

Part III:

Informed consent and artificial intelligence in medicine

Chapter 6: UK tort law

This part asks whether the law imposes an obligation on relevant parties to obtain the consent of patients to the involvement of artificial intelligence (AI) in their care, to inform patients of this involvement and, crucially, to advise them of relevant features of the technology in such a way that meets the autonomy challenges outlined in Chapter 3.

The most direct way in which these questions arise under UK law is *via* the common law causes of action of battery and negligence. These have gained an established role in protecting patient autonomy, either through imposing consent requirements on most forms of medical treatment or through recognising legal obligations to provide patients with information that is deemed relevant to clinical decision making.⁷³⁷

⁷³⁷ This combined role is often demarcated by the use of the more general label of ‘informed consent’. The English courts initially proved reluctant to adopt this American terminology and it was disputed whether, if it were to find application, it would be more suited to denote liability in battery or negligence. Compare *Freeman v Home Office (No 2)* [1984] QB 524, 555 with the statement in *Davis v Barking, Havering and Brentwood Health Authority* [1992] Lexis Citation 2495, that ‘the concept of “informed consent” cannot be made to fit the cause of action for negligence; its place lies in trespass to the person’. Contrast, the recent statement in the context of negligence that: ‘Whatever uncertainty there may have been in the past, the requirement of informed consent to medical treatment is now a fundamental and settled principle of the law in both England and Wales and Scotland’: *Gallardo v Imperial College Healthcare NHS Trust* [2017] EWHC 3147 (QB), [2017] 12 WLUK 198 [1]. For an analysis of the ‘informed consent’ terminology and the law of negligence see: Maclean, ‘The Doctrine of Informed Consent’ (2004) 24(3) Legal Studies p. 386.

Both of these mechanisms are protected through the imposition of *ex post* sanctions on those who breach their legal duties.⁷³⁸ If a patient can make out either tort, then they are entitled to monetary damages. Courts may also be asked to make a declaration on whether a given clinical procedure or type of intervention would constitute a battery and thus whether the medical professionals concerned are legally permitted to proceed with the relevant actions.⁷³⁹ This will technically proceed as an administrative law procedure, but the claim will turn on whether the substantive requirements of the relevant civil law claim are made out.

Following the evaluation of the common law, a specific aspect of the *UK General Data Protection Regulation* and the *Data Protection Act 2018* will also be addressed. Admittedly this falls outside of the immediate remit laid down in Chapter 1, which focussed on the outlined causes of action that have shaped the evolution of this area of the law. It is nevertheless examined as it marks a targeted legislative intervention, which addresses one specific shortcoming in the common law analysis of ML instrumentalisation. In so far as this approach is to be critiqued in the final chapter, it must be done with an understanding of the functioning of this statutory mechanism.

I. Battery

Consent functions to legitimise an act that would otherwise constitute a battery – a criminal offence and a civil wrong under English law.⁷⁴⁰ Judicial pronouncements on battery have established the necessity of consent for most forms of medical treatment and the conditions under which such

738 Battery also constitutes a criminal offence, as some of the cited cases *infra* demonstrate. However, as the mechanisms' requirements are largely consistent, our primary framing will be under the law of tort.

739 This may be at the level of an individual decision: *Re F (mental patient: sterilisation)* [1990] 2 AC 1; *Airedale NHS Trust v Bland* [1993] AC 789. The public law procedure of judicial review may also be utilised to request a formal declaration from the court stating what a relevant private law mechanism requires in certain circumstances. This may affect the legality of a policy to adopt a medical intervention more widely. For a recent summary of the appropriate deployment of this approach see: *Bell v Tavistock and Portman NHS Foundation Trust* [2021] EWCA Civ 1363, [2022] 1 All ER 416 [66]-[90].

740 *Chatterton v Gerson* [1981] QB 432, 442.

consent is valid.⁷⁴¹ These will be drawn upon to determine the information that must be shared with a patient to ensure that the consent they have given to a procedure involving AI is valid.

A. Limitations flowing from the battery doctrine

Before engaging with these specific questions, an assessment of the tort's broader requirements is indispensable in order to understand the limitations upon its operation in the healthcare sphere. To some degree the courts have shown flexibility in this respect, accommodating the unique demands of the clinical situation. For instance, the usual restrictions placed on an individual's ability to consent to certain types of bodily injury have been relaxed in so far as 'reasonable surgical interference' is concerned.⁷⁴² Nevertheless, other conditions stemming from the tort's origin as a mechanism for the prevention of violence persist, without providing much scope for argumentation.⁷⁴³

Most especially, the clinical interaction complained of must involve some form of direct or immediate interference with the claimant's body.⁷⁴⁴ This severely limits the role that battery's consent requirement can play in facilitating patient control of AI use in their treatment. Not only will consent not be pertinent to many decisions not to treat, which involve no relevant contact, but battery's consent and information requirements will likely not be applicable to indirect forms of treatment, such as the prescription of drugs.⁷⁴⁵ Therefore, while some such decisions are likely to benefit from the assistance of intelligent decision-support systems, battery provides no basis on which the doctor is required to inform the patient or obtain their consent for this use.

741 Of course, there are instances where consent cannot be obtained and other justifications are in place to allow for the emergency treatment of an unconscious patient, but these are less relevant for our analysis of AI autonomy violations.

742 *Attorney-General's Reference (No. 6 of 1980)* [1981] QB 715, 719.

743 Brazier and Lobjoit in Erin and Bennett, *HIV and AIDS: Testing, Screening, and Confidentiality* (2001) 185-186.

744 *Scott v Shepherd* (1773) 3 Wils KB 403.

745 Maclean does construct a well-reasoned argument as to why battery should be applicable to drug prescriptions, but he notes how unlikely this is to succeed in English courts: Maclean, *Autonomy, Informed Consent and Medical Law: A Relational Challenge* (2009) 150-152. See also the analysis of Seabourne who reaches a similar conclusion: Seabourne, 'The Role of the Tort of Battery in Medical Law' (1995) 24(3) *Anglo-American Law Review* p. 265, 270-271.

Battery further requires an intentional action by the defendant.⁷⁴⁶ In the past this had generated some uncertainty: must the relevant intention be directed towards inflicting an injury on another? Or is it sufficient that the application of force (the touching) was envisaged?⁷⁴⁷ However, as has been settled in *Wilson v Pringle*, the latter interpretation is the correct one: ‘It is the act and not the injury which must be intentional. An intention to injure is not essential to an action for trespass to the person. It is the mere trespass by itself which is the offence’.⁷⁴⁸ A patient’s action will not be barred under battery because a professional’s intention did not go beyond the touching.

There is also a potential mandate under UK law that, although no intention to injure is required, the behaviour of the defendant must be hostile in nature.⁷⁴⁹ This factor would almost certainly exclude ordinary clinical situations (including those involving AI) from its scope.⁷⁵⁰ Yet it is only a ‘potential requirement’, since its meaning is extremely ill-defined and several judgments have strongly indicated, or have even been decided on the basis, that it has been abrogated, although it is yet to be directly overruled.⁷⁵¹ As requiring hostility in medical interactions would effectively bar the battery action from offering any meaningful protection to patient autonomy, this principle is argued to provide a further, compelling reason for rejecting the questionable validity of this requirement.

In sum, before one even begins to analyse the nature of battery’s consent and information requirements, there are broader difficulties that limit the application of this norm to the use of AI in medical treatment. UK law has arguably maintained an emphasis on rule-specific factors that are separate from, and serve as restrictions upon, the autonomy principle. Most significantly, it is relatively certain that, if a patient is not touched during AI use

746 ‘The least touching of another in anger is a battery’ represents an early indication of this requirement: *Cole v Turner* (1704) 6 Mod 149.

747 *Wilson v Pringle* [1987] QB 237, 248-249.

748 *ibid* 249.

749 *ibid* 246-248.

750 ‘[T]he requirement of hostility, (...) unless it is interpreted in a sense so weak that it collapses into the question of intentional action, is likely always to be a stumbling block to plaintiffs in medical battery’: Seabourne, ‘The Role of the Tort of Battery in Medical Law’ (1995) 24(3) *Anglo-American Law Review* p. 265, 272.

751 Lord Goff’s well-known rejection of this requirement in *Re F* was *obiter dicta* (although it is interesting that this rejection was specifically based on the ‘libertarian principle of self-determination’ in medical treatment): *Re F (mental patient: sterilisation)* [1990] 2 AC 1, 72-73. See also the total omission of this requirement in *Re B (adult: refusal of medical treatment)* [2002] EWHC 429 (Fam), [2002] 2 All ER 449.

or with a view to carrying out an AI intervention, then the professional will have no obligation to advise the patient of their reliance on this technology or its implications.

B. Battery and the nature of valid consent

If the above requirements are fulfilled, then the claimant must next establish that they did not consent to the touching.⁷⁵² Where no consent whatsoever has been given, this claim can be straightforwardly made out.⁷⁵³ This situation is rare, however. Moreover, where an AI/ML device is involved, it has been argued that there is generally a wider intervention to which the patient has agreed. To ground a successful battery action in such circumstances it is necessary to determine a deficiency in the claimant's given consent: either its scope was too limited or, relatedly, it was given without knowing about a certain dimension of the clinical intervention.

The doctrine that shapes this area of the law emerged from *Chatterton v Gerson*, where Bristow J held that consent to a medical treatment must not only be in form but in reality.⁷⁵⁴ This required the patient to be 'informed in broad terms of the nature of the procedure'.⁷⁵⁵ A similar definition of the relevant test can be found in *Re T*, where it was said to be enough if the 'patient knew in broad terms the nature and effect of the procedure to which consent (or refusal) was given'.⁷⁵⁶

While providing the patient with sufficient information on their treatment is one important prerequisite for the validity of their consent, it is clear from such statements that battery is intended to set a low bar for disclosure. Information will not have to be provided on many specific features of medical interventions. For many of these subsidiary aspects the

752 While the Court of Appeal did not address this matter, McCowan J had agreed with counsel at first instance on this point: *Freeman v Home Office (No 2)* [1984] QB 524, 539. For a critique of this position and an overview of how Canada and Australia have framed the matter differently see: McHale in Laing and McHale, *Principles of Medical Law* (Fourth Edition 2017) 424.

753 *Border v Lewisham and Greenwich NHS Trust* [2015] EWCA Civ 8, (2015) 143 BMLR 182 [19]. See also *Cull v Butler* where a hysterectomy was carried out against the express wishes of the patient: *Cull v Butler* [1932] 1 BMJ 1195.

754 *Chatterton v Gerson* [1981] QB 432, 442-443.

755 *ibid* 443.

756 *Re T (adult: refusal of treatment)* [1993] Fam 95, 115.

courts explicitly state a preference for a claim under the negligence cause of action.⁷⁵⁷

Weighty rule-specific considerations are used to support this approach. Foremost among them is a distaste for subjecting medical professionals acting in good faith to this form of liability.⁷⁵⁸ However, commentators have also drawn the logical inference that in those cases where a battery claim is nonetheless successful the courts are responding to a particularly grievous interference with patient autonomy.⁷⁵⁹ That is, the weight of the principle is great enough in those situations to direct legal reasoning towards offering a strong response. Such a differentiated approach further corresponds to the claims made in the theoretical analysis in Chapter 3.

The remainder of this section categorises the significant deficiencies in consent which have given rise to a battery action, yielding three classes: the nature of the procedure, the identity of the actor and the purpose or motivation for the act. Shortcomings in the patient's understanding of these dimensions have proved sufficiently substantial so as to fall outside of their broad consent to diagnosis or treatment. Bearing in mind that AI autonomy violations must be of a comparable gravity to those found in the existing analyses, we turn to evaluate whether any of AI's challenges can be subsumed under these classes.

1. Nature of the procedure

The first class of cases is constituted by a change in the nature of the procedure. This was the type of case explicitly elaborated and examined in *Chatterton*. On its face, the judgment appears tailored to ensure that recourse to a claim in battery is only available for clear-cut misunderstandings of the fundamental physical nature of the procedure. Thus, the court in *Chatterton* referred to the case of a boy who is circumcised instead of re-

757 *Chatterton v Gerson* [1981] QB 432, 442; *Hills v Potter* [1984] 1 WLR 641, 653; *The Creutzfeldt-Jakob Disease Litigation* (1995) 54 BMLR 1, 4-6; *Davis v Barking, Havering and Brentwood Health Authority* [1992] Lexis Citation 2495.

758 '[I]t would be deplorable to base the law in medical cases of this kind on the torts of assault and battery': *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 883.

759 Maclean, *Autonomy, Informed Consent and Medical Law: A Relational Challenge* (2009) 195-196.

ceiving an expected tonsillectomy.⁷⁶⁰ Similarly, in *Devi v West Midlands Regional Health Authority* the sterilisation of a woman was conducted without prior discussion during a postpartum dilation and curettage operation.⁷⁶¹ The Court of Appeal accepted that the action for damages proceeded on the basis of an ‘assault’, even if the issue was not expressly dealt with. The sterilisation was treated as altering the nature of the procedure to such a degree that the patient’s consent had been violated.⁷⁶²

One should also consider *Davis v Barking Havering and Brentwood Health Authority*. Here the court asked whether an aspect of a wider procedure required separate ‘sectional’ consent.⁷⁶³ This is arguably another way of framing the issue of whether a patient knew enough about the overall nature of what was to be done, for their consent to be valid.⁷⁶⁴ However, by enquiring as to the necessity of separate consent, McCullough J provided more nuanced insights into battery’s analysis of this question. In particular, while accepting that separate consent must be obtained in some instances, for example to separate surgeries, the judge objected forcefully to the prospect of complex procedures being broken down into their individual steps and consent being demanded for each of these. Such an approach would give battery an unduly prominent role in medical law. Rather, these were ‘details’ that did not go to the validity of consent.⁷⁶⁵

Obiter statements in the judgment further conveyed the kinds of acts that would be considered to be aspects of a single procedure: the attachment of an ECG to the chest of an unconscious patient, the insertion of a tube into the unconscious patient’s trachea and an injection of morphine while the patient is coming around after the surgery.⁷⁶⁶ *Davis* itself was decided on the basis that an injection of a local (caudal) anaesthetic could not be separated from the patient’s consent to the administration of a general anaesthetic, even though the former carried specific risks that she had not

760 *Chatterton v Gerson* [1981] QB 432, 443.

761 *Devi v West Midlands Regional Health Authority* [1981] Lexis Citation 1417.

762 Cf. *Abbas v Kenney* (1995) 31 BMLR 157.

763 *Davis v Barking, Havering and Brentwood Health Authority* [1992] Lexis Citation 2495.

764 Maclean, ‘Consent, Sectionalisation and the Concept of a Medical Procedure’ (2002) 28(4) *Journal of Medical Ethics* p. 249.

765 *Davis v Barking, Havering and Brentwood Health Authority* [1992] Lexis Citation 2495.

766 *ibid.*

anticipated.⁷⁶⁷ A comparison can also be drawn to *R v Mental Health Commission, ex p X*, where the court relied on the common law's delineation of the concept of consent. It was held that a patient must understand the likely effects of a treatment, but that this must not amount to a detailed understanding of physiological processes.⁷⁶⁸

All in all, these cases suggest that an understanding of the fundamental physical components of a procedure are sufficient for valid consent. It is not necessary to know about all the processes involved in an intervention, whether these have subsidiary physical manifestations or not. Moreover, the UK courts have been clear that the disclosure of the specific risks of a procedure is not something that can properly ground a claim in battery.⁷⁶⁹ Failing to disclose these matters will not vitiate consent.

In so far as ML devices are concerned, it is envisaged that their operation will be closely connected to methods of diagnosis, treatment and prognosis (either human or technological) and the physical manifestation of the procedure will, for the most part, remain singular, apparent and unchanged. A general disclosure of AI use, treating it as a separate procedure, therefore cannot be maintained. Furthermore, although it has been argued that AI may lead to changes in the general risk profile of an intervention, these are matters to properly be considered under negligence.

At the same time, the basis for the courts' delineation of a procedure's nature remains ambiguous and there are indications that it is not entirely restricted to its physical components. Arguably, the effects of an intervention, with their wider significance, must also be conveyed. In *R v Mental Health Commission, ex p X* Stuart-Smith LJ mandated that real consent entails that a patient is aware of the likely effect of a treatment, which here involved the 'effect of reducing male sexual drive', and as stated did not depend upon an awareness of the underlying physiological processes.⁷⁷⁰ This more general association between the nature of an intervention and the

767 Grubb, 'Battery and Administration of Anaesthetic: *Davis v. Barking, Havering and Brentwood Health Authority*' (1993) 1(3) Medical Law Review p. 389.

768 *R v Mental Health Commission, ex p X* (1988) 9 BMLR 77, 86-87.

769 *R v Richardson* [1999] QB 444, 450; *Hills v Potter* [1984] 1 WLR 641, 653; *Chatterton v Gerson* [1981] QB 432, 442-443. Consider also *The Creutzfeldt-Jakob Disease Litigation* as a particularly strong example of this position. There appeared to be absolutely no disclosure of any risk of the clinical intervention, but this was still not seen as a basis for a battery claim: *The Creutzfeldt-Jakob Disease Litigation* (1995) 54 BMLR 1, 3-4.

770 *R v Mental Health Commission, ex p X* (1988) 9 BMLR 77, 83, 86.

wider non-physical significance to the patient can also be found in *Abbas v Kenney*. Here the court considered that the consent to a procedure which removed the patient's reproductive organs was valid, when the patient had agreed to 'the least extensive surgery possible, consistent with saving her life' and 'understood that [the surgeon's] main brief was to save her life, and if possible to save her fertility'.⁷⁷¹ Although the physical manifestations of the surgery were clearly relevant, it was primarily important that the patient was aware of the significance of the procedure for the purposes of saving her life and for maintaining her fertility.

The resulting uncertainty has allowed commentators to argue that the significant implications of a procedure for the patient, shape its nature. This has been extended *inter alia* to the effects of diagnostic tests. A particularly vigorous debate has addressed the question of whether an analysis of a patient's blood for its HIV status falls within a patient's general consent to 'some blood tests'.⁷⁷² The physical nature of this interaction remains unchanged whether the additional analysis takes place or not.⁷⁷³ Moreover, the fact 'that one of the tests is for HIV does not alter the general nature of the procedure as part of a process of therapeutic diagnosis'.⁷⁷⁴ The focus on the physical manifestations of the intervention, as well as the wider clinical processes underlying it, find support in the bulk of the case law analysed above.

771 *Abbas v Kenney* (1995) 31 BMLR 157, 165-166.

