways. In determining what is in the best interests of a patient who lacks capacity, physicians have to take into account the person's current views and feelings. Herring, *Medical Law and Ethics*, 177; Mental Capacity Act 2005, section 4. It is also particularly salient in the Scottish law regarding electro-convulsive therapy (ECT) for patients who lack capacity and who have been diagnosed with a mental disorder. Scottish Government, Mental Health (Care and Treatment) (Scotland) Act 2003 (2003), s. 237–9. <https://www.legislation.gov.uk/asp/2003/13/contents>.

²² Sibley, Sheehan, and Pollard, 'Assent Is Not Consent'.

²³ Faden and Beauchamp, *A History and Theory of Informed Consent*, 276–7.

²⁴ This puts a somewhat simplistic gloss on the hugely complex question of when minors can provide valid consent to medical treatment. For further discussion, see Herring, *Medical Law and Ethics*, 187–94; Hope, Savulescu, and Hendrick, *Medical Ethics and Law*, ch. 10.

²⁵ It would be possible for a physician to be charged with the *criminal* offence of battery in extreme cases, where they have acted maliciously. However, most cases of battery in the medical context are civil rather than criminal cases. Hope, Savulescu, and Hendrick, *Medical Ethics and Law*, 71; Herring, *Medical Law and Ethics*, 150–2.

²⁶ *Chatterton v Gerson* at 265.

to satisfy this informational requirement in disclosure, and the law allows for considerable leeway in its interpretation of the 'broad nature' of medical interventions.²⁷

However, whilst consent grounded by this minimal degree of understanding can be legally valid, and thus invoked to avoid liability to a charge of battery, it is not sufficient to avoid liability to medical negligence. A claim of medical negligence can be raised against a physician if (i) she has failed in her duty of care to the patient and (ii) this failure resulted in the patient suffering a harm. In most cases, actions of battery and negligence can often be distinguished quite easily, because considerations of whether a patient has provided valid consent are often quite distinct from whether the physician harmed a patient by failing to observe her duty of care. However, this distinction can be somewhat muddled by the fact that the physician's duty of care is understood to incorporate a duty to inform her patient of features of the treatment that go beyond its 'broad nature'. For instance, a claim of negligence may be grounded by the fact that the physician neglected to inform the patient of certain risks of a medical procedure, or alternative treatment options. Such information goes beyond that which is required for understanding the treatment in 'broad terms'.

There are interesting questions about these different requirements of information disclosure across these two institutional senses of informed consent, and how they might relate to our understanding of autonomous decision-making. I shall consider these questions in the second half of this chapter. At this point though, we may acknowledge that it is possible to draw a distinction between the 'valid consent' that is grounded by the minimal understanding of the broad nature of a proposed treatment and (ii) the 'substantially informed consent' that can be absent in this case, and which may provide partial grounds for a claim of medical negligence.²⁸

This distinction is not always recognized in bioethical discussions of informed consent.²⁹ However, it is important to clarify it here, as it represents a source of potential confusion given the different ways in which scholars use the language of consent. Following the judge in *Chatterton v Gerson*, some scholars use the term 'real consent' to refer to what I have termed 'valid consent', and the term 'informed consent' to refer to the sort of consent that must be obtained in order to forestall claims of medical negligence.³⁰ This later terminology is somewhat unfortunate, as it may be thought to have the implication that valid or 'real' consent is not 'informed'. As I have explained above though this is a mistake; valid consent must be informed to *some* (albeit lesser) degree.

As such, when I need to refer specifically to the institutional sense of informed consent in my discussion below, I shall instead distinguish the two forms of consent that are operative in discussions of battery and negligence as 'valid consent' and 'substantially informed consent' respectively. Crucially though, in England and

²⁷ Maclean, 'The Doctrine of Informed Consent', 392–3.

²⁸ *Ibid.*

²⁹ For one notable exception, see Walker, 'Informed Consent and the Requirement to Ensure Understanding'.

³⁰ Maclean, 'The Doctrine of Informed Consent'.

Wales, consent can be valid without being what I have called 'substantially informed'.³¹

2. Autonomy-Based Justifications of Informed Consent

The above discussion indicates that it is important to be clear about the sense of consent that one means to invoke, when seeking to justify a particular criterion of informed consent by appealing to the principle of respect for autonomy. In defining the provision of consent as a specific kind of autonomous action in the first sense, Faden and Beauchamp draw an inextricable link between consent and the principle of respect for autonomy. However, this sense is not co-extensive with the second institutional sense of informed consent, and we should not assume that they share the same justification.

Indeed, the development of the institutional sense of informed consent in the legal context raises significant challenges for any philosophical investigation into the topic. As Richard Ashcroft notes, we should not expect to find elusive, abstract philosophical concepts such as autonomy in the law, as the law requires more concrete concepts that can be tested and consistently applied in litigation.³² We might also add to Ashcroft's claims that the philosophical bioethicist has something of an easy way out of complex debates in medical law; upon finding that the law fails to reflect a philosophical principle upon which it purports to be based, they can simply say 'so much the worse for the law'. However, this is not a particularly useful practical avenue for the medical lawyer who has to address these issues whilst working within this framework, which is shaped by a number of competing and conflicting justifications beyond philosophically pure abstract principles. In particular, the institutional sense of informed consent may need to serve a wide range of purposes whilst being constrained by the practical realities of the clinical encounter. In turn, this may legitimize employing lower thresholds for understanding, disclosure, and capacity than might be required for a decision to qualify as an autonomous authorization.³³ However, whilst acknowledging this, the concept of autonomy can still coherently play some, perhaps non-exhaustive role in its justification. It can serve as a guiding value that we may invoke to shape the contours of the institutional requirements of informed consent, so that they better serve the autonomy of patients, amongst the other purposes that they are intended to fulfil.

With that said, I shall begin by framing my discussion by considering the relationship between autonomy and the first sense of informed consent as a form of autonomous authorization. On this first sense, the relationship between informed consent and autonomy is straightforward; to provide informed consent is just to make a certain kind of autonomous decision, a decision to authorize a particular medical treatment. The positive obligation imposed by the requirement to obtain informed consent can be understood to amount to an obligation to help facilitate

³¹ As such, this jurisdiction does not observe the so called 'doctrine of informed consent'. Herring, *Medical Law and Ethics*, 158.

³² Ashcroft, 'Law and the Perils of Philosophical Grafts'.

³³ Beauchamp and Childress, *Principles of Biomedical Ethics*, 123.

autonomous decision-making, and the negative obligation to ensure that medical interventions do not involve the infringement of rights that have not been waived by their holders.

The rationalist account of autonomy that I have developed is meant to supplement the standard account of autonomy. Indeed, one of the strengths of the rationalist approach that I have developed over the course of the preceding chapters is that it can add further explanatory depth to the conditions laid out in the standard account. For instance, in the previous chapters, I have explained how my rationalist approach can provide a principled account of how the different forms of controlling influence captured by condition (3) above undermine voluntariness. Consider also condition (1). Faden and Beauchamp claim that an action must be intentional in order for it to be autonomous, and that ‘... an intentional action is action willed in accordance with a plan’.³⁴ Among the category of non-intentional acts, they include:

... things that persons do inadvertently, certain habitual behaviours, and instances of so called occurrent coercion in which a person is physically forced by another to do something.³⁵

Intentional action understood in this way can also be interpreted as a necessary condition of autonomy on the account that I have defended. To will an action in accordance with a plan, can be understood as willing action in accordance with one’s beliefs about what one has reason to pursue; moreover, the non-intentional actions delineated in the above quotation are also inimical to acting on the basis of this sort of rational deliberation.

The criterion of understanding is also congruous with the account that I developed in the first half of the book. To recap, I have argued that there can be decisionally necessary beliefs, in the sense that one must hold certain true beliefs about central features of one’s choice in order to make an autonomous decision in that context. Furthermore, I have argued that failing to abide by norms of theoretical rationality can undermine the sort of understanding that autonomy requires, in the sense that it jeopardizes one’s ability to assess the extent to which a particular belief coheres with one’s other beliefs about both descriptive and evaluative features of the world.

One of my aims in this chapter is to provide more details about the implications of my theory for the scope of the criterion of understanding. Here though, we may simply acknowledge that Beauchamp and Childress claim that in order to autonomously authorize a medical procedure, the patient must have *substantial*, but not full understanding. They note that the patient’s understanding of ‘diagnoses, prognoses, the nature and purpose of the intervention, alternatives, risks and benefits, and recommendations’ is ‘typically essential’ for such understanding.³⁶ Notice then that the criterion of understanding employed by Beauchamp and Childress for autonomous authorization goes beyond that required by ‘valid consent’ in the institutional sense. Instead, it bears a resemblance to that which is required by the doctrine of informed consent in the institutional sense, or what I have called ‘substantially informed consent’. I shall return to this point below.

³⁴ Herring, *Medical Law and Ethics*, 187–94.

³⁵ *Ibid.*

³⁶ Beauchamp and Childress, *Principles of Biomedical Ethics*, 132.

Despite this significant degree of congruence with the standard account, my rationalist theory of autonomy departs from this view by advocating that we should supplement the standard account with rationality conditions of the sort set out in Chapter 2. Thus, I suggest that the following ought to be understood as a necessary condition of the voluntariness element of informed consent:

Rationality Condition: If an agent is to provide informed consent to some intervention, then they must also endorse their desire to undergo that intervention with a personally authorized preference.