772 To my knowledge no case dealing with this issue ever came before the UK courts. Yet the Public Health Laboratory Service did carry out HIV tests on anonymised blood samples, even though they had been obtained from patients for other purposes: Coghlan, 'Could HIV tests land doctors in court?' (19.1.1994) <<https://www.newscientist.com/article/mg14119100-600-could-hiv-tests-land-doctors-in-court/>> accessed 9.3.2021; Brazier and Lobjoit in Erin and Bennett, *HIV and AIDS* (2001). It appears to have been the fact that several professional bodies sought legal advice on the matter which sparked the wide-ranging academic debate: Sherrard and Gatt, 'Human Immunodeficiency Virus (HIV) Antibody Testing: Guidance from an Opinion Provided for the British Medical Association' (1987) 295(6603) *British Medical Journal* p. 911; Kennedy and Grubb, 'Testing for HIV Infection: The Legal Framework' (1989) 86(7) *Law Society Gazette* 30-35; Keown, 'The Ashes of Aids and the Phoenix of Informed Consent' (1989) 52(6) *The Modern Law Review* p. 790; Grubb and Pearl, *Blood Testing, AIDS, and DNA Profiling: Law and Policy* (1990).

773 Grubb and Pearl, *Blood Testing, AIDS, and DNA Profiling: Law and Policy* (1990) 6.

774 Keown, 'The Ashes of Aids and the Phoenix of Informed Consent' (1989) 52(6) *The Modern Law Review* p. 790, 796.

On the other hand, it may be thought that the patient must have knowledge 'of the underlying purpose of the procedure (ie HIV testing) which is vital in order for the patient to have the necessary understanding of the *quality* of the touching'.⁷⁷⁵ This appeals to the courts' reference to the wider significance and effects of a procedure. To concretise their argument for disclosure, its proponents have exhibited a pronounced reliance on the differential implications for patient autonomy. Kennedy and Grubb state explicitly that 'the consequences for a patient that flow from a positive HIV test (or in some instances even a negative test) are so serious for him that a court would consider it to be contrary to public policy to regard consent to testing as including a procedure with such far reaching implications for the patient'.⁷⁷⁶ Grubb and Pearl also highlight the 'grave and adverse personal and social consequences' that accompany a diagnosis of HIV. Further they cite: the widespread discrimination that such a finding may result in (especially since it cannot always be ensured that confidentiality will be maintained in practice), the effect on the patient's ability to obtain life insurance coverage and the fact that the patient would be deprived of the opportunity to receive appropriate counselling before their diagnosis.⁷⁷⁷

All of these factors indicate that undertaking a HIV diagnosis without the patient's understanding seriously impairs their autonomy. As explored in Chapter 3, it determines an aspect of their medical care, prevents them from reconsidering their position in a reflective manner and directs their care according to external non-subjective values. Invoking these dimensions of autonomy to ascertain the significance of a procedure, is what subsumes the hypothetical HIV-testing scenario under the existing case law.

Moreover, as Keown has argued: if consent is found necessary for HIV-testing, then it would have far-reaching consequences. For instance, one would expect it to be required for many other tests, such as those used in the diagnosis of cancer.⁷⁷⁸ This is what, in turn, connects this line of argumentation with our analysis of ML use.

As noted in Chapters 2 and 3, ML devices can be used to pursue broad, ill-defined goals and in a way that is partially determinative of the clinic-

775 Grubb and Pearl, *Blood Testing, AIDS, and DNA Profiling: Law and Policy* (1990) 6.

776 Kennedy and Grubb, 'Testing for HIV Infection' (1989) 86(7) *Law Society Gazette* 30-35.

777 Grubb and Pearl, *Blood Testing, AIDS, and DNA Profiling: Law and Policy* (1990) 6-7.

778 Keown, 'The Ashes of Aids and the Phoenix of Informed Consent' (1989) 52(6) *The Modern Law Review* p. 790, 797.

al decision. The most relevant example that was given, is AI's ability to reach serious, surprising diagnoses. Without more information, it will not be obvious to the patient that a relevant evaluation is intended or what condition(s) will be tested for. In doing so, the ML device may be striving to accomplish an end that is not the patient's own and on which they have had no influence.

It can even be argued that the autonomy violation is more serious in the novel, AI scenario. The machine does not only provide a positive/negative result for one well-defined condition but may run a more complex diagnosis for a number of illnesses. As such, it is capable of generating a range of surprising insights and the infringement of the patient's autonomy is greater because they are deprived of the opportunity to understand and control a decision that is substantively wider and more nuanced. In such situations there is a strengthened argument that the nature of the procedure is changed by the reliance of the ML device.

A further factor that has emerged from the discussions surrounding blood analysis, and which should certainly not be discounted, concerns the timing of the doctor's intention to proceed with HIV testing. Regardless of one's understanding of the nature of the procedure, this must be determined at the time of the direct and intentional intervention.⁷⁷⁹ Therefore, if AI use is to vitiate the patient's consent for the reasons outlined above, then the doctor who is interacting with the patient must have formed an intention to proceed with an AI analysis before engaging in the requisite contact.

This will place many instances of AI use outside of battery's scope. For instance, although pathology represents a field where the prospects for implementing AI are particularly promising, pathologists would generally have no contact with the patient and would only form an intention to analyse material *via* a certain methodology involving AI after it has already been obtained by another professional. In other scenarios it will also not be clear that AI use was contemplated before contact – a subsequent analysis of the patient's data may be an easy thing, requiring no additional testing. Even if a claim will not be barred here, this requirement will nevertheless create evidential difficulties for the patient.

In sum, it is possible to claim that subjecting a patient to certain unsolicited AI analyses can alter the nature of a procedure and vitiate the validity of their consent. These are analyses where the AI has a degree of independ-

779 Grubb and Pearl, *Blood Testing, AIDS, and DNA Profiling: Law and Policy* (1990) 21.

ence and is capable of arriving at conclusions with significant implications for the patient. The argument from the case law, in combination with the autonomy principle, has been stated to be stronger than that in the established debate on HIV testing. Nevertheless, the argument must remain of uncertain strength given the limited jurisprudence on the matter and the substantial ambiguity in the precedents there are.

2. Identity of the professional

Another class of cases indicates that a patient's consent can be invalidated if they are not informed of the identity of a (purported) professional carrying out a procedure on them.⁷⁸⁰ In general, these are not cases where one person has pretended to be someone else or where one doctor performs a procedure when the patient expected it to be performed by another. These could be seen as true instances of mistaken identity. Rather, liability has arisen primarily where the claimant believed the defendant to have certain attributes, qualifications or a certain status, when in fact they did not. These are instances of confused identity. The following will consider how both of these categories are to be applied to AI.

Regarding true cases of mistaken identity, *Michael v Molesworth* is an illustrative case from the medical sphere. Here a claimant succeeded in suing a house surgeon for battery because an apprentice had carried out the relevant surgery to practise his skill, even though this was done entirely competently.⁷⁸¹ However, it is unlikely that this case would serve as precedent in the context of the modern National Health Service (NHS). As McHale has noted, the NHS standard contract reflects the fact that there is generally no expectation that one particular professional will be responsible for one's care in an NHS hospital and it would likely require exceptional circumstances for the patient to elevate such a factor to a condition for their

⁷⁸⁰ In one of the cases within this class, the court's finding was framed solely in terms of the 'quality' of the act, explicitly distinguishing the 'identity' of the actor: *R v Tabassum* [2000] Lloyd's Rep Med 404. However, this framing does not appear to be based on a substantive distinction. Rather, it can be explained by a need to distinguish the decision in *R v Richardson* [1999] QB 444. As we will see, the courts now appear to accept that arguments can proceed straightforwardly under the head of identity, even where purely the status, attributes or qualifications of the 'professional' are concerned: *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063.

⁷⁸¹ *Michael v Molesworth* (1950) 2 BMJ 171.

consent and legitimate treatment.⁷⁸² Similarly, it appears to be implicit in the more recent case of *R v Richardson* that true mistakes as to identity may only occur where an individual is quite literally in error about who it is that is interacting with them.⁷⁸³

Following this line of argumentation, there will generally be little scope for mistakes of identity to occur in relation to AI. It has not been claimed that ML devices possess legal personhood and it has been argued that they will not fully replace human involvement in any given type of medical treatment. It is true that, similar to what occurred in *Michael*, an AI could take over one aspect of a procedure that involves multiple human professionals. In Chapter 2 this was subsumed under the category of ‘devices that partially replace pre-existing cognitive capabilities’. Yet, as this kind of substitution falls far short of the replacement or impersonation of a known individual, a claim in battery is not envisaged.

It is more promising to argue that AI involvement may lead to confusions of identity and that they would vitiate consent on this basis. Several cases go directly towards this issue. In *R v Tabassum* the defendant purported to show a number of women how to examine their own breasts for the purposes of detecting breast cancer.⁷⁸⁴ The women believed that he had the requisite medical training to do this. However, while he did have minimal experience in this field – having prepared several leaflets for self-examination and having learned about several diseases during his work as a medical representative – he had neither relevant training nor qualifications to give such instructions.⁷⁸⁵ The centrality of the defendant’s identity to the finding

782 The patient would have to reach an agreement with the relevant organisation and professional that the consent is conditional on the treatment being undertaken by a particular consultant: McHale in Laing and McHale, *Principles of Medical Law* (2017) 437. See also Brazier and Cave, *Medicine, Patients and the Law* (Sixth Edition 2016) 133. A different view is expressed by Donnelly, who draws on ‘Lord Hope’s obiter statement in *Chester* that a patient has a right to be informed “as to whether, and if so when and by whom, to be operated on”’: Donnelly, *Healthcare Decision-Making and the Law: Autonomy, Capacity and the Limits of Liberalism* (2010). Yet the influence that such an *obiter* statement, which was made in a case dealing solely with negligence, can have on the law of battery is to be doubted.

783 Otton LJ cited a definition of identity as ‘the condition of being the same’ and considered that cases of mistaken identity must be of the same nature as mistakes that could lead a woman to have sexual relations with a man that she believes to be her husband: *R v Richardson* [1999] QB 444, 459-450.

784 *R v Tabassum* [2000] Lloyd’s Rep Med 404 [4].

785 *ibid* [18].

that the women's consent was indeed vitiated in these circumstances is palpable, even if the court relied on a notion of 'qualitative deception' to decide the issue.⁷⁸⁶ According to one commentator, the patients' consent was vitiated because 'the Court of Appeal considered that the women's consent was to "breast examination for the purposes of preparing a medical software package by a medically qualified person" rather than simply "breast examination for the purposes of preparing a medical software package"'.⁷⁸⁷

This convoluted approach appears to have resulted from a desire to distinguish the decision from the aforementioned *R v Richardson*. In *Richardson* a claim was rejected that a dentist's status or attributes – specifically her disqualification from practicing – went to her identity, vitiating the consent of her patients.⁷⁸⁸ In spite of the different terminology, these decisions are not easy to reconcile. Aspects of a defendant's identity are relevant to the courts' consent analysis, but the relevant arguments have partially been couched in terms of the quality of the act.

Fortunately, the recent case of *R v Melin* has clarified the legal position. The Court of Appeal upheld a finding that the consent of the claimant had been vitiated by aspects of the defendant's identity. To reach this result, the Court had no need to rely on the reasoning of *Tabassum*, but rather proceeded by giving a narrow interpretation to *Richardson*'s restrictive comments on identity.⁷⁸⁹

Specifically, the issue in this case was whether the consent of several women to a series of Botox injections was vitiated by the fact that they believed the defendant to possess medical qualifications, when in fact he did not. It was decisive for the affirmation of the defendant's liability that his identity was inextricably bound up with his claimed status as a doctor. And the two were inextricably bound up with one another because that status was operative in the complainant's mind when they gave their consent.⁷⁹⁰ The significance of certain attributes for the purposes of giving a valid consent was thereby expressly acknowledged. Citing *Smith, Hogan and Ormerod*'s

786 *ibid* [37]-[38]. See also *R v Dica*, which assessed the 'quality of the act' to determine whether there had been consent to sexual intercourse and to the risk of the transmission of HIV where the defendant had not disclosed his positive status: *R v Dica* [2004] EWCA Crim 1103, [2004] QB 1257 [38]-[39].

787 Maclean, *Autonomy, Informed Consent and Medical Law: A Relational Challenge* (2009) 162 (my emphasis).

788 *R v Richardson* [1999] QB 444.

789 *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063 [34]-[35].

790 *ibid* [33], [40].

Criminal Law it was stated that ‘it could be that the attribute is actually more important than the identity. For example, would a patient visiting a general practitioner and being told that a new doctor is taking the surgery be more concerned as to the “status” of the person or his “identity”’.⁷⁹¹

It was also held that the legal requirements for the level of qualifications necessary to undertake a particular procedure, were not determinative of the factual inquiry related to identity.⁷⁹² Therefore, for the purposes of the consent enquiry in *Melin*, it did not matter that Botox injections were not legally required to be undertaken by medically qualified practitioners. What mattered was that, as a matter of fact, the patient’s consent was predicated on the medical qualifications of the individual treating her and that she was mistaken to this extent.

The issues addressed in these cases will become relevant to those situations where AI partially replace human cognitive capabilities. Especially where such devices reduce the level of human expertise that is brought to bear on a given decision. The example of IDx-DR was provided in Chapter 2 in this respect. The key questions that are posed in light of the outlined legal context are this: is an individual taking advantage of such systems to carry out clinical tasks obligated to inform their patients of the fact that they are extending their expertise by relying on the aid of an AI? Are they misleading the patient about possessing a relevant status, qualification or attribute if they do not?

Several grounds can be advanced to support of an affirmative answer. In the first instance, as *Tabassum* and *Melin* illustrate, the treating professional’s expertise and qualifications will often be a factor that is operative in patients’ minds. By contrast, the possibility that a healthcare worker would only have the expertise to deal with their condition as a result of technical assistance is currently unlikely to occur to most patients.⁷⁹³ It is therefore arguable that in circumstances where patients would expect a procedure to be carried out only by a professional with a certain qualification and/or level of expertise, they are only consenting on the basis of this assumption. If they are not informed that the professional does not themselves possess

791 *ibid* [30], citing: Ormerod and Laird, *Smith, Hogan, & Ormerod’s Criminal Law* (Fifteenth Edition 2018) 672.

792 *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063 [33].

793 Distinguish the ordinary situation where the human has the requisite capabilities but merely relies on technical assistance (e.g. an X-Ray machine) as a tool that provides the information upon which these capabilities are exercised.

these attributes, there is a *prima facie* case that this may vitiate their consent.

Melin has further made clear that this analysis is not prejudiced by the wider legal framework. So that, even if an AI is an approved medical device and if professionals are legally permitted to undertake more specialised tasks with its help, this will not mean that the patient's consent to such a procedure will automatically be informed and valid. This is a question of fact.

Another supportive aspect of the existing law is the repeated insistence that what ultimately matters is not the source of the mistake, but the fact of its occurrence.⁷⁹⁴ The ancillary and novel nature of the AI technology make it likely that the patient would be mistaken about its involvement as a result of an omission to inform them, not because of positive misrepresentation or intentional fraudulent conduct. The cases are now clear that the presence or absence of such conduct on behalf of the defendant is not determinative of whether the patient's consent has been vitiated. The decisive factor remains whether there was a mistake of the requisite kind, which was operative in the patient's mind. In consequence, so long as the initial analogy to cases of confused identity is deemed convincing, the wider legal requirements (allowing AI use) or the absence of fraud will not defeat a claim in battery.

What, then, are the factors that make the drawing of this analogy contentious? First, there are grounds for distinguishing medical practice with AI assistance from the claims that have successfully asserted confused identity. Most notable is the fact that, in the established judgments, the defendants evinced a fundamental lack of relevant medical expertise and qualifications. As MacLean has commented, in a case like *Tabassum* the only relevant factor appears to have been 'that the accused was not medically qualified, which emphasises the courts' reluctance to find doctors liable for battery'.⁷⁹⁵ In other words, given the nature of the tort and the case law on confused identity, the fact that a user of AI is likely to have *some* medical

794 'The common law is not concerned with the question whether the mistaken consent has been induced by fraud on the part of the accused or has been self induced': *R v Richardson* [1999] QB 444, 450; 'it would be undesirable for the law to treat all false or fraudulent representations as vitiating consent': *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063 [29].

795 Maclean, *Autonomy, Informed Consent and Medical Law: A Relational Challenge* (2009) 162.

experience and qualification is a distinction that weighs against a successful claim in battery

The exact force of this consideration is difficult to pin down. On the one hand, it appears hard to argue that the total lack of qualification is part of the *ratio* of these cases. For, in *Tabassum* the defendant possessed some rudimentary forms of healthcare experience and in *Melin* some consideration was given to the differing levels of expertise that the defendant claimed to possess.⁷⁹⁶ On the other hand, it is conspicuous that in *Richardson* – the one case dealing with a qualified individual, albeit practising without a licence – a claim in battery was rejected. It seems that the broadly hostile position that the courts have adopted in relation to the application of battery to professionals cannot easily be discounted.

Second, there is the fact that the AI in question will be designed with the purpose of replacing and thereby raising the expertise of the utilising professional to the requisite specialist level. As our analysis of medical AI demonstrates, the technology is intended to compensate for a lack of human experience or qualification. By contrast, the outlined judgments only deal with situations where the individuals, in spite of holding themselves out as having a certain level of expertise, did not themselves possess it and did not and could not acquire it through assistance.

This raises the question of whether the courts will consider that the *relative* proficiency of the professional remains the crucial aspect, or whether the patient should only be informed if there are variations in the *aggregate* levels of expertise that are brought to bear on their treatment. This is an open question under the existing jurisprudence. Consequently, it is at this juncture that one must consider the differentiated approach that the British courts have taken to responding to interferences with patient autonomy.

Our analysis from Chapter 3 suggests that the different nature of AI-generated knowledge poses a problem for a patient's procedural autonomy. In a sense, even when an AI achieves functioning comparable to that of a human expert, it can never represent the same kind of expertise – it never provides the same aggregate level of skill. This deficiency will be more pronounced in relation to some AI uses than others.

Nevertheless, unless there is a substantial unenvisioned issue with the AI's operation (effectively subjecting the patient to unskilled treatment in much

⁷⁹⁶ The defendant had variously been introduced as a doctor in the Turkish army who had specialised in facial surgery, a cosmetic surgeon and a nurse: *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063 [14], [16].

the same way as the individuals in *Tabassum* and *Melin*), the resulting autonomy violation would be on a much smaller scale. The patient's ability to reason with the requisite information (their theoretical rationality) may be impaired, but, without more, they are not prevented from achieving their goals through the human-AI hybrid form of assistance.

Ultimately, in cases where there is no aggregate loss of expertise, the weight of the autonomy principle is more limited and, in spite of it receiving increasing attention, it is unlikely to outweigh the entrenched factors that have deterred the courts from imposing liability on professionals in cases of mistaken and confused identity. It would take a considerable development of the case law to establish liability for qualified medical staff. Attention to the autonomy principle is arguably insufficient to bring this about.

3. Non-therapeutic motivations

The final category of cases to be addressed are those that deal with the non-therapeutic motivations of professionals (or purported professionals). In particular, one can see that in cases where the primary motivator for offering the patient an intervention has been sexual or financial, rather than clinical, the courts have had no trouble entertaining claims that the patient's consent was vitiated and a battery perpetrated – even by qualified medical professionals.⁷⁹⁷

From the perspective of legal policy this is understandable and commendable. A concern that an injustice is being done to the defendant is evidently misplaced in such circumstances, nor is it plausible to maintain that a finding of battery would have far-reaching, negative implications on medical practice. However, since these cases are only tangentially related to AI use, they will be dealt with only briefly.

It is true that, in several respects, non-clinical motivations will be relevant to AI decision-making. This was highlighted in Chapter 3, in our discussion of AI's goal-directed action. Irrespective of how concerning these hidden and ancillary motives may be, a claim in battery is unlikely to

797 *R v Williams* [1923] 1 KB 340 and *Appleton v Garrett* (1997) 34 BMLR 23 respectively. The defendant in the former was a singing instructor who had sexual intercourse with his student under the pretence of an operation that would improve her breathing, while the latter concerned a qualified dentist who carried out unnecessary procedures for financial reasons.

succeed because of the aforementioned targeting of this mechanism. In the existing case law, the motive in question is one that is highly personal to the relevant professional, shaping the nature of their actions and providing the basis for their liability. By contrast, where an AI is involved, the non-clinical motive will often be unknown to the AI user, especially if it is sufficiently problematic. Even if such a purpose is known, it will, as was highlighted in Chapter 2, be difficult to assert that divergent AI goals then actually drove the medical decision or supplanted the therapeutic objective of the human professional.

In sum, this class of case seems the least likely to establish liability in battery for non-disclosure because the actor who is targeted by battery (the professional) and the actor who is covertly introducing the non-medical motive into the individual medical decision (the AI and/or the AI developer) come apart. The analogy with this category of case law must break down and it cannot constitute a basis for the duty to disclose AI involvement.

C. Summation

Following the above analysis, there would be one class of case where the tort of battery requires the disclosure of AI-specific information in medical treatment. A duty to disclose arises where the AI alters the nature of the relevant procedure, by introducing a potential, which is not easily foreseen, to diagnose serious conditions in the patient or to bring about other significant effects without their knowledge. If this use of AI and its implications are not disclosed to a patient, then their consent will arguably become invalid and a claim in battery will lie – subject to the fulfilment of the other outlined requirements. This finding is supported by a differentiated application of the autonomy principle.

Regarding the other classes of information, the claimant will have to turn to the negligence action. However, it must be borne in mind that negligence presents a much more restricted means of vindicating patient autonomy. Unlike the tort of battery, negligence requires that damage eventuated and was caused by the defendant's act. Whereas a touching is actionable *per se* and damages can be claimed for all direct consequences in battery, bringing a claim in negligence ordinarily calls for a patient to have suffered physical

harm and it imposes further limitations in terms of causation and the remoteness of damages.⁷⁹⁸

II. Negligence

To construct an obligation to advise patients of AI characteristics under negligence, one must begin with an analysis of the tort's broader elements. As with battery, these will shape its operation in relation to information provision.

Negligence is concerned with the obligation of individuals and organisations to exercise due care and skill in their interactions with others. Specifically, the elements of the action are: 'a duty of care, a breach of that duty and consequent damage'.⁷⁹⁹ Giving due weight to the fact that 'consequent damage' concerns both an element of damage and an element of causation, our analysis can be split into four categories. First, an actionable form of damage must be suffered by the claimant. Second, the defendant must owe a duty of care of the requisite scope to the claimant. Third, this duty of care must be breached. Fourth, this breach must have caused the damage. It will be seen that each of these elements has been impacted by the principle of patient autonomy in the disclosure context. In addition, it will be seen that it has influenced the monetary compensation, the 'damages' recoverable under the tort.

It should therefore come as no surprise that, as has been discussed above, negligence constitutes the focus of judicial and academic pronouncements on the tortious obligation of a medical professional to obtain the patient's *informed* consent. The common law of England has developed negligence to impose broad informational duties on individuals and organisations, duties that go far beyond the requirements of obtaining a patient's valid consent to treatment.

Indeed, the common law has been prepared to adapt and stretch its doctrinal integrity to accommodate the principle of patient autonomy. Some argue that this has been to the breaking point, questioning the extent to

798 For a more general discussions of the advantages and benefits of the two claims and their relation to one another see: Feng, 'Failure of Medical Advice: Trespass or Negligence?' (1987) 7(2) Legal Studies p. 149; Brazier, 'Patient Autonomy and Consent to Treatment: The Role of the Law,' (1987) 7(2) Legal Studies p. 169.

799 *Burton v Islington Health Authority* [1993] QB 204, 224.

which informed consent actions can still be classed as negligence actions.⁸⁰⁰ The following assesses whether this flexibility would provide an adequate form of protection in the face of the new challenges posed by ML technologies.

It should be highlighted that this action lies in tort. Claims that will not be examined further in relation to English law are those of ‘contractual negligence’ and claims based on contract more generally. This stems partly from the fact that, in the great majority of cases, AI will be applied to the treatment of patients in the context of the NHS, where there is no contract between patient and provider.⁸⁰¹ It is also informed by the view that the courts will be reluctant to imply that contractual agreements raise the negligence-based standard of care.⁸⁰² So it is assumed that a separate consideration of contracts would contribute little to our analysis.

A. Actionable damage

Although often neglected in practical and theoretical argumentation, the requirement that actionable damage must be suffered by the claimant is a precondition for liability in negligence.⁸⁰³ Unlike battery, negligence is not a tort that is actionable *per se*. As it has been said: damage is the gist of the action.⁸⁰⁴ This requires the eventuation of legally specified kinds of harm. Under UK law personal injury, property damage, economic loss

800 Purshouse, ‘The Impatient Patient and the Unreceptive Receptionist: Darnley v Croydon Health Services NHS Trust [2018] UKSC 50’ (2019) 27(2) Medical Law Review p. 318, 329; Nolan, ‘Negligence and Autonomy’ [2022](2) p. 356, 381.

801 ‘[T]he arrangement between doctor and patient under the aegis of the NHS is statutory rather than contractual’: *Reynolds v Health First Medical Group* [2000] Lloyd’s Rep Med 240, 242. See also *Pfizer Corporation v Ministry of Health* [1965] AC 512, 535-536, 544-545, 548, 552-553, 571. Here the House of Lords outlined how the existence of a statutory obligation, in this case for the provision of prescription drugs, negates the existence of a contractual relationship in the NHS context.

802 Pattinson, *Medical Law and Ethics* (Sixth Edition 2020) 52-53, citing *ARB v IVF Hammersmith* [2018] EWCA Civ 2803, [2020] QB 93 and *Thake v Maurice* [1986] QB 644.

803 ‘[D]amage is an essential element in a right of action for negligence’: *Dulieu v White & Sons* [1901] 2 KB 669, 673. See also Nolan, ‘New Forms of Damage in Negligence’ (2007) 70(1) The Modern Law Review p. 59, 59-61.

804 *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 883-884.

and psychiatric harm currently constitute the well-established categories.⁸⁰⁵ By contrast, the courts have been consistent in rejecting distress, inconvenience or discomfort as actionable types.⁸⁰⁶

Without identifying the actionable damage that provides the basis for the claim first, one risks confusing the analysis of later elements.⁸⁰⁷ This is particularly true for the current analysis of the law's response to clinical AI's challenges. A traditional category of harm, specifically physical injury, may sometimes accompany problematic uses of ML devices in this context. But it is ultimately an ancillary possibility. In comparison, as will be elaborated, individual autonomy is not a traditionally recognised category of legally cognisable injury. Yet, the coherent protection of the outlined principle will depend upon the judiciary's preparedness to compensate its violation. This will have repercussions throughout the assessment: whether autonomy is compensable as such can affect analyses of duty, breach, causation and damages.

1. Personal injury

It is beyond contention that personal injury is a legally recognised form of damage under UK negligence law. Often this will mean that the law is in a position to address grievous autonomy violations and to offer a redress that is consistent with this value.⁸⁰⁸

In the case of the challenges posed by AI this can be concretised. In a relatively ordinary medical non-disclosure case, it may happen that the use of a specific AI tool poses a risk to the patient that is not properly disclosed

805 Albeit the latter two are subject to additional conditions and thus actionable only in some circumstances, see: Nolan, 'New Forms of Damage in Negligence' (2007) 70(1) *The Modern Law Review* p. 59, 60-61.

806 'If his negligence has caused me neither injury to property nor physical mischief, but only an unpleasant emotion of more or less transient duration, an essential constituent of a right of action for negligence is lacking': *Dulieu v White & Sons* [1901] 2 KB 669, 673.

807 Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 264-265; Purshouse, 'Judicial Reasoning and the Concept of Damage: Rethinking Medical Negligence Cases' (2015) 15(2-3) *Medical Law International* p. 155, 163-164.

808 '[N]egligence protects autonomy as a second-order value because the kinds of injuries that ground negligence claims almost inevitably have a negative impact on the plaintiff's ability to live the life she would choose to live': Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 357.

(as will be specified in our analysis of the breach element) and that this leads to a situation where the patient suffers precisely that harm, which they were not warned against. The example cited in Chapter 3 was the *Acumen Hypotension Prediction Index Software* giving a false positive reading, which results in an unnecessary intervention. This will straightforwardly interfere with the patient's ability to direct their lives according to their values and, potentially permanently, prevent them from living their life as they had planned.

Alternatively, it may be the case that an ML-related autonomy violation would have weighed heavily enough in a patient's reasoning to alter their conduct. The hypothetically altered decision may have avoided a harm that then eventuated incidentally, potentially entirely unrelated to the AI's functioning. Here too a personal injury claim will protect the patient's autonomy to some extent.

Ultimately, the patient's autonomy can be protected through negligence in such situations because there is a physical harm which flows from the defendant's non-disclosure. This will have repercussions both for the breach element, as certain forms of non-disclosure are more easily associated with the eventuation of physical harm, and for the causation element of negligence, as only certain forms of non-disclosure will be significant enough to bring about a different practical decision.

2. Loss of autonomy

Unlike physical injury, it is not a straightforward task to identify whether the tort of negligence views autonomy as a separate head of damage. As outlined, it is clear that the law recognises multiple categories of damage. However, there is no generally determinative ground for identifying these.⁸⁰⁹ Following Nolan, we may say that the technical legal nature of the term 'damage' gives rise to some circularity: the defendant's action must simply violate one of 'various rights (to bodily integrity, property and so on) which are protected against negligent interference'.⁸¹⁰

809 Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 265-267.

810 *ibid* 268. This can be related to the position that 'the concept of damage is at the end of the day not a factual, but a normative concept': *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [45].

Established legal doctrine must therefore play a leading role in identifying whether a diminishment of autonomy is a harm that is recognised at law. It would be difficult to make this argument without the support of prior legal documentation. This ensures the condition's role as a controlling and stabilising factor for the negligence action.⁸¹¹

In evaluating the relevant documentation, one can begin with the position that there is no well-established common law tradition according to which loss of autonomy is perceived as actionable damage in negligence.⁸¹² Yet this tradition is itself mutable. Purshouse provides the historical examples of harms that have been removed from the accepted catalogue, including: the seduction of a daughter, or the enticement of a wife.⁸¹³ Conversely, the UK courts have been prepared to develop a class of recognised psychiatric harm.⁸¹⁴ Since the categories of loss remain fluid, and especially given that the interest in autonomy has gained relatively recent prominence in the bioethical, societal and legal discourse, the absence of a fortified common law position recognising autonomy damage should not be treated as prohibitive for the recognition of a relevant injury. A more recent case of the United Kingdom's highest court, has arguably brought about just such a development, acknowledging individual autonomy as an independent injury and permitting a claimant to bring an action on this basis.

This is the judgment of the House of Lords in *Rees v Darlington Memorial Hospital NHS Trust*.⁸¹⁵ The House of Lords began by rejecting one form

811 Purshouse, 'Autonomy, Affinity, and the Assessment of Damages: *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 and *Shaw v Kovak* [2017] EWCA Civ 1028' (2018) 26(4) Medical Law Review p. 675, 687. Note how such an argument fits with our general observations of private law reasoning in the common law in Chapter 4.

812 Purshouse, 'How Should Autonomy Be Defined in Medical Negligence Cases?' (2015) 10(4) Clinical Ethics p. 107, 107-108; 'the right to make an informed choice is not a right that is traditionally protected by the tort of negligence': *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [88].

813 Purshouse, 'Judicial Reasoning and the Concept of Damage' (2015) 15(2-3) Medical Law International p. 155, 158.

814 Compare *Dulieu v White & Sons* [1901] 2 KB 669 and *Victorian Railway Commissioners v Coultas* (1888) 13 App Cas 222. See generally: Law Commission, 'Liability for Psychiatric Illness' (Law Commission Consultation Paper No 137, 1995) 4, 8. Likewise, Priaulx argues that 'the action for wrongful conception (...) clearly demonstrates the law of tort's ability to embrace a widening ambit of harms under its cloak': Priaulx, 'Joy to the World! A (Healthy) Child Is Born! Reconceptualizing Harm in Wrongful Conception' (2004) 13(1) Social & Legal Studies p. 5, 6.

815 *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309.

of actionable damage. Namely, maintaining the position that had been laid down in *McFarlane v Tayside Health Board*⁸¹⁶ that a claimant who had been negligently sterilised was not, when they became a parent, entitled to the costs of raising their child. These costs were not an actionable form of damage.⁸¹⁷ At the same time, the parent was considered the victim of a legal wrong and it was maintained that this wrong would not be adequately remedied by an award that covered only the immediate personal and economic consequences of pregnancy and birth.⁸¹⁸ To respond to this wrong, the majority in *Rees* therefore departed from the earlier *McFarlane* decision – adding ‘a gloss’ to it, as it was described in the case⁸¹⁹ – and awarded a conventional sum of £15,000 to the parent.⁸²⁰ Understanding the nature of this sum is crucial for the purposes of establishing whether loss of autonomy constitutes a standalone head of damage.

In order for the award to support the proposition that autonomy is a new actionable head of loss, which can coherently support the remainder of the negligence action, two things must be true: (1) the award must be compensatory (2) the compensatory award must be for a suitable interest in autonomy.⁸²¹ The legal position with regard to both will be examined in turn, considering the interpretations and developments that can be found in the subsequent case law.

i. The nature of the award

First, the sum must constitute compensation for a damage (often described as a loss).⁸²² This has proved somewhat controversial, as it has been argued

816 *McFarlane v Tayside Health Board* [2000] 2 AC 59.

817 I follow Purshouse in this view: Purshouse, ‘Judicial Reasoning and the Concept of Damage’ (2015) 15(2-3) *Medical Law International* p. 155, 160-166.

818 *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [8].

819 *ibid* [7], [17].

820 *ibid* [8], [10], [17], [19]. This picked up upon Lord Millet’s suggestion in *McFarlane* (which was not adopted by the majority) to award a lesser sum of £5,000: *McFarlane v Tayside Health Board* [2000] 2 AC 59, 114.

821 For a similar distinction in the analysis of *Rees*, see the decision of the Singaporean Court of Appeal in: *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [111].

822 As was pointed out, Nolan outlines the technical nature of the term, in particular that it cannot be treated as equivalent with terms such as harm, injury or loss: Nolan, ‘Damage in the English Law of Negligence’ (2013) 4(3) *Journal of European*

that the award in *Rees* is best conceived of as a vindication of the right to autonomy.⁸²³

Vindication has a fundamentally expressive – rather than compensatory – function, seeking to affirm certain interests or rights of the claimant.⁸²⁴ The aim is ‘to make it clear to the world, or more precisely to the two parties, that the wrong was a wrong and should never have happened’.⁸²⁵ If this is the case, then the award and its relation to negligence become highly suspect. The vindication of a right does not presuppose any damage and this calls into question a limitation that is central to this broad tort, with its focus on the provision of compensation for careless behaviour.⁸²⁶ Any argument that the decision in *Rees* ought to be affirmed or even extended would correspondingly be weakened.

Those who would claim that the conventional award in *Rees* was vindicatory can point to several supporting factors. Quite explicitly, some of their Lordships framed the award by employing the language of rights.⁸²⁷ However, in our discussion of the autonomy concept we have already

Tort Law p. 259, 266. I will continue to use these terms interchangeably, in line with common usage, without intending to imply that a legal damage must always entail a factual loss or injury.

823 Witzleb and Carroll, ‘The Role of Vindication in Torts Damages’ (2009) 17(1) Tort Law Review p. 16, 38. The same is implicit in Foster’s claim that: ‘Previously the claimant came into court and said: “The defendant owes me a duty. He has breached it. I have suffered loss. I want compensation.” Now the claimant can say instead: “I have a right. It has been violated by you. You should not have violated it. I may or may not have suffered real loss, but I want compensation”’: Foster, *Choosing Life, Choosing Death: The Tyranny of Autonomy in Medical Ethics and Law* (2009) 127-128.

824 Varuhas, ‘The Concept of ‘Vindication’ in the Law of Torts: Rights, Interests and Damages’ (2014) 34(2) Oxford Journal of Legal Studies p. 253, 258-260.

825 Smith, ‘Duties, Liabilities, and Damages’ (2012) 125(7) Harvard Law Review p. 1727, 1753-1754. Mulligan adopts this approach in the UK context: Mulligan, ‘A Vindicatory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction’ (2020) 40(1) Legal Studies p. 55, 63.

826 Nolan frames this in terms of maintaining potential defendants’ freedom of action: Nolan, ‘New Forms of Damage in Negligence’ (2007) 70(1) The Modern Law Review p. 59, 79. Varuhas also makes the restricted nature of negligence clear: ‘[Negligence’s non-vindicatory nature] explains the focus or emphasis within negligence on actionable damage, the nature of the defendant’s conduct, and compensation for factual harm, as well as the generalized nature of the concepts governing liability’: Varuhas, ‘The Concept of ‘Vindication’ in the Law of Torts’ (2014) 34(2) Oxford Journal of Legal Studies p. 253, 260.

827 For example Lord Millett argued that the ‘right to limit the size of their family’ is ‘increasingly being regarded as an important human right which should be protected

highlighted the imprecision often involved in such statements and that, accordingly, they can hardly be taken to be determinative. More convincingly one can refer to Lord Bingham's statement that 'The conventional award would not be, and would not be intended to be, compensatory. It would not be the product of calculation (...) It would afford some measure of recognition of the wrong done'.⁸²⁸ It is also noticeable that the dissenting Lord Hope despaired at the fact that the award deviated from common law principle and that the majority's position lacked 'any consistent or coherent ratio'.⁸²⁹ In combination with the rights language employed, these statements support the view that the orthodox objective of negligence law, compensation, was not the purpose of *Rees*' award.

Another argument to this effect turns on the relationship between the two outlined cases of *McFarlane* and *Rees*. *Rees* expressly affirmed the earlier decision, and this is said to imply that the claimant did not suffer any other compensable loss – that is, loss going beyond that immediately associated with pregnancy.⁸³⁰ By contrast, if *Rees* found that, exceptionally, one should vindicate the individual's right to autonomy through damages – this issue would lie outside of *McFarlane*'s scope. In other words, a rights-based analysis accords due respect to precedent and presents a more coherent approach.