The condition has significant implications for what we should want the informed consent process to achieve, if it is to facilitate autonomous decision-making. Merely ensuring the sufficient understanding of material information is not enough to facilitate autonomous decision-making, if that understanding remains unconnected to the patient's values. The condition thus speaks against viewing the patient as a passive recipient of disclosure who is 'to be informed'; instead, it speaks in favour of the patient actively contributing to the process, so that physicians can tailor disclosure to what matters to the patient, given her values. Of course, by adding this condition, I am implicitly widening the gap between what informed consent as an autonomous authorization should look like, and what informed consent in the institutional sense currently requires.

Informed consent in this first sense can thus facilitate autonomous decision-making, in so far as the process of informed consent helps to enable individuals to decide to authorize medical procedures in accordance with the above sorts of conditions. However, it also has an important role to play in respecting the agent's autonomous preferences. It has this role by virtue of the negative obligation the requirement to obtain consent imposes on others to refrain from certain kinds of action, in the absence of consent. In this chapter, I shall be interested primarily in the justification of informed consent in the context of medical practice, when that is understood to refer only to the provision of medical therapy but *not* the performance of non-therapeutic research.³⁷ However, by distinguishing the ways in which individuals can face both internal and external impediments to autonomy, and by stressing the importance of the rational endorsement of one's preferences to autonomy, the account that I have delineated serves to highlight an illuminating contrast between the justification of informed consent requirements in the context of non-therapeutic medical research, and the justification of informed consent requirements in medical practice.

It is commonly claimed that there is an important distinction between medical research and therapy due to the primary aims of each activity.³⁸ For instance, the

³⁷ The research context raises different questions about appropriate standards of understanding and disclosure. For analysis, see Sreenivasan, 'Does Informed Consent to Research Require Comprehension?'; Bromwich, 'Understanding, Interests and Informed Consent'; Bromwich and Millum, 'Disclosure and Consent to Medical Research Participation'.

³⁸ See Miller and Brody, 'Clinical Equipoise and the Incoherence of Research Ethics' for a discussion of the different ethical norms governing medical research and therapy. However, see Beauchamp, 'Viewpoint' for an argument against drawing a hard and fast distinction between research and therapy.

Belmont Report defines medical therapies (in part) as ‘interventions... designed solely to enhance the well-being of an individual patient...’ whilst it defines medical research (in part) as ‘an activity designed to test a hypothesis, permit conclusions to be drawn, and thereby to develop or contribute to generalizable knowledge’.³⁹

The prioritization of generalizable knowledge in scientific research suggests that the intended primary beneficiaries of an individual’s participation in a research study are the third parties who stand to benefit from this knowledge. The value of ‘generalizable knowledge’ in this context may be understood as a proxy for the interests of society at large, and future patients. Accordingly, it seems that one of the primary purposes of informed consent in the research context, is that it affords some protection to potential subjects from being put at risk of harm (broadly construed), or having their rights violated on the basis of broadly consequentialist justifications of research.

This is clearest in Phase I trials of a novel medical intervention, in which healthy volunteers are provided with sub-therapeutic doses of an intervention in order to establish its safety. Participants in such trials are put at risk of harm without any real possibility of benefiting from the intervention itself, although they may be benefited indirectly by payment for their participation.⁴⁰ In this way, informed consent requirements can be understood as safeguarding autonomy, in so far as they help to ensure that it is the individual herself who determines whether she wants to waive her claim rights against bodily interference (amongst others), and thereby put herself at risk of the harm that the research might entail for the benefit of others (or, if payment is offered, for indirect benefit). When the individual subject’s autonomy conflicts with the interests of future patients, the former trumps the latter.⁴¹

Informed consent in the context of medical practice partly plays a similar role, in so far as it serves to protect patients from being forced into receiving treatments that might serve another party’s interests rather than their own.⁴² However, informed consent also plays a further role in medical practice that is not applicable in the context of non-therapeutic medical research, in so far as the interventions for which consent is being solicited in the former context are primarily intended to directly benefit the patient herself. Indeed, we may notice that the definition of medical practice quoted above stresses that the aim of a therapeutic intervention is to enhance the recipient’s own *well-being*. Part of the reason that respecting autonomy is important in this context is that it gives the patient a say in the matter of what is really in their own interests, in the light of their own desires and values.

The importance of this is made clear once it is observed that patients can differ from their physicians in their conclusions about what they have strongest self-interested reason to do, even in view of the same relevant descriptive facts. To illustrate, consider this example:

³⁹ The Belmont Report.

⁴⁰ As I explored in Chapter 4, such payment can raise the spectre of coercion; see Emanuel and Miller, ‘Money and Distorted Ethical Judgments about Research’.

⁴¹ World Medical Association Declaration of Helsinki.

⁴² For a relevant legal example of this, see *Appleton v Garrett* – Case Summary.

Suppose that Joe is undergoing an operation to remove a tumour from his diaphragm. An anaesthetist consults him regarding his post-operative analgesia. The efficiency of analgesia in Joe's case is crucial, since if the analgesia is not effective, it is likely that Joe's lung will collapse and lead to the development of a potentially fatal pneumonia. Joe is given the choice between an analgesic that poses a very small risk of spinal cord damage (such as a thoracic epidural), and one that is considerably less effective but which poses no such risk (such as an intravenous narcotic infusion).⁴³

Here, the thoracic epidural is medically indicated; however, Joe might still rationally choose to receive an intravenous narcotic infusion instead, if he places sufficient weight on the value of pursuits that involve physical activity. For example, suppose that Joe is a professional athlete, and believes that his life would not be worth living if he became paralysed. In such a case, the possibility that the more effective analgesia could paralyse Joe might give him reason to believe that the less effective analgesia, which did not pose this risk of paralysis, was the preferable treatment option.

This case suggests that the role of informed consent requirements in medical practice is not merely to protect the patient from competing third-party interests, but also to ensure that the treatment that the patient receives is in accordance with what they want for themselves; and this may or may not coincide with what the physician believes is in the patient's best interests. Even if rational agents agree that they have some reason to pursue an outcome, this does not entail that they will agree on the strength of that reason, relative to their reasons to pursue other outcomes.⁴⁴

This observation goes some way towards explaining the intuitive pull of the claim that informed consent in medical practice can be justified in large part by an appeal to the principle of respect for autonomy, at least on the broadly Millian understanding that I have been developing in this book. Informed consent requirements can be understood as facilitating the patient's self-governance, not just because they protect the patient from controlling forces, or from their being exploited in the interests of others. They also give the patient the power to make their own treatment decisions, in accordance with their assessment of the strength of the reasons that they have to pursue different outcomes. At least part of the importance we attribute to informed consent can thus be construed as a reflection of the Millian thought that I highlighted in the introduction of this book, namely that we value laying out our own mode of our existence, even if our way of doing so is not prudentially optimific by third-party standards.

The view that the moral significance of informed consent is to be justified by an appeal to the value of patient autonomy has often been treated almost as a truism in bioethics. Indeed, as I illustrated above, Faden and Beauchamp simply define informed consent as a type of autonomous authorization, whilst autonomy has elsewhere been described as the 'ultimate moral foundation'⁴⁵ of informed consent. However, it is important to be clear about the extent of this claim, at least as I shall understand it in this chapter. As I suggested above, the moral significance of

⁴³ Savulescu, 'Rational Non-Interventional Paternalism', 327.

⁴⁴ For a discussion of how this point featured in the justification of the Montgomery judgement discussed below, see Herring et al., 'Elbow Room for Best Practice?', 8–9.

⁴⁵ Young, 'Informed Consent and Autonomy', 441.

informed consent in this sense can be grounded in large part by the broadly Millian conception of autonomy that I have been developing in this book. Yet, there are two important mistakes that we must avoid here.

First, it would be a mistake to assume that the Millian conception of autonomy provides the *only* justification for the first sense of informed consent that I have primarily been considering so far. There is a considerable literature that seeks to explore the ways in which Kantian conceptions of autonomy can also ground informed consent requirements. These approaches particularly focus on the role that such requirements play in ensuring that individuals are respected as ends in themselves, in a manner that is consummate with their human dignity.⁴⁶ The Kantian approach thus focuses on the role that informed consent plays in respecting the supreme moral status of rational agents; in contrast, the Millian approach places emphasis on the manner in which informed consent facilitates the agent's pursuit of their own conception of the good life. Nonetheless, these different forms of justification can both plausibly lend support to informed consent in its non-institutional sense.

Second, and perhaps more importantly, it would also be a mistake to claim that either (or indeed both) of these conceptions of autonomy can provide the sole justification of informed consent in its second *institutional* sense. Autonomy-based justifications of informed consent have sometimes been criticized on this score, with philosophers pointing to examples in which it appears that our reasons to abide by informed consent requirements are not best understood as being grounded by considerations of autonomy.⁴⁷ For example, one clear legal function of informed consent in the medical context is that it serves to provide physicians with a record of what has occurred in the course of providing treatment, a resource to which they can appeal to in cases of litigation.⁴⁸ It has also been suggested that informed consent procedures are integral to establishing a relationship of trust between doctors and patients,⁴⁹ or safeguarding a personal sphere of self-ownership that is not best understood in terms of autonomy.⁵⁰

Although there is a great deal of truth in this general criticism, it is somewhat perplexing as an objection to autonomy-based justifications of informed consent. The fact that there may also be non-autonomy based justifications of informed consent in the institutional sense, speaks little against the claim that it may nonetheless derive an integral part of its justification from the fact that it *does* facilitate and enable us to abide by the principle of respect for individual autonomy, in a considerable number of cases.

To illustrate this point, consider David Archard's criticism of the relationship between autonomy and informed consent. Archard offers the example of 'inserting a

⁴⁶ Donagan, 'Informed Consent in Therapy and Experimentation'.

⁴⁷ See Taylor, *Practical Autonomy and Bioethics*, 133; Dworkin, *The Theory and Practice of Autonomy*, 103; Archard, 'Informed Consent'.

⁴⁸ Brock, *Life and Death*, 47–8.

⁴⁹ O'Neill, *Autonomy and Trust in Bioethics*, 145; see also Bok, *Lying*, 11, 26–7, and 63; Jackson, 'Telling the Truth', 491. For recent challenges to this view, see Eyal, 'Using Informed Consent to Save Trust'.

⁵⁰ Archard, 'Informed Consent'.

swab into someone's mouth without her agreement, yet harmlessly, painlessly, and without coercion or deception'.⁵¹ He suggests that a Millian conception of autonomy cannot provide a sufficient explanation of the wrong involved in this failure to observe the requirements of informed consent because 'the value of autonomy is to be found in *the leading of lives*'.⁵² As such, we only have autonomy-based reasons to respect decisions regarding 'critical life-choices'. Accordingly, we must thus appeal to a right against bodily trespass to cash out the wrong involved in the swab example.

I am somewhat sceptical of Archard's suggestion that a putative right to bodily trespass can be entirely divorced from such an autonomy-based justification. This might be true of some rights, particularly those that incorporate the Hohfeldian incident of a *claim* but not the *power* to waive that claim.⁵³ The thought underlying such putative unwaivable rights is that both autonomous and non-autonomous individuals have a very strong interest against certain kinds of interference that might justify affording them this unwaivable claim to protection. However, many rights (including the right against bodily trespass) incorporate both a claim to protection *and* the second-order power to waive that claim, at least in so far as the right is held by an autonomous individual.⁵⁴ Crucially though, it is difficult to see how the *powers* incorporated in these rights can be justified without some appeal to the value of autonomy (even if the claims themselves might be so justified). Indeed, it seems that the power to waive one's claim should only be granted authority if the right-bearer has decided to waive their claim in a locally autonomous fashion. In this way, considerations of local autonomy seem highly relevant to our understanding of aspects of these rights and their justification.⁵⁵

Notwithstanding this point about the relationship between rights and autonomy, even if we agree with Archard that the wrong involved in cases of minor bodily trespass is entirely divorced from considerations of autonomy, it is clear that the requirement to obtain informed consent can be, and indeed *is* invoked in many critical life-choices that do generate strong autonomy-based reasons. The example of Joe above is just one case in point, in which the value of respecting Joe's choice is integral to the value of his leading his own life.

Accordingly, we should qualify the autonomy-based justification of informed consent with the caveat that autonomy is not the *only* justification of the informed consent in its institutional sense, even if it does provide a strong justification for it in many cases. Yet, even when the view is qualified with this caveat, some philosophers have further objected that informed consent in the institutional sense cannot be justified by an appeal to respect for autonomy, because it is not and cannot be sufficient for that purpose. For instance, Neil Manson and Onora O'Neill have argued that this is a decisive problem with justifying informed consent procedures by appealing to the principle of respect for autonomy. The problem is that such requirements fail to ensure that patients will choose autonomously; they only require that physicians respect the choices that the patient actually makes, whether or not

⁵¹ *Ibid.*, 19.

⁵² *Ibid.*, 21.

⁵³ Wenar, 'The Nature of Rights'.

⁵⁴ *Ibid.*

⁵⁵ One explanation of this might be that the power incorporated into the right is grounded by considerations of agential authority, whilst the claim itself is justified by considerations of interest. For discussion of the basis of rights, see *ibid.*

this choice is autonomous or rational.⁵⁶ Not only that, but Manson and O'Neill also claim that if informed consent requirements were reformulated so that they would protect only rational, autonomous choices, then they would become too demanding for the vast majority of patients.⁵⁷

Manson and O'Neill invoke a broad understanding of 'rational choice' here, claiming that theories of autonomy can understand rational choice as 'reflectively evaluated, or endorsed by second order desires';⁵⁸ we may note that this understanding is compatible with many of the theories that I surveyed (and rejected) in Chapter 2. However, it is clear that the standard view of autonomy, which is the primary target of their attack, does not incorporate conditions pertaining to the rationality of the patient's choice, even on this broad understanding. Indeed, advocates of the standard view of autonomy explicitly reject the suggestion that a theory of autonomy should incorporate a condition that requires that the patient's choice be consistent with their reflectively accepted values, because such a condition would, they claim, make autonomy too demanding.⁵⁹

As the preceding chapters of this book should make clear, I believe that Manson and O'Neill are correct to claim that the standard view of autonomy and informed consent is inadequate. If we are to claim that a primary purpose of informed consent requirements is to safeguard patient autonomy, then we should incorporate conditions pertaining to the *rationality* of the patient's choice into our theory of informed consent. In view of the fact that Manson and O'Neill reject this solution (because it would, they claim, make the standards of informed consent too demanding), two strategies are possible. First, we could abandon the project of justifying informed consent requirements by an appeal to the principle of respect for autonomy. This is the strategy that Manson and O'Neill adopt; they argue that we ought to view an agent's provision of consent to a procedure as a waiver of an ethical and/or legal norm against performing the act in question, in limited ways in a particular context.⁶⁰ On the other hand, we might maintain that informed consent requirements are to be justified by an appeal to the principle of respect for autonomy, and supplement informed consent requirements with conditions that will facilitate rational choice, but which will not render informed consent requirements too demanding.

Over the course of the next two chapters, I shall adopt the latter strategy.⁶¹ Again, it is important to be clear about the scope of what I shall attempt to claim. *Contra*

⁵⁶ Manson and O'Neill, *Rethinking Informed Consent in Bioethics*, 21. More recently, in a similar vein, Shlomo Cohen has claimed that 'conceptions of autonomy cannot provide practical ethical guidelines generally or specifically for informed consent'. Cohen, 'A Philosophical Misunderstanding at the Basis of Opposition to Nudging', 39.

⁵⁷ Manson and O'Neill, *Rethinking Informed Consent in Bioethics*, 21. ⁵⁸ *Ibid.*

⁵⁹ Faden and Beauchamp, *A History and Theory of Informed Consent*, 262–4; Beauchamp and Childress, *Principles of Biomedical Ethics*, 103.

⁶⁰ Manson and O'Neill, *Rethinking Informed Consent in Bioethics*, 72. See also 69–84.

⁶¹ Though I cannot argue fully for why I reject Manson and O'Neill's strategy here, I shall sketch my main points of disagreement. First, it is not clear that their theory really divorces autonomy from informed consent in the way that they claim; after all, if consent transactions are meant to signify the patient's waiving a significant legal or ethical norm, it must surely be the case that the patient should still be autonomous with respect to their decision to waive that norm. Manson and O'Neill do warn against the possibility of bogus consent, in which consent is solicited in ways that violate ethical norms; they give the

Manson and O'Neill, requiring that an adequate conception of informed consent must *ensure* that patients make autonomous choices is an unreasonably high bar. Such a view overlooks the fact that the individual herself has an indispensable role to play in their own autonomy; disclosure, and the absence of controlling influences on choice means nothing for the patient's autonomy, if she herself is either unwilling or unable to contribute to the decision-making process. It is thus unfeasible to demand that informed consent in the institutional sense must alone *ensure* that patients make autonomous decisions; the best that we can hope for is that it will *facilitate* the patient's making an autonomous decision, and ensure that the autonomous decisions that she does make are respected.

As I suggested above, the main source of objection to the sort of rationality condition that I have proposed is that it would make the conditions of informed consent too demanding. This is a serious objection that the account must answer. However, it is an objection concerning the *causal* conditions of autonomy and the standards of capacity that they imply, rather than an objection to the rationality condition itself as a *constitutive* condition of autonomy and informed consent. As such, I shall postpone my consideration of it until I am in a position to discuss the issue of capacity in the following chapter. In the next section, I shall turn to further consider the role of autonomy in justifying informed consent in the institutional sense, by considering its relationship to 'valid consent' and 'substantially informed consent', as they are distinguished in the context of medical law governing claims of battery and negligence.

3. Battery, Negligence, Beliefs, and Decisional Autonomy

In the legal domain, questions pertaining to informed consent have been framed not so much in terms of what patients need to *understand* in order to be autonomous with respect to their treatment decision, but rather in terms of what physicians need to *disclose*.⁶² Whilst it is easier to enforce a legal requirement to disclose information than it is to enforce a requirement to ensure understanding, the former receives only limited justification from considerations of autonomy. First, disclosing information to a patient may not be necessary for safeguarding autonomous decision-making in some cases. Most obviously, this can be so if the patient is *already* aware of that information. For example, it seems that I could provide valid consent to a physician's application of a bandage to my wound, without their having disclosed to me that I am bleeding profusely, and that applying a bandage to a wound will help to stop the bleeding.

examples of consent following coercion, force, and duress as examples of bogus consent. See *ibid.*, 92. However, what, we might ask is the basis for *these* ethical norms? Whilst Manson and O'Neill appeal to a principled Kantian sense of autonomy to ground such norms, an alternative plausible explanation would be that coercion, force, and duress are wrong because they violate the patient's autonomy in the Millian sense that I have outlined (see also O'Neill, *Autonomy and Trust in Bioethics*, ch. 4, especially 83–6).

⁶² See Beauchamp and Childress, *Principles of Biomedical Ethics*, 2nd edition, 67.