Beyond an immediate analysis of this case law, commentators have further relied on doctrinal reasons that are related to the fixed nature of *Rees*' award. A compensatory award would usually aim to put the claimant back into the position that they would have been in had the wrong not been perpetrated.⁸³¹ As such, it is rather anomalous to hold that there would be a fixed sum for very different kinds of experiences (some of which Lord Bingham outlined).⁸³² It is hard to maintain that the same award is just

by law': *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [123].

828 *ibid* [8].

829 *ibid* [74].

830 Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) *Legal Studies* p. 55, 64-65.

831 *Livingstone v Rawyards Coal Co* (1880) 5 App Cas 25.

832 'The spectre of well-to-do parents plundering the National Health Service should not blind one to other realities: that of the single mother with young children, struggling to make ends meet and counting the days until her children are of an age to enable her to work more hours and so enable the family to live a less straitened existence; the mother whose burning ambition is to put domestic chores so far as possible behind her and embark on a new career or resume an old one. Examples

or appropriate, even if one considers only the relatively narrow band of wrongful conception cases.⁸³³ In comparison, a vindictory award can readily explain the fixed nature of the conventional award: 'it compensates for the interference with one's autonomy interest, quantum being held constant on the basis that normative damage is assessed objectively'.⁸³⁴

Ultimately, given the undeniable difficulty of reconciling vindictory damages with negligence, the most convincing proponents of the outlined reading of *Rees* argue that it responded to a vindictory impulse.⁸³⁵ This leaves an autonomy-based award as something anomalous and exceptional. As proponents appear to admit, this would make it hard to apply to novel situations, even where a strong analogy can be drawn with *Rees*.⁸³⁶ For our present circumstances it would mean that AI's autonomy infringements would only become actionable *per se* if they led the claimant to have a child that they did not wish to.

This conclusion must be rejected however. Apart from the countervailing statements adduced above, and some admitted variation in their Lordships' statements, the overwhelming tenor of the majority's judgment in *Rees* was compensatory. Lord Bingham spoke directly in terms of the 'real loss suffered' and went on: 'a parent, particularly (even today) the mother, has

can be multiplied. To speak of losing the freedom to limit the size of one's family is to mask the real loss suffered in a situation of this kind. This is that a parent, particularly (even today) the mother, has been denied, through the negligence of another, the opportunity to live her life in the way that she wished and planned': *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [8].

833 Witzleb and Carroll, 'The Role of Vindication in Torts Damages' (2009) 17(1) Tort Law Review p. 16, 38. Mulligan provides a summary of the relevant criticisms: Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) Legal Studies p. 55, 64.

834 Varuhas, 'The Concept of 'Vindication' in the Law of Torts' (2014) 34(2) Oxford Journal of Legal Studies p. 253, 269.

835 *ibid* 269; adopted by Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) Legal Studies p. 55, 65.

836 See especially Mulligan's attempt to extend *Rees* to the relatively similar circumstances of mistakes in reproductive treatment: Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) Legal Studies p. 55, 69, 71. Varuhas and Mulligan both appear to agree that direct and systematic protection of autonomy in such cases would require a standalone action: Varuhas, 'The Concept of 'Vindication' in the Law of Torts' (2014) 34(2) Oxford Journal of Legal Studies p. 253, 270; Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) Legal Studies p. 55, 72-76.

been denied, through the negligence of another, the opportunity to live her life in the way that she wished and planned'.⁸³⁷ Lord Scott similarly supported 'a conventional sum to compensate', placing 'a monetary value on the expected benefit of which [the claimant] was, by the doctor's negligence, deprived'.⁸³⁸ Lord Millet framed the matter in terms of 'A modest award [to] compensate for the very different injury to the parents' autonomy'.⁸³⁹ And it is this characterisation that has primarily been echoed in the literature.⁸⁴⁰

The adduced doctrinal criticisms also seem manageable: a fixed award can be given for a compensated damage in order to render it objective.⁸⁴¹ Doing so may appear anomalous and unjust, but it is not fundamentally subversive or unheard of. In this respect, Purshouse notes how a £200 lump sum was awarded for 'loss of expectation of life' in *Benham v Gambling*,⁸⁴² as the calculation in individual cases proved troublesome.⁸⁴³ He goes on to note that, on a similar account, such an award may be appropriate for autonomy violations where calculating individual losses may cause chaos.⁸⁴⁴

Another commentator has also considered the outcome in *Rees* to be explainable 'on orthodox terms as a fixed amount for non-pecuniary loss'.⁸⁴⁵ In support they adduced a further, later case considering this type of approach: *Shaw v Kovac*.⁸⁴⁶ Here, the Court of Appeal understood the conventional award to 'mark the injury and the loss', being 'designed to

837 *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [8].

838 *ibid* [148].

839 *ibid* [125].

840 Nolan, 'New Forms of Damage in Negligence' (2007) 70(1) *The Modern Law Review* p. 59, 79-80; Purshouse, 'Judicial Reasoning and the Concept of Damage' (2015) 15(2-3) *Medical Law International* p. 155, 166-169; McGregor and others, *McGregor on Damages* (Twenty-First Edition 2021) para 40-300.

841 Varuhas, 'The Concept of 'Vindication' in the Law of Torts' (2014) 34(2) *Oxford Journal of Legal Studies* p. 253, 269.

842 *Benham v Gambling* [1941] AC 157.

843 Purshouse, 'Judicial Reasoning and the Concept of Damage' (2015) 15(2-3) *Medical Law International* p. 155, 167. Cf. Todd, who emphasises the point that this is truly unusual and that only the cited case of *Benham* had previously made such an award: Todd, 'Common Law Protection for Injury to a Person's Reproductive Autonomy' (2019) 135 *Law Quarterly Review* p. 635, 652.

844 Purshouse, 'Judicial Reasoning and the Concept of Damage' (2015) 15(2-3) *Medical Law International* p. 155, 167.

845 McGregor and others, *McGregor on Damages* (Twenty-First Edition 2021) para 40-300.

846 *ibid* para 40-300.

compensate the claimant for the loss of the opportunity to live her life in the way she had wished and planned'.⁸⁴⁷ In short, the unorthodox, anomalous nature of the conventional award must not be overstated.

The need for consistency with *McFarlane*, and the degree to which a vindictory approach ensures such consistency, should also not be exaggerated. Clearly the House of Lords in *Rees* saw itself as adding to that judgment and it is no less consistent to recognise a distinct form of loss, which remained underappreciated in *McFarlane*, than to recognise a novel need for the vindication of a right.⁸⁴⁸ In both cases the core tenet of *McFarlane* – the denial of the costs for raising the child – stands. This is precisely how the Singaporean Court of Appeal in *ACB v Thomson Medical Pte Ltd* conceived of the award. The claim for compensating autonomy does not fix 'on the liabilities arising out of the care of the *unplanned child* (which is the gravamen of the objection against the award of upkeep) but on the *independent interests of the parents which have been transgressed as a result of the negligent act*'.⁸⁴⁹ Clearly, the much larger problem for this evaluative dimension is not the inconsistency that would arise with *McFarlane* if the award were treated as compensatory, but the incoherence that would follow from making a vindictory award in a tort that is fundamentally structured to be compensatory.⁸⁵⁰

Purshouse also points to another fundamental issue for the rights-based approach: it is difficult to conceive of a right that would need vindication, without first determining that negligence has provided a cause of action.⁸⁵¹ This objection can be concretised by reference to the aforementioned case of *Shaw v Kovac*, where the Court of Appeal held that the common law would require a clear and fundamental right before awarding vindictory damages.⁸⁵² Specifically, drawing on the Supreme Court's judgment in *R*

847 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [80].

848 Consider especially in this regard Lord Steyn's statement that the majority in *McFarlane* had considered and rejected a conventional award – even if it was not explicitly discussed by them: *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [41].

849 *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [108].

850 Nolan persuasively outlined this dimension, arguing that a vindictory award would be a fundamental challenge to negligence: Nolan, 'New Forms of Damage in Negligence' (2007) 70(1) *The Modern Law Review* p. 59, 79–80.

851 Purshouse, 'Should Lost Autonomy be Recognised as Actionable Damage in Medical Negligence Cases?' (2015) 65–66, commenting on the fact that there is no independent right to autonomy under orthodox negligence actions.

852 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [50]–[55].

(*Lumba*) v *Secretary of State for the Home Department*,⁸⁵³ Davis LJ highlighted the necessity of an egregious violation of constitutional rights for such an award.⁸⁵⁴ Whatever developments have taken place, the autonomy interest is still not of a type that would meet these demanding requirements.

In sum, it would be much more controversial to recognise a right to autonomy, which could provide the basis for a vindictory impulse, than to acknowledge that autonomy is a head of loss that can be compensated according to established principles. The compensatory interpretation represents an incremental evolution that is consistent with the case law and is situated historically and structurally within negligence's wider requirements. In the following, this interpretation of the conventional award will be assumed. As referred to above, this is likely to allow for the wider relevance of this head of damage – including situations involving medical AI. Yet, the feasibility of such an argument must ultimately rest on the nature of the loss of autonomy that is being compensated.

ii. The autonomy interest

The other feature of the damage that must be addressed is its characterisation as an interference with autonomy. *Rees* itself leaves little doubt that the damage was conceptualised as the parents' – specifically, there, the mother's – loss of autonomy. Several quotes from the judgment have already been adduced to this effect, whether one refers to the parents' ability to live their life in the way they planned (as stated by Lord Bingham),⁸⁵⁵ or more directly to a denial of their autonomy (as outlined by Lord Millett).⁸⁵⁶

853 *R (Lumba) v Secretary of State for the Home Department* [2011] UKSC 12, [2012] 1 AC 245.

854 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [53].

855 *Rees v Darlington Memorial Hospital NHS Trust* [2003] UKHL 52, [2004] 1 AC 309 [8].

856 *ibid* [123].

That this is the type of damage upon which *Rees* was based is also widely reiterated in the commentary⁸⁵⁷ and echoed in the cases.⁸⁵⁸

The truly controversial prospect is the scope of the autonomy head of damage. Different positions have been defended in this regard. Some would restrict the new head of loss to the very specific circumstances of *Rees*.⁸⁵⁹ Others have argued that it was the specific loss of reproductive autonomy that underlay the award.⁸⁶⁰ Others again find that it was explicitly an award for interference with autonomy, albeit it was only a first step towards defining this loss in a manner that would be operational within the law.⁸⁶¹

857 Nolan, 'New Forms of Damage in Negligence' (2007) 70(1) *The Modern Law Review* p. 59, 78; Purshouse, 'Liability for Lost Autonomy in Negligence: Undermining the Coherence of Tort Law?' (2015) 22(3) *Torts Law Journal* p. 226, 229-233; Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) *Legal Studies* p. 55, 67-68; 'Certainly, interference with autonomy can sometimes be compensated in damages': Todd, 'Common Law Protection for Injury to a Person's Reproductive Autonomy' (2019) 135 *Law Quarterly Review* p. 635, 647. See also Priaux, 'Joy to the World! A (Healthy) Child Is Born! Reconceptualizing Harm in Wrongful Conception' (2004) 13(1) *Social & Legal Studies* p. 5, 16.

858 '[The] search for an award to compensate for the 'real loss' culminated in the recognition, in *Rees*, of a novel head of damage: that for a loss of autonomy': *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [107]. See also the claimant's argument in *Khan v Meadows* that the wrongful birth claim at issue 'should not be characterised as pure economic loss but as a mixed claim which combined her loss of autonomy through the continuation of the pregnancy and psychiatric damage incidental to her son's disability as well as her claim for the cost of caring for [him]': *Khan v Meadows* [2021] UKSC 21, [2022] AC 852 [21].

859 Davis LJ appeared to take this approach in: *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [78]-[79].

860 '[I]nterference with reproductive autonomy should be supported as a principled head of damage': Todd, 'Common Law Protection for Injury to a Person's Reproductive Autonomy' (2019) 135 *Law Quarterly Review* p. 635, 648. Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 371-374. If this is the opinion one takes, that autonomy does or can provide the basis for more specific heads of damage, then the subsequent discussion on AI can also be read in this light. Given the multifaceted nature of its capabilities, there are likely to be circumstances where the technology's use touches on the reproductive choices of parents in artificial treatment. As this argument is not accepted here, this aspect will not be pursued further.

861 Keren-Paz in Barker, Fairweather and Grantham, *Private Law in the 21st Century* (2017). The court in *ACB* interpreted *Rees* in this way: 'This search for an award to compensate for the "real loss" culminated in the recognition, in *Rees*, of a novel head of damage: that for a loss of autonomy': *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [107].

It is arguably the third position that can command the most support from *Rees*. Although the factual circumstances before the court led to a framing of the autonomy interest by reference to reproductive choice and parenthood, the language that was used is clearly applicable to a wider sphere of decision making. For instance, the notion that individuals may be deprived of the ability to shape and live their lives in the way they want, has wider, almost pervasive, relevance to medical decisions. This is precisely what was encapsulated under the reflective dimension of decisional autonomy in Chapter 3. Indeed, looking back to *McFarlane* Lord Millet explicitly stated that the parents have ‘lost the freedom to limit the size of their family. They have been denied an *important aspect* of their personal autonomy’.⁸⁶² These analyses understand reproductive freedom as particularly significant *instantiations* of autonomy – as important, as capable of determining the course of a life – they do not purport to establish a separate *type*.

Furthermore, a significant indicator that the head of damage must be wider than merely reproductive autonomy was provided by the House of Lords only a few years later in *Chester v Afshar*.⁸⁶³ As has been noted in Chapter 4, this case concerned the failure of a surgeon to disclose a specific risk to a patient before a surgery was conducted on her spine. The risk eventuated and the patient claimed in negligence, even though she admitted that, had she been warned of the risk, she would still have opted for the surgery at a later date.⁸⁶⁴ Although a single rationale is difficult to discern, the majority’s reasoning appears to have accepted that there was a breach of duty (in the form of a failure to obtain informed consent) that was connected to the personal injury suffered, just not *via* traditional causation principles.⁸⁶⁵ However, given the significance of patient autonomy, their Lordships provided for a limited exception to orthodox causation

862 *McFarlane v Tayside Health Board* [2000] 2 AC 59, 114 (my emphasis).

863 *Chester v Afshar* [2004] UKHL 41, [2005] 1 AC 134.

864 *ibid* [7].

865 Referring to a commentator’s discussion of a comparable Australian case, Lord Steyn stated: ‘Professor Honoré was right to face up to the fact that *Chappel v Hart* — and therefore the present case — cannot neatly be accommodated within conventional causation principles. But he was also right to say that policy and corrective justice pull powerfully in favour of vindicating the patient’s right to know’: *Chester v Afshar* [2004] UKHL 41, [2005] 1 AC 134 [22].

principles to allow for recovery.⁸⁶⁶ Ultimately, the patient was compensated for the full amount of her personal injury.

It is hard to avoid the conclusion that the real damage that was suffered in this case was to the patient's autonomy.⁸⁶⁷ The House of Lords identified that there was an injury to autonomy, stemming from the physician's failure to sufficiently facilitate the making of a significant medical decision, and sought to provide compensation for it. Autonomy must therefore be a head of damage that is wider than the circumstances in *Rees* or the reproductive autonomy of potential parents.

In light of this, the discussion of causation between the breach and the personal injury should have been superfluous. As Keren-Paz has noted, the patient suffers an interference with autonomy (the damage) by being denied an opportunity to consent (in an informed manner) and there is no issue in finding a causal connection with the breach here.⁸⁶⁸ The absence of a conventional award and the overcompensation for the personal injury in *Chester* then appear puzzling – why did the court not simply apply *Rees* to this situation?⁸⁶⁹

Several reasons may be adduced. Simply the fact that the eventuation of damage is rarely examined at any length and often subsumed under the other elements of negligence played its role. The cases' purported focus on, respectively, the quantification of damages and the causation of harm obscured any common ground between them.⁸⁷⁰ This confusion was argu-

866 Lord Steyn maintained that the patient's 'right of autonomy and dignity can and ought to be vindicated by a narrow and modest departure from traditional causation principles': *ibid* [24]. To similar effect Lord Hope stated: 'The function of the law is to protect the patient's right to choose. If it is to fulfil that function it must ensure that the duty to inform is respected by the doctor. It will fail to do this if an appropriate remedy cannot be given if the duty is breached and the very risk that the patient should have been told about occurs and she suffers injury': *ibid* [56].

867 Amirthalingam, 'Causation and the Gist of Negligence' (2005) 64(1) *The Cambridge Law Journal* p. 32, 34-35. Clark and Nolan outline the advantages of this view: Clark and Nolan, 'A Critique of *Chester v Afshar*' (2014) 34(4) *Oxford Journal of Legal Studies* p. 659, 684-685, 688-689. See also: Keren-Paz, 'Compensating Injury to Autonomy in English Negligence Law: Inconsistent Recognition' (2018) 26(4) *Medical Law Review* p. 585, 592-593.

868 Keren-Paz, 'Compensating Injury to Autonomy in English Negligence Law' (2018) 26(4) *Medical Law Review* p. 585, 592-593.

869 Clark and Nolan, 'A Critique of *Chester v Afshar*' (2014) 34(4) *Oxford Journal of Legal Studies* p. 659, 684; *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [114].

870 Purshouse, 'Judicial Reasoning and the Concept of Damage' (2015) 15(2-3) *Medical Law International* p. 155, 157.

ably amplified by the fact that a considerable form of personal injury had occurred in *Chester*, representing a traditional, readily accepted category of damage. It may have appeared convoluted to engage in a reframing of this loss. It has also been argued that the harm suffered in *Rees* was undercompensated because of its gendered nature.⁸⁷¹ Without such specific factors in play in *Chester*, the House of Lords approached the case differently and compensated the loss of autonomy more appropriately.

Given the diverging paths taken by the UK's highest court in these cases, and without a clear rationale, it is arguably not possible to resolve this inconsistency in the law. However, from a perspective that takes the autonomy principle seriously, the most cogent interpretation of *Chester* is that the House of Lords built upon *Rees*. The court thereby recognised an injury in the diminishment of the patient's autonomy, yet opted to compensate it with damages for personal injury.⁸⁷²

This indicates that UK law provides a patient with the ability to bring a claim on the basis of an interference with their general autonomy interest. However, it must also be recognised that a definition of this interest has not been forthcoming and lower courts have since expressed a reticence to invoke it. In particular, the Court of Appeal judgments of *Shaw v Kovac* and *Duce v Worcestershire Acute Hospitals NHS Trust* purport to limit the relevance of *Rees* and *Chester* severely. Both decisions refused to recognise an autonomy loss *per se* as a form of damage.⁸⁷³ Moreover, they purported to limit both cases to their facts or, at least, to a very specific fact pattern.⁸⁷⁴

871 Priaulx, 'Joy to the World! A (Healthy) Child Is Born! Reconceptualizing Harm in Wrongful Conception' (2004) 13(1) Social & Legal Studies p. 5, 10-15.