This example suggests that the sort of understanding that is required for patients to make autonomous decisions can sometimes be implicitly assumed with some justification, and patients can in some cases provide tacit consent that authorizes a procedure, without going through the rigmarole of institutional procedures of informed consent. Of course, given the expertise gap between physicians and their patients, and since physicians cannot be certain about their patients' prior knowledge, patients will often need to have information disclosed to them. However, it would also be a mistake to assume that disclosure of relevant information is *sufficient* to ensure the degree of understanding that is necessary for autonomous decision-making. Even competent patients may struggle to process the information presented to them if it is not adequately explained to them. Indeed, this speaks to a considerable tension in informed consent with which I shall be concerned in this chapter. The disclosure of information that may be material to a patient's decision may also serve to undermine the patient's ability to make an autonomous decision in that same context, for reasons I shall explain.

Disclosure, then, is neither a necessary nor sufficient condition of ensuring the degree of understanding that is required for autonomous decision-making. That said, it is perhaps the most powerful tool at our disposal when thinking about how best to facilitate autonomous decision-making in the institutional context. This is particularly so in the context of medicine, in which there is a significant expertise gap between the parties soliciting and providing consent.

As I explained above, informed consent in the institutional sense draws a distinction between the amount of information that must be disclosed (and understood) in order for a patient to provide 'valid consent', and that which must be disclosed (and understood) in order for the patient to provide 'substantially informed consent' to the intervention. Recall that a physician's receipt of the patient's 'valid consent' to an intervention can be invoked as a defence against battery, whilst a physician must have obtained 'substantially informed consent' in order to defend themselves from claims of negligence.

The torts of battery and negligence can cover quite separate kinds of action, and I do not have space to list the various differences between the two here.⁶³ However, these two legal categories somewhat overlap when it comes to the question of information disclosure in the medical context.⁶⁴ Indeed, the kind of information of which the patient is ignorant when they agree to undergo a procedure, can make all the difference between whether the patient has given valid consent or not. The question I want to consider in the remainder of this section is whether there are any autonomy-based reasons to support the way in which the law stipulates the lower threshold of understanding required for valid consent, in comparison to the standard deployed for claims of negligence.

Without discounting the possibility of other justifications, I want to suggest that considerations of autonomy do lend support to the distinction between what I have called 'valid consent' and 'substantially informed consent'. To recap some arguments

⁶³ For discussion, see Herring, *Medical Law and Ethics*, 152–3.

⁶⁴ For an argument that the Montgomery judgement represents a way in which the legal distinction between risk disclosure and consent is breaking down, see Herring et al., 'Elbow Room for Best Practice?'

I made in the preceding chapters, I have argued that there is a cognitive element of decisional autonomy; if an agent holds certain false beliefs, or fails to hold certain true beliefs, they may be precluded from making certain decisions voluntarily in that context. This will be so if the beliefs in question are decisionally necessary; I also suggested a modal test that may be employed to identify some of these beliefs.

The individual's failure to hold such beliefs undermines their decisional autonomy, regardless of how this came to be the case, be it through intentional third-party deception, or the non-intentional omission of true information. Accordingly, on the account that I have developed, if a physician performs a medical procedure on the basis of a token of consent that is grounded by a lack of understanding of this key information, it can be plausible to say that the physician has not only failed to act in accordance with their duty of care; they have also performed a medical procedure on the basis of a non-voluntary, and thus invalid token of consent.

Crucially though, I argued that the absence of many true beliefs is compatible with decisional autonomy. Our decisional autonomy is not undermined by our lack of true beliefs about future states of affairs, or the actual, rather than predicted, consequences of our choices. Of course, our practical autonomy may be enhanced by holding such beliefs; we will be more likely to achieve our goals if our beliefs about these matters are true. However, our ignorance of these matters does not render our decisions at the point of choice non-voluntary.

In light of the above recap, the key question for my purposes here is whether the legal distinction between valid consent and substantially informed consent maps onto the distinction between beliefs that are decisionally necessary, and those that are not. I believe that it does. To see why, recall that in order for a patient to provide valid consent to a medical procedure, the law states that they must be informed 'in broad terms of the nature of the procedure'. On my view, the reason that such ignorance can be said to be incompatible with valid consent, is that it modally precludes the patient from achieving their ends, by divorcing the decision they make from the values that are actually operative in their particular choice context. A patient who agrees to undergo a vasectomy without understanding that it will render him infertile does not provide valid consent to the procedure, because their understanding of the intervention is so poor that their informational condition precludes them from choosing what to do in accordance with their own values.⁶⁵

In contrast, the sort of further information that is required for 'substantially informed consent' (but not merely 'valid consent') is not captured by this modal test. In previous chapters, I have argued that the concept of valid consent must be compatible with the fact that many patients must make treatment decisions in the light of severely restricted choices. As such, it seems that ignorance about alternative treatment options can be compatible with valid consent, even if not substantially informed consent. Furthermore, most information pertaining to risk will not be captured by the modal test that we can use to identify decisionally necessary beliefs, since there is typically a close relationship between modal possibility and probability.

⁶⁵ This also accounts for why the therapeutic misconception is such a problem for consent in medical research. See Henderson et al., 'Clinical Trials and Medical Care'.

This means that ignorance about risk cannot be understood to modally preclude an individual from acting effectively in pursuit of their ends, as I defined this in the previous chapter. Even though the risk-event that actually occurs in this world might have precluded the individual from achieving their ends in this world, there is another nearby possible world in which this risk-event does not eventuate. Indeed, Duncan Pritchard argues that very low-probability events can nonetheless be modally close.⁶⁶ Thus, one is not modally precluded from achieving one's ends by ignorance about the degree of risk attending different courses of action.

That said, in discussing how beliefs about risk can affect the voluntariness of our decisions, it is important to be clear that 'ignorance about risk' can cover two importantly different types of ignorance. First, it can relate to ignorance that one is assuming *any* risk by engaging in an activity; alternatively, it can relate to ignorance about the *particular degree* of risk that one is assuming. The first sense outlined above is incompatible with decisional autonomy. The reason for this is that in consenting to a medical procedure, a patient is not simply being asked to consent to the procedure itself; they are also being asked to consent to the assumption of risk. If we agree with the Aristotelian claim that ignorance of particulars can undermine voluntariness, then it seems that ignorance that there are risks associated with a medical procedure can undermine the voluntariness of a patient's consent to that procedure, simply by virtue of the fact that this form of ignorance means that they do not understand a significant part of what it is they are consenting to, namely, assuming a risk. However, this does not entail that ignorance about the degree of risk you are taking on, or indeed ignorance about future states of affairs, similarly precludes decisional autonomy.

Accordingly, I believe that the distinction between valid consent and substantially informed consent can be philosophically grounded by considerations of autonomy, by virtue of the fact that some beliefs are decisionally necessary, whilst others are not. Notice though that this is quite compatible with the claim that holding these further true beliefs may nonetheless serve to enhance an agent's autonomy. Indeed, it would be a mistake to conclude from the above discussion that questions pertaining to the disclosure required for 'substantially informed consent' are divorced from considerations of autonomy. Such a conclusion might seem tempting, since the law treats a failure to secure substantially informed consent as a breach in the doctor's duty of care, rather than an issue with the patient's autonomy *per se*. Yet, even if this is so, one must still account for *why* this amounts to such a breach in the doctor's duty of care. To my mind the most plausible answer is that the doctor has a positive duty to facilitate their patient's ability to make autonomous decisions. This includes facilitating autonomy beyond that which is required for the minimum threshold of understanding that is necessary for consent to be valid. In addition to enhancing the agent's decisional autonomy by increasing their understanding of their options, and indeed the number of options available to them, such disclosure can serve to facilitate the patient's autonomy by facilitating their own self-trust.

⁶⁶ Indeed, Pritchard argues that very low-probability events can nonetheless be modally close. See Pritchard, 'Risk'.

However, deciding on how one can best act in accordance with the positive obligation to increase decisional autonomy beyond that which is minimally required for valid consent is complex. Physicians cannot be expected to disclose *all* the available information about the risks attending a medical treatment, and patients cannot be expected to understand it. Medical conditions, procedures, and their attendant risks often admit of exceedingly complex descriptions, which, however accurate they might be, are unlikely to aid the patient in their decision-making.

This has important implications for disclosure. If physicians understood the obligation to facilitate autonomy to require that they aim for full understanding amongst their patients, they would be likely to overwhelm them with an excess of information that they could not reasonably be expected to compute, especially given that many patients will be less able to deal with complex information because of their illness. There is thus an important balance to be struck between (i) providing patients with the information that can enable them to identify the values at stake in their decision, and to form an impression of the strength of their apparent reasons, and (ii) refraining from providing so much information that the patient is unable to utilize it in a process of rational deliberation.

The theoretical apparatus we may appeal to in seeking to strike this balance is the concept of materiality; to adequately facilitate substantially informed consent, and to thus discharge their duty of care to their patient, we may say physicians must only disclose information that is *material* to the patient's decision. I shall conclude this chapter by considering how we should understand this concept of materiality, and the relationship between different standards of disclosure and autonomy.

4. Standards of Disclosure

According to a *physician*-oriented view of materiality, it should be solely up to the physician, in their professional capacity, to decide which information is material to their patient's decision. In turn, the physician's decision here is understood to be governed by a standard of disclosure endorsed by the professional community, according to which information ought to be disclosed if the majority of physicians within that community would customarily make such a disclosure. Until recently, this view was enshrined in the law in England and Wales in the form of the so-called 'Bolam test', named after the judge's assessment in *Bolam v Friern Hospital Management Committee* (1957) that 'a doctor is not guilty of negligence if he has acted in accordance with a practice accepted as proper by a responsible body of medical men skilled in that particular art'.⁶⁷ This test was then deemed to determine the boundaries of the physician's duty to disclose information to their patient in a later case.⁶⁸

⁶⁷ *Bolam v Friern Hospital Management Committee* at 1 WLR 583.