872 Amirthalingam has argued that the 'the physical injury, which turned on the patient's response, merely went to the quantification of the loss': Amirthalingam, 'Causation and the Gist of Negligence' (2005) 64(1) The Cambridge Law Journal p. 32, 33-34. Keren-Paz argues for a similar resolution of this issue, although our approaches as to how the new head of damages should then be defined within acceptable bounds differ: Keren-Paz in Barker, Fairweather and Grantham, *Private Law in the 21st Century* (2017) 428-437. Note also how crucial the difference between actionable damage and damages becomes in this analysis.

873 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [58]-[74]; *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [88].

874 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [61], [79]; *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [87]-[90]. The latter case especially highlighted the exceptional nature of *Chester v Afshar* and the role that policy factors played in determining an unorthodox approach.

From a purely formal perspective these cases cannot overturn the above decisions – they remain bound by the higher courts’ judgments.⁸⁷⁵ If *Rees* and *Chester* do carve out an autonomy interest that is actionable, as it has been argued, then the lower courts’ refusal to apply it constitutes a practical problem. The Court of Appeals’ findings would be more troublesome for the present analysis if they appealed to broader, persuasive normative grounds for rejecting the autonomy claim. But such arguments were not explicitly provided. The *ratio* of these cases has been interpreted to be very narrow.⁸⁷⁶ Furthermore, the broader, *obiter* argumentation is undoubtedly of a questionable calibre.⁸⁷⁷

In fact, the strongest normative case for rejecting the view that an independent autonomy injury was created, is provided by the aforementioned case of the Singaporean Court of Appeal: *ACB*. Although this only constitutes persuasive authority in England, the court arguably detailed criticisms that were implicitly touched upon in *Shaw* and which have carried great weight with British commentators.⁸⁷⁸

In essence, one can identify two kinds of objections in this judgment: one conceptual argument, that goes towards the nature of the legal (non-)identification of autonomy, and one class of more principled objections: going towards the compatibility of a moral-/political- concept with the structure of the common law, and specifically the requirement of damage in negligence. The assumption is that a narrower aspect of that concept (such as reproductive autonomy) better satisfies the requirements

875 Leggatt LJ highlighted that the Supreme Court may need to reconsider this area of the law: *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [92].

876 For a full argument see: Keren-Paz, ‘Compensating Injury to Autonomy in English Negligence Law’ (2018) 26(4) Medical Law Review p. 585, 599-602. Similarly Todd notes that Shaw ‘holds at least that the notion of injury to autonomy as a new head of loss in a personal injury claim should not be supported, but that says nothing about the implications of the decision in *Rees*’: Todd, ‘Common Law Protection for Injury to a Person’s Reproductive Autonomy’ (2019) 135 Law Quarterly Review p. 635, 648.

877 Purshouse, ‘Autonomy, Affinity, and the Assessment of Damages’ (2018) 26(4) Medical Law Review p. 675, 681-684.

878 *ibid* 684-687; Nolan, ‘Negligence and Autonomy’ [2022](2) p. 356, 364-371; Mulligan, ‘A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction’ (2020) 40(1) Legal Studies p. 55, 69-70.

of conceptual clarity and (partially as a result) can better fulfil its function as a head of damage.⁸⁷⁹

Here we need not devote too much time to the basis for the conceptual objection, given Chapter 4's extensive treatment of the issue. It was seen that different types of autonomy are discussed in the legal sources, as well as in the wider literature, and/or are incorporated into them. What is notable about the specific criticism of Phang JA, which is echoed and amplified by Nolan,⁸⁸⁰ is the sense that arriving at a sufficiently settled definition is impossible for autonomy.⁸⁸¹ The positions reflected in the law are simply too diverse and philosophically and politically contested.⁸⁸²

Curiously, there is also a sense that, if a sufficiently settled legal definition could be reached, then it would be a highly individualistic conception of autonomy that would contribute to the second, principled objection.⁸⁸³ In particular, it would reinforce the argument that one cannot objectively assess whether a damage has eventuated, given the inherent subjectivity of autonomy.⁸⁸⁴ I will address these different objections in turn, in light of the procedural principle of autonomy that has been adopted throughout this work.

To begin with, one must acknowledge that, undoubtedly, there is not one uncontested concept of autonomy in UK law. Yet, this does not mean that the common law could not adopt a defensible understanding of its nature and influence, providing an adequate degree of fit with the legal material.

879 On the interrelation between the two types of criticism: *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [119].

880 Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 364-366.

881 *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [116]-[119].

882 *ibid* [119].

883 Purshouse explicitly notes that in English law 'it is the current desire version of autonomy (...) that the courts presently use in a number of contexts': Purshouse, 'How Should Autonomy Be Defined in Medical Negligence Cases?' (2015) 10(4) *Clinical Ethics* p. 107, 111. This approach was relied on in *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [120].

884 Both Purshouse and ACB note that an autonomy violation may (objectively) leave a person no worse, no different, or even better, off: Purshouse, 'Liability for Lost Autonomy in Negligence' (2015) 22(3) *Torts Law Journal* p. 226, 237; *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [120]. Similarly in *Shaw*, it was found objectionable that a patient could recover for damage to autonomy even where the operation 'was a complete success': *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [71]. Nolan provides a convincing refutation of this argument, citing both the fact that damage in negligence does not require being worse off and that it ignores the possibility of the autonomy interest itself being the damage: Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 364.

Chapters 3 and 4 offered one coherent concretisation in this regard. Conceived of as a legal principle, a procedural conception of autonomy was argued to command considerable support in the legal system – particularly with a view to the medical context and the issues of informed consent. As we will return to below, adopting such a definition does not allow for the delineation of a qualitatively distinct concept of reproductive autonomy.

To respond to the objection that autonomy does not fit easily into the damage element, an established legal category, one ought to begin by rejecting the assumption that the law must rely on a purely subjective account of autonomy. This is precisely the kind of conception that generates the problems to which the proponents of a restrictive approach then object. Keren-Paz has already made the argument that autonomy is not self-evidently so subjective as to prevent the identification of an interference.⁸⁸⁵ Verifiable cases of significant violations patently do exist, and this is obscured by both Phang JA's and Purshouse's reference to an allegedly indivisible class of autonomy violations that do not meet a *de minimis* threshold of seriousness.⁸⁸⁶ The task therefore becomes the identification of a yardstick by which the law can limit itself to addressing only these interferences.

As has been pointed out by Chico, an understanding of autonomy that incorporates objective or ideal elements is much more suited to this task.⁸⁸⁷ This is true of our theoretical account. Autonomy violations were to be assessed by reference to objective, or objectively verifiable, elements, including: the need for certain classes of information (necessary true beliefs) and the centrality of the patient's established commitments and character under the reflective dimension of autonomy. To make out the significance of an interference, the patient should be able to adduce these.

885 Keren-Paz in Barker, Fairweather and Grantham, *Private Law in the 21st Century* (2017) 432-433.

886 *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [120]; Purshouse, 'Liability for Lost Autonomy in Negligence' (2015) 22(3) *Torts Law Journal* p. 226, 234-242.

887 Especially: 'the English courts would be more willing to recognize that a novel interest, namely autonomy, might form the basis for legally recognized harm where the interest is amenable to limitation by the courts. An ideal conception of autonomy as imbued with rationality, as opposed to a liberal conception which rests simply on capacity and independence, would allow greater analytical purchase on what autonomy actually consists in, thereby providing potential boundaries to the legal recognition of the interest. In this way, the courts would be able to prioritize those interferences with autonomy which they perceive to be the most deserving of legal recognition': Chico, *Genomic Torts: The English Tort Regime and Novel Grievances* (2010) 47-49.

Moreover, as Keren-Paz has argued, the different consequences of interferences for a patient's practical decision will also shape its significance: if a patient would have made a different decision (as in *Rees*) then the interference is arguably stronger, if they were denied the opportunity to reflectively affirm that position (as in *Chester*) it is arguably weaker.⁸⁸⁸ Such an approach tracks the distinction drawn in Chapter 3 between the practical element of autonomy and the decisional one. Ultimately, our theories' components, and specifically English law's understanding of them, seek to provide a yardstick for the objective identification of significant autonomy violations. If this is accepted, then the necessary legal specification of autonomy damage is not an insurmountable hurdle.

The courts' and commentators' treatment of reproductive autonomy also suggests as much. For, there is a sense that *Rees* was defensibly decided for the purposes of remedying a specific injury to reproductive autonomy.⁸⁸⁹ Hence it is puzzling how critics envision the more specific aspect of autonomy to overcome the conceptual and objectivity objection. Does it really assist in an objective evaluation to say that autonomy can 'underlie'⁸⁹⁰ or 'be used as a justification for'⁸⁹¹ a specific head of damage. A definitional choice is still needed. This should be made consistently and it should comport with the requirements of the law.

The only difference in cases involving reproductive decisions is that judges and commentators can be relatively confident that there has been an interference with the patient's deepest desires,⁸⁹² their long-term plans,⁸⁹³ etc. There is a relatively well-documented process (sometimes involving a contraceptive intervention) leading up to their decision and capturing their commitments. If this is the approach to autonomy, which the present author would endorse as an outgrowth of its reflective dimension, then a limitation to reproductive choice is unsatisfactory and, beyond allusions to

888 See Keren-Paz's analysis on this issue: Keren-Paz, 'Gender Injustice in Compensating Injury to Autonomy in English and Singaporean Negligence Law' (2019) 27(1) *Feminist Legal Studies* p. 33, 44-45.

889 The court in *ACB* were not bound to follow *Rees*, and distinguished it in other ways, but still argued for an award based on an aspect of autonomy: *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [125]-[130].

890 *ibid* [115].

891 Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 371.

892 Keren-Paz in Barker, Fairweather and Grantham, *Private Law in the 21st Century* (2017) 427.

893 *ibid* 415.

practicality, no rationale for it is forthcoming.⁸⁹⁴ As Austin has highlighted: where a relevant clinical decision involving non-disclosure is brought before a court, the claimant's position must be presented, justified and analysed in a way that avoids attributing hindsight to them.⁸⁹⁵ If they can do this outside of the reproductive context, then it should not be for the law to say that there loss is not sufficiently 'objective'.

If an objective delineation of autonomy is possible, then one must further consider its interaction with established rule-specific categories. For example, commentators have pointed out the challenges of disentangling a violation of autonomy from other kinds of damage – the concern here is the prospect of double recovery.⁸⁹⁶ However, double recovery is only objectionable if autonomy is not seen as a harm in itself, which proponents of autonomy would deny, or if courts mistakenly overcompensate it, as occurred in *Chester*.⁸⁹⁷ In addition, this consequence is not inevitable. As Keren-Paz has highlighted (and Nolan submits) it is open to the court to amend *Chester* and to refuse additional recovery for autonomy where an existing head of damage is already compensated.⁸⁹⁸ We will return to this discussion below, in our assessment of available damages.

In the final analysis it is therefore argued that our development of the autonomy principle affirms the position that the English negligence action provides a compensatory award for certain, significant and verifiable, violations of patient autonomy. Nevertheless, in light of the House of Lords'

894 '[W]hereas previously autonomy was only protected indirectly by the law of negligence (via claims for personal injury and so forth), increasingly loss of autonomy appears to be being accorded recognition as a form of actionable damage in its own right. (...) Since autonomy is a very important interest, this development is to be welcomed, although if negligence liability is to be kept within acceptable limits, protection can only be accorded, as at present, to certain derivative autonomy interests – freedom of movement, reproductive autonomy and so forth – rather than to autonomy in the round': Nolan, 'New Forms of Damage in Negligence' (2007) 70(1) *The Modern Law Review* p. 59, 87.

895 Austin, 'Correia, Diamond and the Chester Exception: Vindicating Patient Autonomy?' (2021) 29(3) *Medical Law Review* p. 547, 559.

896 Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 366. See also *ACB v Thomson Medical Pte Ltd* [2017] SGCA 20 [123]-[124].

897 Clark and Nolan, 'A Critique of *Chester v Afshar*' (2014) 34(4) *Oxford Journal of Legal Studies* p. 659, 684.

898 Keren-Paz, 'Compensating Injury to Autonomy in English Negligence Law' (2018) 26(4) *Medical Law Review* p. 585, 590 (fn. 37), 600-601. Much depends on how the autonomy concept is shaped and limitations can still be imposed at the duty of care stage: Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 367.

inconsistent rulings and the Court of Appeal's hesitancy to recognise and implement such an approach, its strength is not beyond question. Proactive steps by the courts, most likely even an intervention by the UK's Supreme court, would be necessary to clarify the legal position. At the same time, given how recently the autonomy interest has emerged as a significant influence on UK jurisprudence,⁸⁹⁹ it should also not be surprising that a concrete approach has yet to be fully specified. For now, our autonomy concept offers one defensible interpretation of the still open-textured legal position.

3. Summation

Going forward it will be argued that UK law can compensate a patient for significant forms of autonomy violations, as well as for the personal injury, that they suffer as a result of AI use. At points it will be necessary to distinguish between the two forms of damage, since it is beyond doubt that a claim for personal injury is more likely to succeed. As our discussion of *Chester* has demonstrated, the recognition of an autonomy interest will also shape the other elements of the analysis, including causation. Overall, loosening the law's insistence on physical damage will allow for a more coherent remedy to be provided in response to violations of the patient's informed consent.

B. Duty of care

A duty of care specifies those relationships where purely inadvertent and unintentional conduct constitutes an appropriate basis for liability. As such it is an integral element in the delineation of the scope of liability of a negligence action. This is encapsulated by the fact that the existence of a duty of care is based primarily upon established categories of case, where a class of defendants has already been found to owe a duty with respect to a certain class of claimants. If the relevant constellation falls within such a category, then this is normally the end of the analysis: a duty of care is

899 'As recently as the 1970s, for example, commentators framed the issue of "informed consent" in the medical context as a question of dignity, rather than autonomy': Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 382.

established. Alternatively, if the court is confronted with a situation that falls outside of existing classifications, then a duty of care may still be made out, but only *via* incremental development. Here the duty is extended beyond existing categories on the grounds that it would be fair, just and reasonable to do so.⁹⁰⁰

In the following we will consider the duties owed by two classes of defendants with respect to AI use: medical professionals and healthcare institutions. These are the actors who are anticipated to instrumentalise AI in the English healthcare system and who are likely to be the primary targets for negligence actions concerning the provision of medical information. Parties that also owe some informational duties of care to the patient, but which will not be investigated further below, are AI developers and manufacturers. Rather than conducting a separate analysis of such duties, this aspect of the designers' responsibility will be considered only in so far as it shapes professional and institutional breaches.⁹⁰¹ This is justified by the fact that the manufacturer's informational, facilitative role to the patient is mediated by healthcare professionals and institutions and that their obligations are more narrowly defined: to ensure that the user of a product is adequately warned of its risks.⁹⁰² It is not arguable that these duties can be extended to cover the unique properties of ML devices.

900 *Caparo Industries Plc v Dickman* [1990] 2 AC 605. That this is the correct order of consideration, focussing first on the established situations where a duty is owed and then potentially seeking to expand upon these incrementally, was stated *obiter* in *Michael v Chief Constable of South Wales Police* [2015] UKSC 2, [2015] AC 1732 [102]-[111] and affirmed in *Darnley v Croydon Health Services NHS Trust* [2018] UKSC 50, [2019] AC 831 [15]-[16].

901 This approach finds the strongest support in existing case law 'Where the treatment involves a medical device there is a duty on the doctor to apprise themselves of the implications of the use of that device together with any risks associated with its use. That again is part of the clinician's duty to the patient but its importance is underlined by the doctor's role as a learned intermediary between the producer of the product and the patient. The producer has no relationship with the patient who relies on the doctor to advise them of the risks associated with it': *AH v Greater Glasgow Health Board* [2018] CSOH 57, (2019) 169 BMLR 120 [61]. *AH* also cited another recent authority that dealt with the disclosure of information required from the manufacturer of a prosthetic hip. Here the court held that 'Parliament has determined that, in relation to such products, information about risks is best relayed to, considered by, and applied and passed on to the patient by the treating surgeon, who must advise that patient as to intervention choices, and seek and obtain that patient's informed consent the particular chosen implant procedure': *Wilkes v DePuy International Ltd* [2016] EWHC 3096, [2018] QB 627 [107].

902 Goldberg, *Medicinal Product Liability and Regulation* (2013) 58-61.

1. Medical professionals

Regarding medical professionals, it is trite law that they owe duties to their patients.⁹⁰³ And it will be seen that, beyond any doubt, these duties encompass demanding informational components to those they are diagnosing, advising and treating.

If a medical professional utilises an ML device, then an enquiry into their informational responsibilities does not call for a reconsideration of the existence of this duty. Rather, it problematises the actions that are necessary, with respect to that technology, to ensure compliance with it. It is primarily a question of whether that duty was breached.

UK courts do not appear to have limited this obligation to a particular class of medical professionals. They have considered whether the treatment team as a whole has obtained the necessary consent from the patient. For example, in *Lybert v Warrington Health Authority* the patient's gynaecologist was found negligent for a failure to warn, although he was not the only individual whom the claimant had 'sought to criticise' and although he was 'entitled to assume, or at least to expect, that an initial warning had been given somewhere along the line'.⁹⁰⁴ Similarly, it was stated in *Montgomery v Lanarkshire Health Board* that:

a wider range of healthcare professionals now provide treatment and advice of one kind or another to members of the public, either as individuals, or as members of a team drawn from different professional backgrounds (with the consequence that, although this judgment is concerned particularly with doctors, it is also relevant, *mutatis mutandis*, to other healthcare providers).⁹⁰⁵

Consequently, where ML devices are utilised in the care process, there may be an obligation on several professionals to obtain the patient's informed consent and, in this respect, no straightforward delineation can be drawn between them at the duty stage.⁹⁰⁶

903 *Barnett v Chelsea and Kensington Hospital Management Committee* [1969] 1 QB 428.

904 *Lybert v Warrington* (1995) 25 BMLR 91, 95.

905 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [75] (my emphasis).

906 As will be discussed under the breach element, liability must of course still depend on the knowledge that the professional in question did have or should reasonably have had.

2. Healthcare institutions

Regarding healthcare institutions, the position is admittedly more complex since there are several routes by which they can be held legally responsible. The nature of organisational liability, and its potential significance in situations where organisations opt to rely on AI, therefore bears more detailed elaboration.

One route for establishing organisational liability in negligence takes the form of vicarious liability. At least since *Cassidy v Ministry of Health* and *Roe v Minister of Health* it has been established that healthcare providers can be held liable for the actions of their medical and non-medical staff in this fashion.⁹⁰⁷ In such an instance it is the relationship between two identifiable individuals, such as healthcare professional and patient, that founds the duty. The organisation is then held strictly liable for the negligence of the tortfeasor because the two are connected in the requisite manner.⁹⁰⁸ This aspect of liability should be noted for its practical relevance but need not be dwelled on here: if informational duties are owed by professionals in respect of AI, then there is no reason the technology would prevent the employer from being held vicariously liable as well.