⁶⁸ *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985]. For discussion, see Miola, 'On the Materiality of Risk'. In the US, see *Robinson v. Bleicher* (559 N.W.2d 473) 1997.

The Bolam test has been widely criticized as a paternalistic standard of information disclosure.⁶⁹ Part of the problem with the test is that a 'responsible body of medical men' might find it proper to omit information on the basis of considerations of beneficence alone, rather than patient autonomy. That is, the professional body could plausibly seek to justify the omission of certain information disclosure on the basis that such disclosure would lead patients to refuse treatments that are in their best interests. Indeed, on this justification, a physician could even refrain from disclosing information that the patient explicitly requested.

Extending the Bolam judgement to the context of risk disclosure was to leave the door open to this kind of paternalism. The natural rebuke to doing so was that this sort of paternalism runs contrary to the underlying autonomy-based justification of informed consent; part of the purpose of the informed consent requirements is to empower patients to make their own decisions about what happens to them, and to avoid paternalistic interference.⁷⁰ Partly on this basis, other patient-centric standards of materiality, which purport to emphasize patient autonomy over the paternalism of the physician-oriented approach, have been mooted.

In their philosophically grounded approach to substantially informed consent, Faden and Beauchamp explicitly advocate a purely subjective account of materiality. They claim that:

... a person must understand those propositions about (some medical intervention) *R* and about authorizing *R* that are germane to the person's evaluation of whether *R* is an intervention the person should authorize. This criterion is entirely subjective.⁷¹

Call this the subjective patient-oriented account of materiality.

The justification for adopting a subjective approach is that different patients are likely to regard different information to be pertinent to their treatment decision. Yet, one problem with the subjective account is that it fails to explain precisely *how* patients are to subjectively assess whether or not certain information is pertinent; Faden and Beauchamp simply point out that a person's long term goals and values can affect how individuals value various act descriptions.⁷² However, it seems that in order for information to qualify as material to a patient's decision, even if only subjectively, there must be some plausible basis upon which the individual understands the information to be pertinent to her decision. Call this the subjective assessment problem.

One way in which the subjective assessment problem raises a difficulty for the subjective account arises when we recall that understanding some information is decisionally necessary. That is, certain information is so fundamental to the nature of a decision that one cannot make that decision voluntarily if one remains ignorant of it. Crucially, this is so regardless of whether an individual deems the information to be material or not. For instance, suppose Jerry feels ill and continually takes antibiotics because she believes that they are the only thing that will cure her; she does

⁶⁹ Brazier, 'Patient Autonomy and Consent to Treatment'; Jones, 'Informed Consent and Other Fairy Stories'; Miola, 'On the Materiality of Risk'; Brazier and Miola, 'Bye-Bye Bolam'.

⁷⁰ Indeed, this anti-paternalistic argument can be found in the judgement in *Chester v Afshar* at 18.

⁷¹ Faden and Beauchamp, *A History and Theory of Informed Consent*, 302. ⁷² *Ibid.*

not believe that the fact that she is suffering from a viral rather than a bacterial infection is material to her decision. Irrespective of her own views regarding the materiality of this information, Jerry's failure to appreciate the significance of this precludes her from receiving an effective treatment. Similarly, it is difficult to imagine how a patient could be autonomous with respect to their decision to undergo an anaesthetic if they failed to understand that undergoing an anaesthetic will render them unconscious.

One might defend the subjective standard on this score by saying that the standard is only meant to apply to information disclosure beyond that which is necessary to secure valid consent. However, the subjective assessment problem raises its head in other ways. In view of the fact that patients are normally not experts in medicine, and may be in a vulnerable state owing to the nature of their condition, they may make mistakes about what information is and is not material to their treatment decision. Indeed, as I shall explain below, empirical evidence suggests that patients are subject to a number of cognitive biases that can distort their understanding of their condition and treatment options. A second problem then with the subjective account, is that a particular patient may attach significance to information, but not in a manner that bespeaks an adequate understanding and processing of that information in rational deliberation.

Finally, it is not clear that a purely subjective account of materiality is practically realizable; first, it will often be difficult for practitioners to know what information a patient believes to be relevant to their decision, and it is the physician who has to decide what information to disclose to their patient. Crucially, since physicians are liable to negligence if they fail to disclose information deemed to be material by the standard invoked in medical law, a requirement that doctors must disclose any and all information that a patient could deem to be material would leave doctors highly vulnerable to litigation. Moreover, we might also be concerned that requiring physicians to disclose *any* risks a patient deems to be material, is to overlook the potentially detrimental effect that informational overload can have on the individual's ability to make autonomous decisions, and potentially overestimating patients' ability to process that kind of information.⁷³

The failings of the subjective account in this regard might be claimed to lend support to a purely objective patient-centric account, which appeals not to what the particular patient deems material, but rather to what a hypothetical reasonable person would deem to be material. The way in which we ought to understand what constitutes such a hypothetical 'reasonable person' has been the subject of considerable debate. Briefly though, we may note two prominent interpretations outlined by

⁷³ It has also been claimed that the subjective element of the Montgomery ruling (discussed below) will encourage defensive medicine, by requiring doctors to disclose more information (to ensure their legal protection) than will actually facilitate their patient's decision-making. Chan et al., 'Montgomery and Informed Consent', are correct to point out that here the Montgomery ruling simply brings medical law into line with the GMC recommendations regarding the importance of communication. However, this simply means that a similar charge may be raised against these recommendations.

Dunn et al.⁷⁴ They note that one might interpret the reasonable person standard to mean that information ought to be disclosed simply if the *majority* of people think it ought to be disclosed in a specified set of circumstances. However, as the authors note, there is little support for this interpretation in medical law. Not only do judges lack empirical support for what most people think about the significance of different kinds of risks, but a significant reason for invoking the concept of reasonableness in one's criterion of materiality is to avoid reliance on capricious, and potentially unreasonable opinions held by the population at large.⁷⁵ Accordingly, Dunn et al. favour an interpretation of the 'reasonable person' as invoking a concept of 'reasonableness as normatively justifiable'.⁷⁶ On this interpretation, information ought to be disclosed if it concerns information that warrants weighting in a rational agent's decision-making process, by virtue of its significance for individual well-being.

I agree with these authors that the 'normatively justifiable' interpretation of reasonableness in the 'reasonable person' standard is the most convincing; indeed it is congruous with the concept of impersonal reason-giving facts that I discussed in Chapters 1 and 2. But the reasonable person standard, so construed, cannot alone provide a standard of disclosure that is sufficient for adequately facilitating patient autonomy. The problem with such a proposed standard is that a particular patient's preferences may seldom be exhausted by what is normatively justifiable in this way. Even if we assume that there are certain impersonal goods that all rational agents have reasons to pursue, to presume that material information is wholly constituted by information that a reasonable patient (in this sense) would want to know presumes that our self-interested reasons are exhausted by a certain set of our impersonal reasons. As I explored in Chapter 2, this sort of claim seems to implicitly assume an overly objective conception of well-being. Rational individuals can and do disagree about the weight of the reasons that they have to pursue different goods, and individuals can have personal self-interested reasons to instrumentally value certain goods which may not be shared by a hypothetical rational patient.

There are thus significant gaps in both the subjective and objective patient-centric standards of disclosure that suggest that they are each insufficient for adequately facilitating patient autonomy. Of course, an obvious answer to this problem is to combine the two approaches in a hybrid approach. This is just the approach taken by

⁷⁴ Dunn et al., 'Between the Reasonable and the Particular'. For other detailed discussions, see Miller and Perry, 'The Reasonable Person'; Gardner, 'The Many Faces of the Reasonable Person'; Kennedy, 'The Patient on the Clapham Omnibus'.

⁷⁵ Dunn et al., 'Between the Reasonable and the Particular', 8. The authors also consider an interpretation of the reasonable person criterion according to which 'it is reasonable to inform a patient of risk when there is logical coherence between the patient's values concerning risk and the patient's beliefs about the significance of the risk in these circumstances'. Their reason for dismissing this interpretation is that on this interpretation, the reasonable person criterion would simply serve as a constraint on the exercise of the 'particular patient' limb; yet, since the reasonable person criterion represents a separate limb to the Montgomery test of materiality, the authors conclude that this interpretation is not what the judges had in mind. Nonetheless, this interpretation could potentially serve as an objective standard in its own right.

⁷⁶ Dunn et al., 'Between the Reasonable and the Particular'.

the UK Supreme Court judgement on *Montgomery v Lanarkshire*.⁷⁷ Paragraph 87 of the judgement states that physicians are:

... under a duty to take reasonable care to ensure that the patient is aware of any material risks involved in any recommended treatment, and of any reasonable alternative or variant treatments. The test of materiality is whether, in the circumstances of the particular case, a reasonable person in the patient's position would be likely to attach significance to the risk, or the doctor is or should reasonably be aware that the particular patient would be likely to attach significance to it.⁷⁸

The Montgomery judgement thus offers a two-pronged disjunctive test of materiality that incorporates both objective and subjective patient-centric elements. The benefit of this hybrid approach is that it allows the judgement to overcome the respective insufficiencies of the purely subjective and purely objective patient-centric approaches noted above. Furthermore, we may notice that by appealing to what 'the doctor is or should reasonably be aware that the particular patient would be likely to attach significance to', the subjective limb of this test avoids concerns about the practicalities of Faden and Beauchamp's subjective condition of materiality. On the Montgomery standard, physicians can only be charged with negligence if they are aware, or could reasonably be aware, that a patient would attach significance to the information in question.