In addition to vicarious liability, liability may alternatively be established on the basis that an institution owes a direct duty of care to the patient. For this mechanism, accountability does not depend on the mediation of an identifiable human professional with a relevant duty.⁹⁰⁹ For example, it is settled law that a healthcare institution can be held directly liable where there is some organisational or managerial failure; where the negligence can be ascribed to ‘no one and nothing but the system itself’.⁹¹⁰ Thus in *Bull v Devon Area Health Authority* the system for summoning obstetricians

907 *Cassidy v Ministry of Health* [1951] 2 KB 343; *Roe v Minister of Health* [1954] 2 QB 66.

908 Deakin, ‘Organisational Torts: Vicarious Liability Versus Non-Delegable Duty’ (2018) 77(1) *The Cambridge Law Journal* p. 15, 17-18.

909 Syrett distinguishes between direct and non-delegable duties: Syrett in Laing and McHale, *Principles of Medical Law* (Fourth Edition 2017) 379. I have in mind the former.

910 *Bull v Devon Area Health Authority* (1989) 22 BMLR 79, 100. In a similar vein Lord Phillips MR, in another case, spoke of ‘A duty to use reasonable care to ensure that the hospital staff, facilities and organisation provided are those appropriate to provide a safe and satisfactory medical service for the patient’: *A (a child) v Ministry of Defence* [2004] EWCA Civ 641 [32].

for the provision of maternity services was found wanting and a negligence claim against the hospital was made out.⁹¹¹

As addressed in Chapter 2, AI implementation will generally be premised on a degree of expert mediation. Typically, informational duties would therefore fall to be considered under the duties owed by human professionals. Primary institutional liability will only gain relevance in a restricted set of circumstances. Namely, where the institution deploys AI to determine an aspect of the patient's medical care. The example offered for this in Part I. was the use of AI in triage. ML devices would be giving patients medical and/or organisational data – such as a preliminary assessment of their condition and information on which further healthcare services to access (if any) and how to access them. If any informational duties are to be asserted here, then the role of the corporate person becomes paramount.

A recent case of the UK Supreme Court, *Darnley v Croydon Health Services*, suggests that institutions do owe some such duties.⁹¹² Specifically it was held that providing the patient with information on the timeframe within which they could expect to be seen in an accident and emergency department fell within an NHS Trust's primary duty.⁹¹³ This case emphasised the corporation's general responsibility to ensure that there is an adequate system of information provision in place. In this respect, Lord Lloyd-Jones JSC expressed agreement with a dissenting judge of the Court of Appeal, McCombe LJ, indicating that: 'The failure to impart the reality of the triage system to the claimant on his arrival was, on the facts of this case, a breach of duty by the hospital'.⁹¹⁴ It was seemingly irrelevant that the actionable misinformation – that the patient would be seen in four-to-five hours by a doctor, rather than in 30 minutes by a triage nurse – was the result of an individual receptionist's mistake and it was explicitly stated that the duty could be discharged *via* medical staff, non-medical staff, pamphlets or notices.⁹¹⁵

911 *Bull v Devon Area Health Authority* (1989) 22 BMLR 79.

912 *Darnley v Croydon Health Services NHS Trust* [2018] UKSC 50, [2019] AC 831. The Supreme Court also drew on an earlier Court of Appeal finding that an ambulance driver was liable for provided misleading assurances concerning the arrival time of an ambulance: *ibid* [18]-[20], citing *Kent v Griffiths (No.3)* [2001] QB 36.

913 *Darnley v Croydon Health Services NHS Trust* [2018] UKSC 50, [2019] AC 831 [16] - [17], [21].

914 *ibid* [13], [19].

915 *ibid* [26]-[27].

Consequently, an institution owes a direct duty to patients that covers one aspect of information provision. However, it would be a mistake to equate this with the obligation of medical professionals to obtain informed consent.⁹¹⁶ At the risk of conflating the duty and breach analyses, it is worth clarifying already at this stage that the obligation defined in *Darnley* is not suited to the task of alerting the patient to the unique aspects of AI treatment.

Such a limitation emerges from the Supreme Court's focus on information with the potential to cause physical injury.⁹¹⁷ The court explicitly framed the duty to be 'one to take reasonable care not to cause physical injury to the patient'.⁹¹⁸ Similarly, it was held that 'it is not the function of reception staff to give wider advice or information in general to patients', the NHS Trust must simply 'take care not to provide misinformation to patients'.⁹¹⁹ It appears that it was the close relationship between the provision of misleading information and a direct danger to the patient's health that shaped the Supreme Court's finding.

In the case of an ML triage system, such a relationship is arguably lacking, or exists only to a very limited extent. Indeed, it is where a patient is not using an AI, that it is arguable that a healthcare institution must alert the patient to the availability of such an automated triage system, which could enable faster access to the necessary care. In this sense the ML device would assume the role of the triage nurse to whom the patient in *Darnley* ought to have been alerted.⁹²⁰ By contrast, it cannot be maintained that a patient who uses an AI, or is prepared to do so, must be advised of its specific risks or even its risk profile – never mind its ability to influence value judgments or its relationship to human expertise. To enable the patient to avoid suffering physical harm it would clearly suffice to warn them not to treat the AI output as definitive and, if concerned, to seek further medical assistance. As with AI manufacturers' and developers' duty to warn, such an obligation is a patently unsatisfactory basis for the promotion of patient autonomy and informed decision making.

916 Purshouse, 'The Impatient Patient and the Unreceptive Receptionist' (2019) 27(2) Medical Law Review p. 318, 329.

917 For a similar framing of the matter see: Armitage, Charlesworth and Percy, *Charlesworth & Percy on Negligence* (Fifteenth Edition 2022) para 10-178.

918 *Darnley v Croydon Health Services NHS Trust* [2018] UKSC 50, [2019] AC 831 [16].

919 *ibid* [19].

920 *ibid* [26].

Limited support for a marginally more extensive institutional disclosure duty can be found in the earlier case of *Lybert v Warrington Health Authority*. Here the patient had not been warned of the risk that a sterilisation procedure performed on her may not have been effective and Otton LJ was prepared to find the health authority primarily liable in the circumstances.⁹²¹ It was held that there was ‘a duty upon those responsible for the conduct of this unit to ensure that there was a proper and effective system for giving a proper warning at some stage during her time as a patient’.⁹²²

This duty is arguably broader than the one found in *Darnley*. It requires the affirmative provision of a class of information which did not present an immediate danger to the patient’s health. However, it is again noticeable just how narrowly the duty is framed. It essentially constitutes a duty to warn of specific risks and therefore still maintains a strong connection to the patient’s physical well-being. As will be seen below, a duty to disclose AI’s specific risks does not in any way begin to address the technology’s novel autonomy challenges. Furthermore, to fulfil this obligation, it was sufficient to avoid ‘having a lax system in place that allowed the plaintiff to “slip through the system”’.⁹²³ In consequence, *Lybert* does not do much to advance the argument that institutions bear responsibility for facilitating the patient’s informed decision making. Instead, they primarily bear responsibility for the protection of their health.

In conclusion, there is some basis for finding that, in England, institutional liability does extend to minimal informational obligations, at least in so far as they are relevant to the protection of the patient’s physical well-being. If appropriately developed, direct institutional liability could provide one useful legal mechanism for imposing obligations on institutions, which are utilising relatively independent ML models, to let patients know that they are being processed by an AI with certain goals, capabilities etc. Given the recent changes to the breach element in non-disclosure cases, there is a possibility that a recalibration may occur in the future. However, currently there are no existing indications that the autonomy principle has shaped the evolution of this duty. Quite the opposite: it is a concern for the patient’s beneficence that has determined a narrow framing of relevant obligations. In light of this, it would be difficult to construct a disclosure

921 *Lybert v Warrington* (1995) 25 BMLR 91, 92.

922 *ibid* 95.

923 Grubb, ‘Failed Sterilisation: Duty to Provide Adequate Warning’ (1995) 3(3) Medical Law Review p. 297, 298.

obligation that could ensure the protection of the patient's autonomy in the AI context.

3. Summation

For most uses of AI, the duty we should consider is that of healthcare professionals providing their services to a patient with the aid of an ML device. Here it is well established that informational duties are owed to patients. The existence of a duty is certainly not called into question by reliance on a new technology, although it may shape its discharge. By contrast, in the scenario where an AI device functions systematically as an information conveyer, determining an aspect of the patient's care, it has been argued that the institutional provider owes only limited direct duties of care and that these are not suitable vehicles for the protection of the patient's autonomy.

C. Breach

When a duty has arisen, the next question goes towards the circumstances under which the relevant party will fall short in complying with it. It is a defining aspect of negligence liability that the defendant is not strictly liable but must, in their actions, meet a certain standard of care. This standard has long been defined by reference to the reasonable person:

Negligence is the omission to do something which a reasonable man, guided upon those considerations which ordinarily regulate the conduct of human affairs, would do, or doing something which a prudent and reasonable man would not do.⁹²⁴

In other words, in order to succeed a claimant must usually show that a defendant fell short of the standards expected of this objective reasonable person.

Given the difficulty of judging the reasonableness of the actions of specialist actors, the British courts have traditionally provided a deferential test

924 *Blyth v Birmingham Waterworks Co* (1856) 156 ER 1047, 1049.

to make this assessment in clinical negligence.⁹²⁵ Under the seminal case of *Bolam v Friern Hospital Management Committee* it is enough for a medical professional to show that they acted in accordance with one responsible body of medical opinion to demonstrate that they met the reasonableness standard.⁹²⁶ Following *Bolitho v City and Hackney Health Authority* one may add the caveat, in line with the reasonableness standard, that the professional position must withstand logical scrutiny by the courts.⁹²⁷

The specific question for our purposes is: what kind of information must be provided about the involvement of AI in a patient's care, in order to avoid breaching a healthcare professional's duty of care? Here, the relevant obligations have been, and continue to be, subsumed under the ordinary negligence action.⁹²⁸ However, the value of patient autonomy has brought about an alteration in the standard of care governing cases that concern the patient's informed consent. This section defines this standard, identifies its requirements in operation and applies these to medical AI/ML.

1. The informed consent standard

The prevailing standard of care in nondisclosure cases was, for a long time, the same as the aforementioned *Bolam* standard applicable to medical negligence in general, establishing a uniform approach. This was the finding in *Sidaway v Board of Governors of the Bethlem Royal Hospital*.⁹²⁹

925 As has been argued: this standard 'does not represent a departure from the ordinary standard of the reasonable person, but is merely an attempt to apply that standard in the light of the fact that judges lack the knowledge and expertise usually required to choose between competing professional opinions': Nolan, 'Varying the Standard of Care in Negligence' (2013) 72(3) *The Cambridge Law Journal* p. 651, 654-655.

926 *Bolam v Friern Hospital Management Committee* [1957] 1 WLR 582.

927 *Bolitho v City and Hackney Health Authority* [1998] AC 232, 243.

928 It has been argued however that the novel emphasis on the autonomy of the patient in the definition of the standard of care (see *infra*) should be recognised as having created a *sui generis* cause of action: Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 376-382; Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) *Legal Studies* p. 55, 71-74. Even if this would be a preferable way to frame the courts' anomalous approach to medical non-disclosure cases, this is not purported to represent current practice.

929 'In English jurisprudence the doctor's relationship with his patient which gives rise to the normal duty of care (...) has hitherto been treated as single comprehensive duty covering all the ways in which a doctor is called upon to exercise his skill and

Only in 2015 did the Supreme Court, in *Montgomery v Lanarkshire Health Board*, overrule this established norm on the strength of the autonomy principle.⁹³⁰ By appealing to the patient's need for material information, this case established a patient-centred standard of disclosure, under which the courts are 'developing separate rules to those governing the rest of medical negligence'.⁹³¹

i. The meaning of reasonable disclosure

Montgomery v Lanarkshire Health Board sets the current standard for the steps that a health professional must take to advise a patient about the nature of their treatment. In an oft-cited passage, the majority held:

An adult person of sound mind is entitled to decide which, if any, of the available forms of treatment to undergo, and her consent must be obtained before treatment interfering with her bodily integrity is undertaken. The doctor is therefore 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.⁹³²

This statement encapsulates the remarkable shift in the influence attributed to patient autonomy in negligence actions dealing with issues of informed consent. A concern for this principle led to a transformation from a defendant friendly standard of care – where breach was analysed primarily by

judgment (...) This general duty is not subject to dissection into a number of component parts to which different criteria of what satisfy the duty of care apply, such as diagnosis, treatment, advice (including warning of any risks of something going wrong however skilfully the treatment advised is carried out.): *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 893. For an application of the *Bolitho* caveat to disclosure, see also: *Pearce v United Bristol Healthcare NHS Trust* (1999) 48 BMLR 118.

930 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430.

931 Purshouse, 'The Impatient Patient and the Unreceptive Receptionist' (2019) 27(2) Medical Law Review p. 318, 329.

932 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [87].

reference to the actions of a responsible body of professional opinion – to one where the standard of care is set ‘by reference to *the reasonable person in the position of the claimant*’.⁹³³ And, even more remarkably, by reference to the information that the specific patient deems significant.

In defining this novel approach, the Supreme Court employed rhetoric that, as outlined in Chapter 4, aligns with our understanding of procedural autonomy. Recognising an important prerequisite for the exercise of the patient’s practical autonomy, the court framed patients as individuals with the power to make medical decisions. They are ‘widely treated as consumers exercising choices’ and ‘widely regarded as persons holding rights’.⁹³⁴ These choices and these rights are not subject to the discretion of the medical profession but are located within a bureaucratic framework, subject to public law supervision.⁹³⁵ All in all, the patient is properly regarded as having a decision to make, an action to perform, and the court evidently recognised that it is their unique commitments that shape this decision.

Moreover, it is the medical profession that bears responsibility for facilitating the process leading up to this decision. In this respect, Lord Kerr and Lord Reed appealed to the positive, facilitative dimension of practical autonomy when they rejected indications from *Sidaway* that the disclosure standard could be determined by the patient’s preparedness to ask specific questions:

the more a patient knows about the risks she faces, the easier it is for her to ask specific questions about those risks, so as to impose on her doctor a duty to provide information; but it is those who lack such knowledge, and who are in consequence unable to pose such questions and instead express their anxiety in more general terms, who are in the greatest need of information. Ironically, the ignorance which such patients seek to have dispelled disqualifies them from obtaining the information they desire.⁹³⁶

Clearly, these remarks are based on the fact that the patient is in need of assistance to realise their values in the medical contexts – obtaining ‘the in-

933 Arvind and McMahon, ‘Responsiveness and the Role of Rights in Medical Law: Lessons from *Montgomery*’ (2020) 28(3) *Medical Law Review* p. 445, 476 (my emphasis). Despite the reference to a treatment interfering with the patient’s bodily integrity, it is clear that the concern for autonomy takes a wider form in the context of negligence.

934 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [75].

935 *ibid* [75].

936 *ibid* [58].

formation they desire'. The statement also implicitly delineates between the function that different classes of information play in a patient's autonomous decision making. Some information is necessary for the patient to concretise, and therefore to decide and act upon, their preferences – without this they cannot even ask the necessary questions. More information may then be provided to improve the position of patients who already have some knowledge, i.e. enough to ask pertinent questions about their care.⁹³⁷ This suggests that autonomy requires a tapered form of information disclosure, as discussed in Chapter 3.

Simultaneously, the Supreme Court viewed the patient, not as someone who was uncritically reliant on professional guidance, but as an agent with the capabilities to participate meaningfully in the clinical process. Most stringently, the majority was clear that a view of patients as 'medically uninformed and incapable of understanding medical matters' could not be 'the default assumption on which the law is to be based'.⁹³⁸ The laws on the labelling of pharmaceutical products and on the provision of information sheets assumed as much.⁹³⁹ More widely, societal developments had also done their part to solidify the fact that patients were not 'wholly dependent upon a flow of information from doctors'.⁹⁴⁰ Patients obtained information about 'symptoms, investigations, treatment options, risks and side-effects' *via* media like the internet (where reliable sources could be distinguished from unreliable ones), they formed support groups and had recourse to materials provided by healthcare providers.⁹⁴¹

Put briefly, patients were not only competent to make medical decisions in the barest, minimal sense, but were to be understood as actors that engaged meaningfully in the decision-making process. They gathered evidence, evaluated the reliability of sources, applied them to their own circumstances, etc. This is arguably the dimension that was captured under our theoretical approach by the reference to rationality. Moving from the abstract to the more specific, one therefore also sees that a procedural

937 See in this regard: *Ollosson v Lee* [2019] EWHC 784 (QB), [2019] 3 WLUK 562 [156].

938 It is also noticeable that the Supreme Court distinguished this capability from the much more demanding characterisation posited by Lord Diplock in *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 894-895, which classed only certain patients as 'highly educated men of experience': *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [76].

939 *ibid* [76].

940 *ibid* [76].

941 *ibid* [76].

autonomy principle has a significant role to play in the law's understanding of a medical professional's disclosure obligations. Especially as it will be seen in the next sections that UK law still leaves many concrete areas of the standard undefined.

ii. The operationalisation of reasonable disclosure

The evaluated rhetoric of *Montgomery* suggests a strong affinity with the procedural autonomy concept and it therefore appears that the developed standard could offer a response to the ML challenges, which were conceived under this umbrella. Indeed, certain commentators assume that the disclosure required by the law is now directly related to that principle.⁹⁴² Hereby, a professional advising a patient would be liable, not in virtue of their failure to meet a duty of care in disclosure, but in virtue of their failure to achieve the outcome of protecting the patient's autonomy.⁹⁴³

Some of the rhetoric used in *Montgomery* offers support to this argument. For instance, in defending the new duty, the judgment of Lord Kerr and Lord Reed framed it as, fundamentally, relating to what 'the respect for the dignity of patients requires'.⁹⁴⁴ Likewise Lady Hale cited Jonathan Herring, who asked: 'whether there was enough information given so that the doctor was not acting negligently *and* giving due protection to the patient's right of autonomy'.⁹⁴⁵

Under this interpretation, reasonable disclosure would be equated with autonomy-compliant disclosure. An operationalisation of this standard would simply require the courts to elaborate upon patients' decision-making needs; the professional would be obligated to meet these.⁹⁴⁶ For example, as *Montgomery* itself stated, this requires a degree of comprehensib-

942 An overview can be found in: Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 380.

943 *ibid* 378-379.

944 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [93].

945 *ibid* [108] (my emphasis), citing Herring, *Medical Law and Ethics* (Fourth Edition 2012) 170.

946 As a result, 'the Supreme Court arguably abandoned negligence analysis altogether, in that for liability to arise it may no longer need to be shown that the defendant acted unreasonably in all the circumstances of the case': Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 378.

ility that is not achieved ‘by bombarding the patient with technical information which she cannot reasonably be expected to grasp’.⁹⁴⁷

However, when the Supreme Court directly addressed the doctrine underlying the newly fashioned disclosure obligation, it clearly understood itself to be operating within a ‘traditional framework of negligence’.⁹⁴⁸ Holding that ‘the doctor’s duty of care takes its precise content from the needs, concerns and circumstances of the individual patient, *to the extent that they are or ought to be known* to the doctor’.⁹⁴⁹ In operation therefore, a connection to the reasonable beliefs and actions of the professional is maintained. An untempered protection of the autonomy principle is not foreseen.