The Montgomery judgement has been described as a triumph of autonomy over paternalism, in that it epitomizes a shift of the balance in medical law regarding disclosure away from the paternalism of the Bolam approach, towards the protection of patient values.⁷⁹ However, in the remainder of this chapter, I want to suggest, from a philosophical perspective, that the manner in which the Montgomery ruling frames both the subjective and objective elements of its test of materiality is somewhat problematic, at least if the goal of the ruling is to facilitate patient autonomy.⁸⁰ This is not intended to be a criticism of the judgement as a legal ruling. Such rulings have to strike a careful balance between ethics, jurisprudence, and practical realities; in particular, it has to set the boundaries of the physician's liability to negligence. However, I shall argue that it might be possible to better capture the spirit of Montgomery, of prioritizing patient values, by reconceptualizing its two-pronged test of materiality.

⁷⁷ In the following I shall be interested in the Montgomery judgement's implications for risk disclosure. For a more general overview see Herring et al., 'Elbow Room for Best Practice?'

⁷⁸ *Montgomery (Appellant) v Lanarkshire Health Board (Respondent) (Scotland)*, paragraph 87. Emphasis added.

⁷⁹ Bolton, 'The Montgomery Ruling Extends Patient Autonomy'; Heywood, 'R.I.P. Sidaway'; Edozien, 'UK Law on Consent Finally Embraces the Prudent Patient Standard'; Farrell and Brazier, 'Not so New Directions in the Law of Consent?'

⁸⁰ Interestingly, Dunn et al. point out ways in which elements of the Montgomery judgement cannot be adequately grounded by an autonomy-based justification, and argue that its justification is better understood in terms of the value of patient-centric care. Dunn et al., 'Between the Reasonable and the Particular'. Whilst I agree with these authors that parts of the Montgomery judgement do not optimally facilitate autonomy, I believe that standards of disclosure have greater significance for patient autonomy (as opposed to patient-centric care alone) than these authors envisage, for reasons that I have outlined over the course of the last two chapters.

As I suggested at the end of the previous section, any standard of disclosure has to strike a delicate balance between providing too little information to adequately facilitate patients' autonomous decision-making, and providing too much for that purpose. At the same time, the standard also has to establish practically realistic boundaries of the doctor's duty of care. This is a tall order, and I suggested that the purely objective and subjective patient-centric accounts erred in the first way; they risk providing patients with insufficient information for autonomous decision-making. In order to avoid this error, a hybrid approach incorporating both objective and subjective elements is necessary. In the final section, I shall suggest that the manner in which the Montgomery phrases the subjective element of its hybrid approach means that it is in danger of erring in the second sense, of providing patients with too much information.⁸¹

5. Rational Materiality

The problem with the subjective limb of the Montgomery test is that it does not avoid the subjective assessment problem. True, it does stipulate *some* basis upon which the subjective assessment must be made; the patient must 'attach significance' to the information. However, it is far from clear that this is a suitable basis for identifying information that is material to the kind of rational decision-making that autonomy requires.

I shall explain this point in the remainder of this chapter. At the outset of this discussion though, I want to suggest an alternative test of materiality grounded in the theory that I have developed over the course of the book, one that draws an explicit link between the materiality of information, and rationality.

Rational Materiality: Information is material to a particular patient's treatment decision if the physician is, or should reasonably be aware that:

- (i) the patient's understanding of that information is necessary for adequately appreciating facts that are likely to give that patient considerable reasons to choose or reject a certain treatment option;
- and
- (ii) average human decision-makers would not be incapable of understanding and incorporating that information in a rational deliberative process.

Even before elaborating on this account of materiality in more detail, it should be clear that this account departs from the subjective account of materiality in two ways. First, by virtue of (i), certain information will be material to the patient's decision, *regardless* of the patient's own assessment of the materiality of that information. Furthermore, information that a patient mistakenly believes to be relevant to their treatment decision will not be material if it does not concern reason-implicating facts.

Criterion (ii) acknowledges that some information disclosure may be detrimental to autonomous decision-making, even if the information in question meets other

⁸¹ I shall follow Dunn et al. in interpreting 'reasonableness' to mean 'normatively justifiable' in the objective limb of the Montgomery judgement. Dunn et al., 'Between the Reasonable and the Particular'.

conditions of materiality. I shall explore this further below. Here though we may note that this criterion of materiality links my account to considerations of decision-making capacity, in so far as the latter requires the ability to understand, use, and weigh material information. My suggestion here, to be fleshed out below, is that an adequate test of materiality should acknowledge that the degree of capacity required to make a decision can be influenced by the nature and degree of the information that is deemed to be material to it.

First though, let us consider criterion (i). In so far as reason-giving facts can pertain to our impersonal reasons, the above test of materiality can also be understood to incorporate the objective limb of the Montgomery test, if reasonableness in that context is interpreted in terms of normative justifiability. However, by explicitly focusing on the rational content of the information to be disclosed, rather than the rationality of a hypothetical subject of that disclosure (as per Montgomery), the objective element of my account avoids some important ambiguities with the reasonable person criterion. For instance, even when we agree that the ‘reasonableness’ criterion in the objective limb of Montgomery should be interpreted in terms of normative justifiability, it remains ambiguous as to how we should understand the nature of the individual for whom the information must be normatively justifiable. Must it be normatively justifiable to a patient who is able to rationally process all and any relevant information relevant to well-being, or should we interpret it to mean a patient who is able to engage in rational deliberation, but who is nonetheless limited in their capacity to process such information? On the latter understanding, must the information be normatively justifiable to one who is *aware* of the limitation to their capacity to rationally process such information? My approach avoids these ambiguities by appealing to the reason-giving content of the information itself, whilst acknowledging (in the second criterion) ways in which disclosure can threaten an agent’s rational decision-making.

I will begin by exploring how this aspect of my test of materiality would apply in medical contexts, before explaining the manner in which it departs from the Montgomery judgement in further detail. On my proposed test of materiality, information pertaining to two broad aspects of a patient’s treatment decision will concern reason-implying facts. First, information pertaining to the *nature* of the proposed intervention will be material. For instance, the fact that an intervention will be painful or invasive provides patients with reasons not to choose that treatment (although these reasons will often not be decisive). This sort of information is also captured by the modal test that I outlined in previous chapters; patients must understand the nature of what they are consenting to if they are not to be modally precluded from acting effectively in pursuit of their ends when making treatment decisions. Second, facts pertaining to the probability of an intervention’s bringing about some outcome will also be reason-implying. Such facts will include not only those pertaining to the risks attending the intervention and possible side-effects, but also those pertaining to the probability of an intervention’s ameliorating the patient’s condition.⁸²

⁸² In their discussion, Herring et al. note that a positive aspect of the Montgomery judgement is that it emphasizes the disclosure of benefits as well as risks. Herring et al., ‘Elbow Room for Best Practice?’

I have thus far identified different *types* of information that will be material to a patient's decision. Consider now the *extent* to which patients should be made aware of these different aspects of their treatment. Whilst some information about the nature of the treatment might concern reason-implicating facts (for instance, the fact that the intervention is painful), a great deal of information about the treatment will not. For example, information concerning the exact biological mechanism that explains why an antibiotic helps to destroy a bacterial infection will normally not be material to a patient's decision to choose to take antibiotics; such information does not itself concern facts that provide agents with self-interested reasons. However, corollaries of that information may; for instance, in Jerry's case above, the fact that antibiotics are not an effective treatment for viral infections is material to her decision.

Although a great deal of specific information about the nature of a patient's condition or treatment options is unlikely to be material to their decision, it might be claimed that information concerning the foreseeable outcomes of their treatment options and their attendant possibilities is *always* likely to be material to a patient's decision. The Montgomery ruling comes close to making this sort of claim when it states that physicians need to make patients aware of *any* material risks of the proposed procedure, and that:

...the assessment of whether a risk is material cannot be reduced to percentages. The significance of a given risk is likely to reflect a variety of factors besides its magnitude: for example, the nature of the risk, the effect which its occurrence would have upon the life of the patient, the importance to the patient of the benefits sought to be achieved by the treatment, the alternatives available, and the risks involved in those alternatives. The assessment is therefore fact-sensitive, and sensitive also to the characteristics of the patient.

I argued above that in order to provide valid consent to *x*, the patient must understand that they are assuming *some* risk in agreeing to undergo *x*. That you are assuming a risk in performing some action can be a reason-giving fact. However, the strength of the reason it connotes depends on the gravity of the risk in question; very low risks generate very weak reasons.

The problem that this raises is that autonomous decision-making in a medical context can sometimes be threatened by the disclosure of such risks, even if they generate (weak) reasons. The disclosure of the nature and magnitude of very small risks can serve to hinder, rather than promote the patient's autonomy, because it feeds into well-known cognitive biases that serve to distort the agent's perception of the strength of the reason that they have. It is for this reason that I suggest that the rationalist account of materiality must incorporate criterion (ii) above, as I shall now explain.

If it were the case that patients were always able to understand the nature of small risks, and to incorporate them into a rational decision-making process that facilitated the pursuit of their own goals, then perhaps we should endorse the Montgomery ruling's implicit suggestion that physicians ought to disclose even minute risks, in order to increase the patient's autonomy with respect to their treatment decisions. However, research on cognitive biases suggests that patients are not able to compute information about risks in such an unbiased manner.⁸³ As Cass Sunstein points out,

⁸³ See Ingelfinger, 'Informed (but Uneducated) Consent'. See also Levy, 'Forced to Be Free?', and Sunstein, *Risk and Reason* for analyses of relevant empirical evidence.

when people have to make a decision in an emotionally charged context such as health care, they:

... tend to focus on the adverse outcome, not on its likelihood. That is, they are not closely attuned to the probability that harm will occur.⁸⁴

This is particularly problematic when patients are being asked to choose two potential means of achieving the same valuable end. To illustrate, consider the following:

Being fit or active has been associated with a greater than 50% reduction in risk of all-cause mortality.⁸⁵ However, vigorous exercise has also been associated with a very small acute risk of suffering a cardiac event during exercise; one study reported one such event per 2,897,057 person-hours of physical activity amongst healthy adults.⁸⁶

This is the so-called paradox of exercise.

Suppose a patient visits his doctor, and reports that he is petrified of suffering a heart attack after a friend recently died following one. The patient also has young children, and is terrified of not being able to see them grow up. The physician consults with the patient, and observes that he is at moderate risk of a heart attack. The patient is scared and says he will do anything to reduce his risk. The doctor could suggest that the patient modifies his diet; however, she notes that the patient's diet is not particularly bad, and modifying it is unlikely to lead to great improvement for this patient. The doctor could prescribe statins; however, she notes that these drugs may have some unpleasant side-effects.

She believes that the best way in which the patient can reduce his risk is to engage in an exercise program. However, she is also aware of the 'paradox of exercise', and that engaging in vigorous exercise will transiently cause an extremely small increase in the patient's acute risk of a heart attack, even though the long-term benefits to the patient's cardiac health far outweigh this small increase in their transient risk. Nonetheless, she prescribes an exercise plan without mentioning this risk, and tells the patient to take things slowly.⁸⁷

In this case, the doctor has grounds for believing that the patient would attach significance to the information about the small acute risk associated with bouts of vigorous exercise. The patient has stated that he will do anything to reduce his risk of a heart attack, and the information in question pertains to the transient risk of a cardiac event associated with the doctor's recommended course of action. The Montgomery standard thus seems to speak in favour of disclosure of this risk. The problem with this is that it is far from clear that the patient would take such information to be significant because of the strength of the reason it implies. Rather, the information is significant for the patient because of the particular emotional salience he attributes to the (highly unlikely) adverse outcome.

To exacerbate matters, the patient is likely to attribute even greater (yet unwarranted) salience to this risk simply by virtue of the fact that it has been disclosed to

⁸⁴ Sunstein, 'Probability Neglect', 62.

⁸⁵ Warburton, Nicol, and Bredin, 'Health Benefits of Physical Activity'.

⁸⁶ Malinow, McGarry, and Kuehl, 'Is Exercise Testing Indicated for Asymptomatic Active People?'; Thompson et al., 'Exercise and Acute Cardiovascular Events'.

⁸⁷ Maron, 'The Paradox of Exercise'; I discuss ethical implications of this paradox and exercise prescription in Pugh, Pugh, and Savulescu, 'Exercise Prescription and the Doctor's Duty of Non-Maleficence'.

them by an expert authority; 'if the doctor is telling me this, she must think it's important!' This may serve to radically distort the patient's perception of the strength of the reasons that the information is intended to convey, and instead lead the patient to make their decisions on an emotional response to that information that is not adequately grounded in the reality of their situation. The reality of the situation is that adhering to an exercise programme that slowly progresses towards more vigorous forms of exercise will significantly reduce their long-term risk of a heart attack, despite the small increase in the patient's acute risk that attends vigorous exercise for previously sedentary individuals.

Crucially, the doctor's reluctance to disclose in this case need not be born from a paternalistic motive of the sort that the Bolam approach enabled, and for which that standard was criticized. The doctor's concern about disclosure here may be grounded in doubts about whether the patient will use the information in a process of rational deliberation, and not simply in a concern that disclosure will lead the patient to choose contrary to what she believes to be medically indicated. Her concern may be that disclosing this information will not help the patient to make a decision that will help them best pursue their own goals and values. Instead, it might result in him forgoing exercise that will facilitate his pursuit of the goal he wants to achieve, in order to avoid a minute, but emotionally salient risk of the outcome that he wants to avoid.

The problem that this example raises is not simply that disclosing information about low levels of risk will harm patients by causing them psychological distress. If that were the case, the omission of this information could plausibly be justified by an appeal to the so-called therapeutic exception, an exception typically grounded by considerations of non-maleficence.⁸⁸ Rather the point that this example raises is that the provision of this information, and the unwarranted alarm it causes the patient, may actually be detrimental to their autonomy, their ability to make a rationally grounded decision, and to choose in accordance with what they actually value, regardless of whether or not the disclosure significantly harms the patient.

Accordingly, when thinking about disclosure and facilitating autonomous decision-making beyond the standard of voluntariness required for valid consent, we are faced with a significant tension. Extending the scope of disclosure requirements in accordance with the subjective limb of the Montgomery judgement will mean that patients are given more opportunities to make decisions about the risks they are and are not willing to take. However, increasing the opportunity for such decisions will only serve to increase patient autonomy if patients are able to make those decisions in an autonomous manner. This may sound almost trivially true, but it contains an important truth that the subjective limb of the Montgomery judgement overlooks: whilst our understanding of risks is sometimes crucial to making decisions about what we have most reason to do, the disclosure of precise details about risks can also play into the hands of our irrational biases. These biases can lead us to make

⁸⁸ The Montgomery judgement does include a clause allowing for the therapeutic exception, but notes that this is only justified if disclosure would be seriously detrimental to the patient's health. *Montgomery (Appellant) v Lanarkshire Health Board (Respondent) (Scotland)*, paragraph 88.

decisions that run contrary to what we believe we have reasons to pursue, and to acting in pursuit of our long-term plans.

How can we resolve this tension? It might be argued that we could seek to resolve it without adding further criteria of materiality to criterion (i) in the rational materiality test. For instance, perhaps we could take measures to partly mitigate some of the cognitive biases to which patients are subject. Perhaps if disclosure about a particular risk is to be worthwhile, patients should be helped to contextualize that risk to the kinds of risk they take in their day-to-day life.⁸⁹ Physicians should also be informed of the various cognitive biases that they and their patients are prone to exhibit, and should seek to mitigate their influence; this lends support to the dialogical approach that the Montgomery judgement endorses.⁹⁰ We might also advocate the introduction of ‘informed consent specialists’ who have received specialist training in human rationality, to act as ‘middle-men’ between the physician and their patient in complex cases.⁹¹

Each of these proposals has merit, but they are unlikely to act as a panacea solution for the types of irrationality that I have discussed here. Even where it is possible to try to mitigate the influence of these biases, it is not clear that such efforts will always be successful.⁹² Accordingly, rather than appeal to what patients (actual or hypothetical) attach significance to, as the Montgomery ruling implies, or even appealing to the rational content of the information alone, the test of materiality I have suggested maintains that we should acknowledge this tension and seek to resolve it by appealing to the underlying purpose of information disclosure in our test of materiality. If the purpose of information disclosure is to facilitate an individual’s ability to make decisions that are an accurate reflection of their evaluative judgements rather than their irrational biases, our decisions about what to disclose should be sensitive to the kinds of abilities that typical humans have.

Let me pre-empt to two potential objections to the objective phrasing of criterion (ii) above, which seeks to capture this thought. First, it might be argued that it is the individual patient’s ability to understand, weigh and use, information that is relevant in this context, rather than what standard human decision-makers are able to do. However, phrasing the criterion in subjective terms would entail a problematic circularity between definitions of materiality and decision-making capacity. The reason for this is that assessments of decision-making capacity are typically partly grounded by the individual’s ability to understand material information; accordingly, our definition of the latter cannot incorporate considerations of subjective capacity without circularity. An objective phrasing of criterion (ii) avoids this circularity; moreover, as I shall explain in response to the second objection, an objective phrasing captures a key element of the relationship between capacity and information disclosure.

A second objection to the objective phrasing is that it appears to contradict my earlier rejection of interpreting the reasonable person standard of disclosure as appealing to claims about what the majority of people think. However, there is no

⁸⁹ Sunstein, ‘Probability Neglect’, 92.

⁹⁰ Herring et al., ‘Elbow Room for Best Practice?’.

⁹¹ Levy, ‘Forced to Be Free?’, 299.

⁹² *Ibid.*, 297.

contradiction here; my criterion concerns what the majority of people are able to do, not what they *value*. The claim that I am making is that the information that we decide to disclose to patients should be information that most people are *able* to understand, weigh, and use; I am not claiming that our decision should be determined by what most people value.

In fact, the claim that I am making simply extends a point that is implicitly accepted by anyone who accepts the need for an account of materiality to limit the scope of disclosure requirements. There is a tight connection between the information that we deem to be material for making a decision, and the standard of decision-making capacity that will be relevant to that decision, since decision-making capacity requires the ability to understand, weigh, and use material information. It is for this reason that *full* understanding is not deemed to be necessary for decisional autonomy; to claim otherwise would be to preclude individuals from having decision-making capacity, since we humans are not capable of processing all of this information. This then, forms part of the motivation for a theory of materiality; we want to strike a balance between ensuring informed choice, and ensuring that most patients will be capable of understanding the information that we stipulate as being necessary for providing an autonomous authorization to treatment.

The point I am making here is that similar considerations arise with respect to information concerning very low risks; declaring this information material threatens to raise the standard of capacity required to make a decision beyond that which standard human decision-makers are capable.