In particular, regarding the types of disclosure with which the *Montgomery* case was concerned – disclosure of risks and alternatives – more concrete indicators were provided. Namely, the legal standard required one to ask whether the professional acted reasonably in identifying and providing material information. A patient had to be informed of the risks and benefits that the professional ‘anticipated’ and of the alternatives that were deemed ‘reasonable’.⁹⁵⁰ Regarding the resultant communication of this information, Lord Kerr and Lord Reed also stated that they were imposing ‘a duty on the part of doctors to *take reasonable care to ensure that a patient is aware* of material risks of injury that are inherent in treatment’⁹⁵¹ and that it was the ‘aim’ of the doctor’s advisory role to engage in a dialogue to impart such information.⁹⁵²

This approach, defining categories of information as appropriate subjects of disclosure, was a further, less noticeable manner in which the UK’s highest court sought to reconcile their unorthodox focus on the patient’s need with negligence’s narrower analysis of the reasonableness of the defendants’ actions. By implication, it was only certain types of information that a doctor could realistically expect the reasonable patient to require. As was noted above, this may be interpreted to coincide with the finding that a patient’s autonomy is not served by bombarding them with information. Unfortunately, it was left relatively open what such categories of disclosure

947 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [90].

948 *ibid* [82].

949 *ibid* [73].

950 *ibid* [90].

951 *ibid* [82] (my emphasis).

952 *ibid* [90].

are, beyond the firmly established classes of anticipated, i.e. known, risks and reasonable alternatives.

Below we will take a casuistic approach, one focussed on these and other categories that the established case law has considered. However, also recognising the openness of the standard laid down by the Supreme Court, we will appreciate the demands of the autonomy principle. This arguably has the potential to reshape certain dimensions of the medical professional's obligation to act with reasonable care, specifically in the AI context.

What ought to be mentioned before moving on, and which was clear even before *Montgomery*, is that negligence imposes broader obligations than battery to convey information to the patient. A doctor's duty is not limited to interferences with the claimant's bodily integrity. They may be negligent for the information they provide on the drugs that they prescribe⁹⁵³ and for the information they give regarding decisions not to treat at all.⁹⁵⁴ Nor is the assessment of the professional's reasonableness limited to disclosure that precedes a specific intervention, but encompasses the disclosure that they ought to make at later stages as a part of the ongoing relationship with the patient.⁹⁵⁵

iii. Summation

The 'material information' standard of disclosure laid down in the seminal case of *Montgomery* appeals to the figure of the reasonable patient, but also to the reasonable professional and to the information that such a professional ought to know a specific patient would attach significance to. In laying down this standard, the court left open many questions, but

953 In *Kennedy v Frankel* a breach of duty was found for a failure to warn about certain side effects of an oral dopamine agonist: *Kennedy v Frankel* [2019] EWHC 106 (QB), [2019] 1 WLUK 216 [50]. See similarly: *Blyth v Bloomsbury Health Authority* [1987] 2 WLUK 77 (although no breach of duty was found).

954 *Webster v Burton Hospitals NHS Foundation Trust* [2017] EWCA Civ 62, (2017) 154 BMLR 129. In *Mordel v Royal Berkshire NHS Foundation Trust* full information had to be given about a treatment that was positively declined: *Mordel v Royal Berkshire NHS Foundation Trust* [2019] EWHC 2591 (QB), (2020) 172 BMLR 106.

955 *Spencer v Hillingdon Hospital NHS Trust* [2015] EWHC 1058 (QB), [2015] 4 WLUK 354; *Gallardo v Imperial College Healthcare NHS Trust* [2017] EWHC 3147 (QB), [2017] 12 WLUK 198.

evidently gave great weight to the kinds of factors that were argued to shape the principle of procedural autonomy.

In the following, the challenges associated with clinical ML devices will be analysed from the perspective of established categories of information disclosure found in English law. Where the law is open-textured or uncertain, the principle of autonomy and the materiality test will be appealed to.

2. The risks of medical AI

Montgomery affirmed the general position that a patient must be advised of the risks associated with a procedure.⁹⁵⁶ Its novelty arose from the standard it set in determining which risks are disclosable: material risks that a doctor knows or ought to know a reasonable patient, or this particular patient, would attach significance to. In this section the goal is to determine whether AI's autonomy challenges pose distinct risks and/or whether they give rise to interrelated factors that must be disclosed.

i. Specific risks

If one considers whether an AI challenge constitutes a specific disclosable risk, analogous to the risk of shoulder dystocia in *Montgomery*, then one should begin with the judiciary's usage of the concept. The courts have not explicitly defined risk for the informed consent context, but they have assumed the relevance of two functions: 'the degree of likelihood of [the event] occurring and the seriousness of the possible injury if it should occur'.⁹⁵⁷ The judgments also indicate that medical evidence is relevant to the definition of a risk, going towards both of these dimensions.⁹⁵⁸

A concrete example of the approach can be given by reference to the already discussed decision of *Chester v Afshar*. Here there was a recognised 1% to 2% risk of neurological damage during an operation on the claimant's

956 This had already been clear, even under less generous interpretations of the *Bolam* standard: *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 894-895.

957 *ibid* 888.

958 *Montgomery* cited Lord Scarman's approach, *ibid* 888, with apparent approval: *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [48], [83].

spine.⁹⁵⁹ Although this was relatively small, the consequences of its occurrence were correspondingly serious: including the risk of paralysis.⁹⁶⁰ The duty to disclose this specific risk, with a small probability but grave significance, was not in dispute.⁹⁶¹

In relation to AI this means that, where the professional knows that the technology alters the severity or probability of a specific physical harm befalling the patient, a straightforward case can be made for disclosure. And some of the AI risks, outlined in Chapters 2 and 3, will satisfy this test. For instance, it was seen that the technology will assist in providing critical care in certain settings, such as for the identification of brain haemorrhages. In addition, even if a device exhibits good initial accuracy (going towards the likelihood of an adverse event occurring), we saw that such assessments are likely to require adjustments for real-life clinical settings. Consequently, where AI is related to a sufficiently likely and sufficiently serious event, it poses a type of risk that the doctor ought to know would be significant to the reasonable patient, or sometimes, to the specific, individual patient.

However, some distinct problems still emerge for cases involving AI-generated risks of this kind, given that negligence requires the defendant to meet a certain standard in their behaviour. *Inter alia* a defendant professional must have been aware of the absent information, or ought to have been aware of it, in order to be liable for its non-disclosure: 'A doctor cannot be held negligent for failing to make a disclosure of matters of which he had no knowledge, unless that knowledge can be imputed to him (ie, unless it can be concluded that these were matters he ought to have known about)'.⁹⁶² As was discussed in Chapter 2, the nature of AI's specific risks is less likely to be known or quantifiable in this way. Particularly as techniques for scientific validation are more difficult to apply and evaluate.

Another issue arises in relation to the concern, expressed in *Montgomery*, that the patient can be overwhelmed by disclosure of technical information and, specifically, that risk disclosure should not be broken down to percentages.⁹⁶³ It is possible to provide the patients with too much information about a procedure and its risks. This is illustrated by *Ollosson v Lee*. It was

959 *Chester v Afshar* [2004] UKHL 41, [2005] 1 AC 134 [11].

960 *ibid* [39].

961 *ibid* [55].

962 *AH v Greater Glasgow Health Board* [2018] CSOH 57, (2019) 169 BMLR 120 [55]; *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [42].

963 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [89]-[90].

sufficient that it was communicated to the patient that there was ‘a risk of long term persisting pain which could range from mild to severe’ and that, in terms of magnitude, this risk was ‘small, adding that it was greater than the rare/remote risks of early or late failure’.⁹⁶⁴ A more detailed accounting of the risk’s size and effects did not have to be provided.

In relation to the clinical deployment of ML technologies, these findings suggest that there will be specific risks raised by AI that are material to patient autonomy, but do not need to be disclosed separately. More generally it suggests that professionals need to be mindful of not overburdening patients with technical information on AI and the way in which its functioning shapes these risks. This analysis buttresses the observations made in Chapter 3 that the disclosure of specific risks is not the correct vehicle for conveying the unique challenges posed by AI to patients.

ii. Risk-relevant status

Beyond AI’s alteration or generation of specific risks, it was also claimed that the technology constitutes a risk-relevant characteristic of an intervention. This characteristic was likened to that involved in the deployment of innovative procedures or of unlicensed or off-label uses of devices. When combined with the autonomy principle, this analogy provides an argument for the disclosure of the technology’s use and status.

In establishing the relevance of such a risk-related status under UK law, it is notable that, in defining the disclosure standard, *Montgomery* deemed a multiplicity of features relevant to the materiality of a risk:

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.⁹⁶⁵

This marks the adoption of an open-ended approach by the Supreme Court. Certainly, it indicates that a patient must be advised of a broad

964 *Ollosson v Lee* [2019] EWHC 784 (QB), [2019] 3 WLUK 562 [151]-[152].

965 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [89].

range of factors that can reasonably be anticipated to contribute to their assessment of a risk.

One such factor, which has been given considerable attention by legal commentators, is the innovative nature of a procedure. This is not merely an aspect of the seriousness or magnitude of risks that the patient is facing in their case. Instead, this classification conveys generic information to the patient. For example, it indicates the lack of scientific testing of a procedure, whereby one is not proceeding according to the ordinary standards of validated scientific knowledge.⁹⁶⁶ Relatedly, there is an assumption that there is greater degree of uncertainty in such situations about the nature and degree of physical harms to which the patient is exposed.⁹⁶⁷

Mills v Oxford University Hospitals NHS Trust provides a rare judicial examination of the issue. The case applied *Montgomery's* materiality test to demand disclosure of the fact that an operation involved the use of 'a novel technique, still in its evolution and not well established'.⁹⁶⁸ In the course of this, the court emphasised particularly the uncertainty involved in such innovative procedures. Stating: 'if [the surgeon] had referred to the series [of novel operations] he had undertaken up to that point, it would have been important to emphasise that it was uncertain what the complication rate for a longer series would be and, as with any new technique, the risks and benefits were not yet clear'.⁹⁶⁹ The absence of adequate validation of the professional's performed procedure was also brought up in the claimant's successful argument: the surgery was not 'a standard well tested technique for resection of a brain tumour'.⁹⁷⁰ It appears that *Mills* accepted a procedure's wider risk-related status as a piece of information

966 '[I]n offering innovative treatment, the physician is working on a hunch or scientific theory that has not been adequately investigated or researched': Chan, 'Legal and Regulatory Responses to Innovative Treatment' (2013) 21(1) *Medical Law Review* p. 92, 94.

967 'Such medical procedures are administered for the benefit of a specific patient but have uncertain outcomes because they have not been adequately tested': *ibid* 94. The 'great uncertainty' of a proposed innovative treatment was also a dominant theme in: *Simms v Simms* [2002] EWHC 2734, [2003] Fam 83.

968 *Mills v Oxford University Hospitals NHS Trust*, [2019] EWHC 936 (QB), (2019) 170 BMLR 100 [197], [201].

969 *ibid* [204].

970 *ibid* [15].

that is material to the reasonable patient or at least material to certain risk-averse individuals.⁹⁷¹

Beyond *Mills*, commentators have also regarded the uncertainty generated by innovative treatments as a significant, reasonable concern. In this respect, Cockburn and Fay have noted that: ‘While by definition it is impossible to quantify unknowns as statistical risks, ensuring the patient realises that s/he is effectively rolling the dice by undergoing innovative treatment is an important aspect of obtaining consent’.⁹⁷² Similarly, O’Neill surmises that ‘Although under UK law, surgeons cannot be expected to warn of unforeseeable unknown risk, unknown harm is foreseeable in the case of innovative medical treatment’ and one can require the disclosure of this factor.⁹⁷³

Given the limited authority on the matter of innovative interventions, it is also worth considering the relevance of another status explored in Chapter 3. Namely, the unlicensed or off-label use of medical products and devices. While there is no exact connection to an intervention’s risk-profile, these categorisations were also understood as something significant for the patient’s risk-related decision-making.

Jones v Taunton and Somerset NHS Foundation Trust supports this analysis in the British context. This case concerned the novel, unlicensed drug (Nifedipine), which had not been subject to ‘convincing double-blind studies’ for the relevant use, and its risks and benefits were not fully determined.⁹⁷⁴ Nevertheless, it was becoming the drug of choice over a licensed drug that was associated with marked side-effects. The question before the court was whether the prescription of this drug, in these circumstances was negligent.⁹⁷⁵

971 *ibid* [215]. Contrast the prior situation under the *Bolam* standard, which was incorrectly maintained in *Grimstone v Epsom and St Helier University Hospitals NHS Trust* [2015] EWHC 3756 (QB), [2015] 12 WLUK 749. On the problematic nature of the case, see especially Austin’s commentary: Austin, ‘Grimstone v Epsom and St Helier University Hospitals NHS Trust: (It’s Not) Hip to Be Square’ (2018) 26(4) Medical Law Review p. 665.

972 Cockburn and Fay, ‘Consent to Innovative Treatment’ (2019) 11(1) Law, Innovation and Technology p. 34, 46.

973 O’Neill, ‘Lessons From the Vaginal Mesh Scandal: Enhancing the Patient-Centric Approach to Informed Consent for Medical Device Implantation’ (2021) 37(1) International Journal of Technology Assessment in Health Care e53, 1-5, 4.

974 *Jones v Taunton and Somerset NHS Foundation Trust* [2019] EWHC 1408 (QB), [2019] 6 WLUK 193 [135].

975 *ibid* [135].

Consequently, *Jones* is not strictly speaking concerned with informed consent obligations. Still, it supports the approach taken to the classification of risk-related statuses in *Mills*. For one, the judge endorsed the relevance of the unlicensed nature of a drug to the clinical assessment of risks. He acknowledged that a licensed drug should not be abandoned until the risks and benefits of an unlicensed one have been established.⁹⁷⁶ But he was also clear that this relationship was not determinative of relevant assessments: ‘the fact that Nifedipine was unlicensed is merely one factor, and I find, not a strong factor, in the light of the other evidence’.⁹⁷⁷ In addition, it was mentioned that the patient’s informed consent must be obtained with respect to the unlicensed nature of the drug and that this was a prerequisite for the finding that the drug’s use was non-negligent.⁹⁷⁸

Even if the innovative or off-label nature is not determinative of a risk/benefit balance, the case law therefore suggests that it provides the patient with important information that they should be provided with. Such a view gains added support from the General Medical Council’s (non-legally binding) guidance. This specifically indicates that the doctor should share information with the patient on ‘whether an option is an innovative treatment designed specifically for their benefit’.⁹⁷⁹ Here the principles of informed consent do not require a detailed explanation of the (lack of) scientific evidence or approval, or the mechanisms associated with this, but they do require the disclosure of the procedure’s type and the general way in which this is likely to impact the risk calculus.

Such a disclosure obligation may further be modelled on the finding in *Webster v Burton Hospitals NHS Foundation Trust*, in which a specific risk was identified. Namely, ‘the increased risk of perinatal (the period around birth) mortality, including *ante partum* (before delivery) mortality’ from a rare combination of conditions.⁹⁸⁰ However, this risk had to be inferred from a ‘small (or extremely small) statistical base’.⁹⁸¹ While the court was therefore not strictly speaking dealing with an innovative procedure, this

976 *ibid* [139]–[140].

977 *ibid* [135].

978 The defence expert accepted that prescription was non-negligent on the supposition that the patient was told ‘that it is an unlicensed drug’ and was given a full explanation: *ibid* [131].

979 General Medical Council, ‘Decision Making and Consent’ (London 2020) para 13.

980 *Webster v Burton Hospitals NHS Foundation Trust* [2017] EWCA Civ 62, (2017) 154 BMLR 129 [38].

981 *ibid* [39].

case also exemplifies judicial dealings with a form of uncertainty that exists at a high level of generality. Crucially for the present argument, the suitable informed consent standard was stated at a corresponding level of generality: ‘that there was “an emerging but recent and incomplete material showing increased risks of delaying labour in cases with this combination of features’.⁹⁸² *Webster* arguably strikes a reasonable balance between the patient being provided with sufficient information to make a practical decision and not providing them with too much information, thereby overwhelming their decision-making capabilities. This balance is crucial for the preservation of procedural autonomy.

It is argued that, on the basis of the factors discussed in Chapters 2 and 3, an argument can be constructed for UK law requiring AI characteristics to be disclosed at a comparable level of abstraction. In particular, the patient should be told that there may be limitations in the way in which an ML output can be connected to established scientific knowledge and that, relatedly, the use of AI involves substantial uncertainties, similar to those associated with novel or unlicensed procedures. The provision of such information would avoid the aforementioned issues with the disclosure of specific risks. While health professionals may not know the precise nature of the dangers inherent in AI deployment, they are (or certainly should be) aware that an ML device is being used in the patient’s care and that certain structural uncertainties accompany this.

In the final analysis, while some statements in existing case law, in combination with our normative autonomy-based analysis, suggest that such an obligation can plausibly be formulated under UK common law, its strength must necessarily remain uncertain given the limited nature of precedent.

3. Alterations of expertise

Under English law, the analysis of obligations to disclose (the lack of) professional expertise appears to have been couched almost exclusively in terms of battery. The limitations of such an approach were assessed above. Particularly regarding AI, it is not straightforward to subsume the technological supplementation of human capabilities under the existing case law.

982 *ibid* [40].

This raises the question of the role that negligence can play in this realm. Theoretically, the greater emphasis placed by this mechanism on *informed* consent, rather than merely valid consent, should grant patients a more comprehensive protection. The finding in *Jones v Royal Devon and Exeter NHS Foundation Trust* grants some support to this position. In this case, there was a discrepancy between the surgeon that the patient believed would operate on her spine and the one who in fact did. The patient had been admitted under the care of a particular surgeon and was given to understand that her operation was scheduled to be performed by him.⁹⁸³ She was only told of the change in her surgeon's identity as she was going into theatre, at which point she could not effectively provide her consent.⁹⁸⁴ Under these circumstances the court found that there was a breach of the defendant's informed consent obligations.⁹⁸⁵

At the very least, *Jones* recognises that a change in the specific identity of a medical professional can defeat the claimant's consent. This was in spite of the patient being told that it could not be guaranteed that her procedure would be performed by one particular person, but only that it would be a team member with appropriate experience.⁹⁸⁶ Moreover, the claimant had withdrawn the argument that the surgeon who performed the procedure was not adequately qualified.⁹⁸⁷ As such, the court arguably found that, regardless of comparable expertise, a breach of informed consent obligations arises where the personal identity of a known professional changes. This would be similar to the position assumed by the earlier battery cases. As AI would not generally bring about such changes, being deployed by an identifiable human individual, *Jones* would not demand the disclosure of changes in the expertise of that professional.