Criterion (ii) broadly concurs with one professional standard of disclosure according to which decisions about disclosure should give consideration to features of the 'man on the Clapham omnibus'. However, it departs from this (flawed) standard by stressing that the relevant consideration here is what the man on the bus is capable of, and not what he values. Let me now turn to consider the implications of this criterion. How can we tell whether average human decision-makers are capable of understanding, weighing, and using a particular piece of information? This is ultimately an empirical question; however, a rough plausible heuristic here may be that we should tailor our disclosure about risk to the strength of the reason it connotes, on the basis of the assumption that cognitive biases may plausibly have less of a hold on us when we are weighing considerations that we take to have considerable reasoning force. Explicitly attending to strong reasons is one way in which we can reduce the influence of these biases.

The strength of a reason associated with information pertaining to a particular outcome of a medical procedure is a function of the (dis)value of the outcome it concerns, and its likelihood. The higher the (dis)value, and the more likely it is to occur, the stronger the reason. However, since risk is a comparative concept,⁹³ when we are thinking about the strength of an all things considered reason to assume a certain risk in undergoing a medical procedure, we must weigh both the disvalue and probability of the risked outcome, against the probability and value of the hoped-for

⁹³ Herring et al., 'Elbow Room for Best Practice?', 10.

outcome of the procedure which provides the rational justification for the assumption of risk.

Indeed, the Montgomery judgement implicitly reflects this in stressing that the material benefits of alternative procedures must be disclosed in addition to the material risks.⁹⁴ To illustrate, the strength of your reason provided by a 2 per cent mortality risk of a medical intervention depends on what you stand to benefit by assuming that risk in undergoing the intervention. This risk might give Patient A a comparatively weak reason not to undergo surgery, if the risk is posed by a surgery that is necessary for saving her life. In contrast, the same degree of risk might give patient B a much stronger comparative reason to not undergo the surgery, if the surgery is an elective procedure that is less valuable for that patient, such as the alleviation of considerable, but tolerable pain. Finally, Patient C might have very strong comparative reasons not to expose themselves to even small risks of catastrophic outcomes in order to undergo cosmetic surgery.

For this reason, my proposal does not entail that physicians should never disclose information about low risks, or that patients cannot incorporate such information into a rational decision-making process. The point is that risks of the same probability can imply reasons of different comparative strength in different contexts, and decisions about information disclosure should be sensitive to this point. We may also notice that although the likelihood with which an outcome will occur is an objective matter, the strength of some of the reasons represented by facts about aspects of particular treatments will depend significantly on the patient's *values*. More specifically, it will depend on the weight that the patient assigns to the reasons they have to pursue different goods. In order to determine whether information about a low degree of risk represents a reason that is sufficiently strong to be deemed material to the patient's decision, physicians must thus be aware of the patient's own values, and their own general attitudes towards risk. This does not imply that the responsibility for facilitating the patient's autonomous decision-making lies solely with the physician. On the contrary, the patient is in a far better epistemic position with regards to her own values; she has an important role to play in facilitating her own ability to make autonomous decisions by engaging with the physician about the values that are central to how she wants to live her life over the course of the clinical encounter, enabling the physician to tailor the information they disclose. Patients thus share in the responsibility to facilitate their own autonomy; as I argued in the previous chapter, this includes a degree of doxastic responsibility.⁹⁵

The disclosure of material information pertaining to impersonal reasons should thus serve as a *starting point* of the physician's disclosure; patients should be made aware of the salient aspects about the nature of their procedure, and also reminded that medical procedures, like many other activities in life, unavoidably involve the assumption of some low risks. However, disclosure concerning the precise nature and degree of risks that connote very weak reasons can serve to hinder rather than facilitate the patient's autonomy.

⁹⁴ *Montgomery (Appellant) v Lanarkshire Health Board (Respondent) (Scotland)*; Herring et al., 'Elbow Room for Best Practice?', 5.

⁹⁵ Foster, *Choosing Life, Choosing Death*, 104; Kukla, 'How Do Patients Know?'

As such, an adequate informed consent procedure requires that the physician and patient engage with each other to establish the patient's values and their attitudes towards risk, and in doing so conjointly establish the threshold level of strength of reason that a risk would have to connote, in order to warrant full disclosure. Further disclosure of risk should be justified by the decision reached following such a dialogue. For this reason, it makes little sense to stipulate that risks of a certain percentage probability must always (or never) be disclosed; decisions about disclosure of risk have to be sensitive not just to the nature of the outcome to which they pertain, but also to how and in what way that outcome matters for the particular patient in question.

In the exercise case above, because the doctor is aware that the overall goal that the patient places value on is avoiding a heart attack in order to be there for his young family, it is clear that the transient minuscule increase in acute risk of a cardiac event associated with exercise connotes an extremely weak reason for that patient. It is a risk that is clearly outweighed by the fact that the exercise programme will significantly decrease their long-term risk of such an event. Of course, as the degree of risk associated with a procedure under consideration increases, so too will the strength of the corresponding reason; in such cases, physicians may be required to engage further with their patient about the kinds of risk that they are willing to accept to achieve certain outcomes. Rational agents can come to different conclusions with regards to the appropriate attitude to take towards risk. Notwithstanding this point, discussions about attitudes towards risk need to be placed in the context of the kinds of risk to which the patient is already exposed, by virtue of her medical condition, and also those that she encounters simply by carrying out everyday activities.

The Montgomery ruling also advocates the importance of dialogue, in claiming that the doctor should act in an advisory role, and that this involves engaging in a dialogue with the patient that aims to:

... ensure that the patient understands the seriousness of her condition, and the anticipated benefits and risks of the proposed treatment and any reasonable alternatives, so that she is then in a position to make an informed decision.⁹⁶

However, the account of autonomy that I have defended supports a stronger approach to this dialogue before the disclosure of risk. In order for informed consent to become a truly two-way informational transaction, the physician should not be limited to merely facilitating understanding. Rather, facilitating the individual's autonomy requires the elicitation and defence of the patient's values, and the physician advocating their own view, drawing on their own professional experience, about the kinds of risk disclosure that will facilitate the patient's autonomous decision-making.⁹⁷

In abandoning the Bolam test, the Montgomery ruling denied the significance of the medical professional body's expertise in delimiting the scope of material

⁹⁶ *Montgomery (Appellant) v Lanarkshire Health Board (Respondent) (Scotland)*, at paragraph 90.

⁹⁷ Savulescu, 'Liberal Rationalism and Medical Decision-Making'. Decision aids can be a valuable tool in this model of the doctor-patient relationship. See O'Connor et al., 'Decision Aids for Patients Facing Health Treatment or Screening Decisions'.

information.⁹⁸ However, obviating *any* reference to medical opinion in thinking about disclosure fails to acknowledge that the medical profession does have something significant to offer to the facilitation of autonomous decision-making. Medical professionals have a great deal of experience in helping patients in medical contexts, and are aware of the effects that information disclosure can have on patient decision-making. The rationalist approach to materiality that I have developed here allows for this experience to be brought to bear on disclosure decisions by establishing a threshold strength of reason that is appropriate for a particular patient, without the professional body determining the boundaries of materiality in the overly objective manner of the Bolam test.

It might be objected that the approach I have advocated here is impractical, given the lack of resources available to health services. As I acknowledged at the outset of this discussion, my intention here has not been to criticize the Montgomery judgement as a legal instrument that has to balance a plethora of ethical, jurisprudential, and practical considerations. Rather, my aim here has been to use the Montgomery judgement as a legal model of disclosure that serves as the best starting point to think about the effects of disclosure for autonomous decision-making. The approach that I have outlined is naturally somewhat divorced from practical realities in a way that the Montgomery judgement cannot afford to be. The model outlined in the Montgomery judgement requires more time and resources than simply giving the patient a consent form to read and sign, and my proposed model arguably goes further still.

If resources do not permit this approach, this gives us strong reason to refrain from treating my approach as necessary to informed consent in the institutional sense outlined at the beginning of this chapter. Moreover, doctors should clearly not face sanction for failing to meet this standard if the resources do not allow them to spend the sort of time with their patient that this approach requires. However, this concern does not speak against my approach as the standard to be met in order for consent to play an operative role in facilitating autonomous authorization of treatment. Existing mechanisms are notoriously poor at ensuring adequate levels of understanding amongst patients,⁹⁹ and facilitating rational decision-making. We have reasons to try to improve upon this, and one step on the path to doing so is to think about how we could meet the challenges of facilitating autonomous decision-making *without* considerations of restrained resources. This is not to say that there are not other morally relevant considerations in play when we think about whether it would be justifiable to spend more scarce resources in order to better facilitate autonomous decision-making in the patient population, over other worthy goods.¹⁰⁰ These are important and difficult questions. However, these are questions about the weight we should attribute to different values in health care, and not simply a discussion about the nature of autonomy per se.

⁹⁸ Herring et al., 'Elbow Room for Best Practice?', 4.

⁹⁹ Flory and Emanuel, 'Interventions to Improve Research Participants' Understanding in Informed Consent for Research'.

¹⁰⁰ Chan et al., 'Montgomery and Informed Consent'.

Conclusion

In this chapter, I have begun to situate my rationalist account of autonomy in our understanding of the structure and justification of informed consent requirements. Having considered the cognitive element of decisional autonomy and its implications for understanding, and standards of materiality and disclosure, I shall in the following chapter continue my investigation of rationalist informed consent by considering the question of decision-making capacity. In doing so, I shall seek to respond to the demandingness objection that I highlighted in section 2 of this chapter.