Other comments in the case, however, suggest that it was a shift in expertise that was the operative factor about which the claimant should have been informed. It was emphasised that the desired professional 'enjoyed a very high reputation both locally and nationally' and that he was

983 *Jones v Royal Devon and Exeter NHS Foundation Trust* [2015] Lexis Citation 3571 [28], [31].

984 *ibid* [37]. We will return to the possibility of giving informed consent under certain pressures below. What is crucial here, is that there was no effective disclosure to the patient of the surgeon's identity and/or experience.

985 *ibid* [37].

986 *ibid* [23].

987 *ibid* [12].

a surgeon of considerable seniority.⁹⁸⁸ Furthermore, the patient's GP had made the medically informed recommendation that it would be preferable for her to wait (in spite of considerable pain) to be operated on by this surgeon.⁹⁸⁹ Lastly, the patient made her own assessment of the surgeon's reliability, based upon the operations she knew he had performed on her acquaintances.⁹⁹⁰ In light of these statements, it appears that the court found a breach of duty because the patient had reason to believe that one level of professional expertise was being brought to bear on her care (that of a senior, very experienced surgeon) when in fact a professional with a lower, albeit adequate, capability was involved.

This line of argument would be extendable to AI. In some circumstances a patient may have grounds for believing that a human professional involved in their care is an expert with a high degree of skill, as they are performing a demanding procedure. This impression may be misleading if an ML device is in fact supplying a substantial amount of the relevant capabilities. Even though these capabilities may be adequate, Chapter 3 illustrates their different nature.

Given the ambiguity in *Jones* regarding the significance of shifts in expertise *per se*, rather than changes in human identity, two further arguments can be advanced. First, that the more recent battery analyses, holding that individuals may attach significant weight to the status and expertise of their carers, and limited *dicta* in other negligence actions, reinforce the recognition of related forms of actionable non-disclosure in negligence. Second, *Montgomery's* test of material information and the weight of the autonomy principle indicate the correctness of the same conclusion.

Regarding the findings of the battery analysis, one can recall the case of *R v Melin*. Here the court strongly affirmed the significance that a patient would attach to the experience-related status of their professional. In particular, the example was cited of a new doctor taking over a general practice and it was implied that their status, not their identity, would be the most important piece of information for the patient to know.⁹⁹¹ This evinces a novel receptiveness to the patient's concerns about the professional's expertise, one which had previously been obscured by a focus on quite drastic changes of identity.

988 *ibid* [6], [65].

989 *ibid* [29].

990 *ibid* [64].

991 *R v Melin* [2019] EWCA Crim 557, [2019] QB 1063, citing Ormerod and Laird, *Smith, Hogan, & Ormerod's Criminal Law* (Fifteenth Edition 2018) 672.

To supplement this finding, there is also limited *dicta* in the lower courts that reasonable patients may attach significance to the automated, rather than human, nature of an important decision. Specifically, in considering the patient's refusal to enter a clinical trial that involved a randomised computer-based selection of participants, the judge in *C v Colchester Hospital University NHS Foundation Trust* noted: 'Mr C, understandably, refused to enter into that study (...) he wanted a clinician to dictate and determine when he was to be operated upon and not to be dealt with, as he saw it, by computer'.⁹⁹² In effect this amounts to a recognition that the full automation of an important clinical decision constitutes a legitimate concern for a patient.

Second, one can consider the role of the informed consent standard delineated in *Montgomery*, which had not yet explicitly been applied in *Jones*. As was noted, this has been interpreted to impose a broad duty to disclose material information to the patient. That is, information which would be significant to the reasonable patient, or which the doctor knows or ought to know is significant to the particular individual. The outlined findings suggest a willingness on behalf of the courts to recognise that shifts in expertise, and to a lesser extent a shift between human and computerised clinical decision making, are significant factor for a patient. Consequently, it can be argued to be a natural step, under *Montgomery's* only relatively recently redefined approach, to deem certain shifts in expertise to be material information.

In the case of ML devices, this would mean disclosing certain instances where the technology is covertly substituting human capabilities, which are expected to be of a high standard, to a meaningful degree. For example, both the battery case law and the *dicta* cited above are referring to substantial, meaningful replacements of expertise; what is significant to the patient is a change in 'status' or the comprehensive automated randomisation of a decision. Similarly, the distinctions drawn in Chapters 2 and 3 illustrate that a full automation of a process is much more problematic from the perspective of patient autonomy than the enhancement of a human decision.

Relevant differences also emerge when one considers whether the cumulative amount of human expertise, which is brought to bear on a decision, is decreased or whether the AI is merely providing a 'second opinion'. To the extent that *Montgomery* mirrors such autonomy concerns, the courts should apply the fact-sensitive materiality standard to determine exactly

992 *C v Colchester Hospital University NHS Foundation Trust* [2018] 2 WLUK 850 [17].

those kinds of substitutions that would be deemed significant by patients. Focussing on instances where there is a partial replacement of pre-existing human capabilities, as outlined in Chapter 2, could serve as a useful starting point for this purpose.

In sum, *Jones* indicates that, where a patient has an expectation that a high degree of expertise will be brought to bear on their care, not disclosing shifts in this expertise may constitute a breach of duty. Drawing on an analysis of other case law and the *Montgomery* materiality standard, we have added the caveat that relevant shifts must be sufficiently substantial. Where AI significantly reduces the overall level of expertise brought to bear on a procedure – substituting itself for the decision making of an experienced human – a breach may consequently occur. It is more questionable whether lesser shifts, or indeed simply a shift from human to computerised capabilities, must also be disclosed to the patient.

4. Information concerning the choice of goals

In the context of AI, the most striking effect on the patient's autonomy is perhaps not its risk-profile or its alteration of the balance of human-machine expertise, but the novel way in which it sets and pursues its own objectives. This generates two related issues that were outlined in Chapter 3 and touched upon under the battery analysis above.

First, certain AI can partially determine an aspect of the patient's care and therefore pose a particularly significant challenge to a patient's ability to select the objectives that they wish to realise. Second, ML devices can exhibit a propensity to direct human decision making through nudging. These were problematic to a lesser extent – subjecting patient decision-making to systematic external influences. It is argued here that legal categories dealing with alternative procedures and with external influences on patient choice can partially respond to these interferences with patient autonomy.

i. Understanding choices

One class of disclosure cases seeks to identify those procedures that a patient must be informed about in complex healthcare interactions to

make appropriate choices in their care. As outlined, *Montgomery* expressly requires that the patient be made aware of ‘any reasonable alternative or variant treatments’.⁹⁹³ The reference to treatment should not be read restrictively. Even prior to the development of the patient-centred standard of care, the courts had required the disclosure of alternative modes of diagnosis.⁹⁹⁴ Accordingly, the term procedure will be used as a catch-all term for both treatment and diagnostic scenarios.

One ought to begin by establishing the kinds of choices that the law has protected by requiring the disclosure of alternative options. One type of case is relatively straightforward in this regard: where one readily available alternative represents a distinct risk-benefit balance to the professional’s chosen intervention. This normally constitutes the most important factor that the courts have considered in determining whether an alternative represents a choice that must be discussed with the patient.⁹⁹⁵ *Montgomery* itself bears testament to this fact. In determining the necessity of disclosure, Lord Kerr and Lord Reed considered that there was a stark difference between the substantial risk of shoulder dystocia occurring during vaginal delivery and the miniscule risks of harm resulting from a caesarean section.⁹⁹⁶

Birch v University College London Hospital NHS Foundation Trust relied on a similar analysis. Here the doctors carried out their investigation of the patient’s symptoms with an invasive cerebral catheter angiogram, when the performance of an MRI scan constituted another non-invasive option.⁹⁹⁷ A choice was made in terms of sequencing. Both procedures were modes of assessing the patient for a similar range of conditions and the angiogram was more accurate but also riskier in some respects.⁹⁹⁸ The professionals opted for the former *instead of* the latter and this had implications for the specific risks that the patient was exposed to.

993 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [87].

994 *Birch v University College London Hospital NHS Foundation Trust* [2008] EWHC 2237 (QB), (2008) 104 BMLR 168.

995 ‘[J]udgements made about information about risk that should be disclosed to a patient will often have a direct impact on the range of alternative or variant treatments that a patient should be offered accordingly’: Dunn and others, ‘Between the Reasonable and the Particular’ (2019) 27(2) *Health Care Analysis: Journal of Health Philosophy and Policy* p. 110, 112.

996 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [94].

997 *Birch v University College London Hospital NHS Foundation Trust* [2008] EWHC 2237 (QB), (2008) 104 BMLR 168 [39].

998 *ibid* [50]-[52].

Such a straightforward analysis can also be applied to choices between procedures with risk-relevant characteristics. Consider the aforementioned judgment in *Mills v Oxford University*. The case concerned the use of a novel surgical technique over a pre-existing standard procedure. Neither the experts nor the court questioned the claim that the standard procedure was a relevant, alternative choice.⁹⁹⁹ This was accepted as a matter of course, in spite of certain clinical factors that favoured the new procedure, including: decreased short-term mortality and better cosmetic results.¹⁰⁰⁰ Apparently the uncertainty surrounding the comparative risks and benefits of the new technique over the older one sufficed to frame the professional's choice as being between material alternatives.

This solidifies the argument that the patient must be given information on the AI's risk-related characteristics and be offered a choice regarding the technology's use. It is one alternative in this respect. However, this approach would only grant the patient information and control over the ML device's objectives in so far as they are related to risks of physical harm.¹⁰⁰¹ The necessity of maintaining this connection was made clear

999 *Mills v Oxford University Hospitals NHS Trust*, [2019] EWHC 936 (QB), (2019) 170 BMLR 100 [195]-[197].

1000 *ibid* [215].

1001 This depends somewhat on how closely one associates the anticipated benefits of a procedure with the risks of action or inaction. For example, Veatch has noted that 'almost any medical procedure will involve a mixture of potential benefits and potential harms. Almost any procedure or therapeutic agent has potential side effects. Even determining that an effect is a "side" effect rather than a "benefit" involves value judgments that are complex': Veatch, 'Doctor Does Not Know Best' (2000) 25(6) *The Journal of Medicine and Philosophy* p. 701, 706. See similarly Feng, who holds that the nature-risk 'distinction is also fallacious as it assumes that there is an inherent difference in terminology and substance between nature of treatment and risks inherent in treatment (...) If there is a duty in the tort of negligence on the part of a doctor to inform his patient of risks inherent in treatment there must surely be the same duty on the doctor to disclose to his patient the patient's ailment, treatment (including nature of treatment), benefits of treatment, and other matters connected with medical advice': Feng, 'Failure of Medical Advice' (1987) 7(2) *Legal Studies* p. 149, 156-158. The courts' approach to the disclosure of benefits remains ambivalent however. In *Montgomery*, the benefits and the seriousness of the condition were considered alongside risks – apparently going towards the nature of the risk, rather than constituting a consideration in their own right: *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [89]-[90]. Lord Scarman in *Sidaway* seems to treat them as separate but complementary: 'The doctor's duty can be seen, therefore, to be one which requires him not only to advise as to medical treatment but also to provide his patient with the information needed to enable the patient to consider

in *Plant v El-Amir*. Here the limited likelihood of achieving the patient's highly valued outcome (to be able to read again) through an eye operation was only deemed disclosable to the extent that it impacted her readiness to run the risks of the procedure.¹⁰⁰²

To contend that the disclosure of AI's ability to pre-determine choices is independently necessary, one can appeal to the argument of Section II.A.: that negligence does not strictly require the eventuation of physical injury and recognises harms to autonomy. But, under the breach element, one must bolster this argument by additionally examining the rationale underlying the disclosure of alternatives.

This rationale can arguably be found in the disconnect that an advising professional may cause between a patient's decision making and the realisation of their goals, echoing the autonomy challenge conceptualised in Chapter 3. In the UK legal context, this was put most forcefully by Lady Hale in *Montgomery*, when she concretised the majority's observation that a choice between alternatives will often depend upon non-clinical considerations.¹⁰⁰³ She observed that the doctor's decision to withhold information interfered with the patient's evaluative judgment – itself evincing the view that vaginal delivery was 'in some way morally preferable to a caesarean section'.¹⁰⁰⁴ By making this judgment the doctor was, precisely, 'depriving the pregnant woman of the information needed for her to make a free choice in the matter'.¹⁰⁰⁵ Turton has similarly emphasised this in her analysis of *Montgomery*, stating that: 'The duty is not limited to enabling the patient to decide which risks she is willing to accept, but is also focused on protecting her right to make that decision in line with her own values'.¹⁰⁰⁶

and balance the medical advantages and risks alongside other relevant matters': *Sidaway v Board of Governors of the Bethlem Royal Hospital* [1985] AC 871, 886. Similarly, Lord Hope stated in Chester that 'whether the risk is worth running for the benefits that may come if the operation is carried out': *Chester v Afshar* [2004] UKHL 41, [2005] 1 AC 134 [86]. Given the significance of certain choices for autonomy, independently of involved risks, this section considers this aspect in isolation.

1002 *Plant v El-Amir* [2020] EWHC 2902 (QB), [2020] All ER [82]-[87].

1003 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [82], [109].

1004 *ibid* [114].

1005 *ibid* [114].

1006 Turton, 'Informed Consent to Medical Treatment Post-Montgomery' (2019) 27(1) *Medical Law Review* p. 108, 129.

A later first instance decision, *C v Colchester Hospital University NHS Foundation Trust*, also bears out this analysis of the considerations underlying the disclosure of alternative options. In this case the claimant had been subjected to a course of treatment that involved chemoradiotherapy (CRT) in order to shrink his tumour, before being subjected to a life-changing surgery to remove it completely. The patient argued that he had not been informed that a proportion of people were cancer-free after undergoing only CRT and, had he been so informed, he would have refused surgery.¹⁰⁰⁷ Walden-Smith DJ found that CRT in combination with no further treatment was not an alternative to surgery, at least not for a patient who had made it clear that they wanted to be cured.¹⁰⁰⁸ Given the patient's desire and the state of medical knowledge, it was fallacious to argue that CRT constituted a separate type of treatment. The information the patient lacked in no way disconnected them from the choice to pursue their admitted objective.

In consequence, one can see that the information about alternatives is envisaged to place patients in a position where they can make a decision that realises their personal values. Only certain deficiencies will altogether prevent them from occupying this position.

In the AI context, Chapter 3 identified a comparable deficiency in one particular subset of cases: where the technology had a relatively wide discretion to select non-personalised goals and partially determined a clinical decision for the patient. This pre-empts them from making a choice by reference to their own considerations, even if a professional is not the one deploying an evaluative standard or foreclosing the patient's options.

Two ML scenarios were taken to exemplify sufficiently serious interferences of this type. Namely, ML uses that pre-determine a patient's choice to undergo a serious diagnostic procedure and AI uses that triage a patient's access to health care.¹⁰⁰⁹ Given the outlined case law, it is arguable that information about these objectives would be considered significant by the reasonable person in the patient's position and by many specific patients.

Under *Montgomery's* standard it is therefore arguable that, when a patient is being subjected to these AI uses, they must be made aware that they are committing themselves to the pursuit of non-personalised objectives

1007 *C v Colchester Hospital University NHS Foundation Trust* [2018] 2 WLUK 850 [36].

1008 *ibid* [47].

1009 Assuming for the sake of argument that one can identify a person with a relevant duty for the triaging uses of AI.

– whether these possess specific or generalised risk characteristics or not. This is material information. At the same time, such a disclosure is then closely aligned with the discussion of alternatives. The aim is to give the patient a choice between courses of action and one must consider the wider limitations that the law has placed on the disclosure of different options.

The jurisprudence encapsulates these limitations under the already mentioned concept of *reasonable* alternatives.¹⁰¹⁰ This qualifier represents a concession to a standard orientated towards the defendant.¹⁰¹¹ It also coincides with the intuition that, if there is no realistic decision for the patient to make, then a failure to advise them causes no possible violation of their practical autonomy. Consequently, even where an ML's non-personalised commitments amount to material information, a duty to disclose will not be breached if the patient did not have a reasonable choice to make on whether to submit to these commitments or not.

The legal standard for the determination of reasonableness is somewhat uncertain in the wake of *Montgomery* because, while introducing the terminology, the case stopped short of specifying its meaning in any detail. The resulting controversy has been discussed in *AH v Greater Glasgow Health Board*, with Lord Boyd distinguishing between two approaches that arguably constitute the extreme positions on a spectrum. At one end of possible interpretations, a reasonable alternative would be one that is determined by the *Bolam* test – i.e. one that 'no ordinarily competent clinician, exercising ordinary skill and care, would have failed to offer' – and, at the other end, it would be one that draws on *Montgomery's* patient-friendly

1010 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [87], [90]; *Bayley v George Eliot Hospital NHS Trust* [2017] EWHC 3398 (QB), [2017] 12 WLUK 670 [56]; *AH v Greater Glasgow Health Board* [2018] CSOH 57, (2019) 169 BMLR 120 [39].

1011 For example, a doctor cannot be expected to disclose (material) information if they did not know and did not reasonably have to be aware of it. This proved fatal in *Bayley* where 'physicians would not have been familiar with the stenting procedure for this condition in 2008' ... 'I am not satisfied that a reasonably competent vascular surgeon would or ought to have known about the availability or potential use of this treatment in the second half of 2008': *Bayley v George Eliot Hospital NHS Trust* [2017] EWHC 3398 (QB), [2017] 12 WLUK 670 [64], [99]. Cf. *Webster*, where it was the doctor's carelessness that caused his lack of knowledge and thus did not exclude the information from consideration, he 'had failed to inform himself about the implications of the rare combination [of conditions]': *Webster v Burton Hospitals NHS Foundation Trust* [2017] EWCA Civ 62, (2017) 154 BMLR 129 [38].

approach – determined by what ‘a patient might find reasonable after a full discussion of all the treatments whether or not these are available’.¹⁰¹²

The former approach would be unusually limiting in the wake of *Montgomery*. For AI the application of this standard would mean that, if at least one body of medical opinion believed that the patient should only be offered the AI-related procedure and this view withstood logical analysis, then this would be the end of the court’s assessment. No disclosure obligation would lie.

However, the courts have not generally embraced such a drastic limitation, which would ‘simply be reinstating [the *Bolam* standard] by the back door’.¹⁰¹³ What emerges from the cases is that clinical suitability, including availability (matters to be determined with the assistance of medical experts) will indeed perform an important gatekeeping role in the determination of which options are reasonable.¹⁰¹⁴ For example, in *Malik v St George’s Hospital Trust* the court accepted that a non-surgical intervention was not a realistic option *inter alia* because of the long-waiting times involved in this method of treating the patient’s acute needs.¹⁰¹⁵

To this extent the approach is moved more towards the *Bolam*-end of the continuum, but this does not constitute the final assessment. In light of the force of the autonomy principle as asserted in *Montgomery*, the courts will not simply defer to one body of medical opinion. Patient-specific factors,

1012 *AH v Greater Glasgow Health Board* [2018] CSOH 57, (2019) 169 BMLR 120 [8].

1013 Sutherland, ‘The Law Finally Reflects Good Professional Practice: *Montgomery v Lanarkshire Health Board*’ [2015](123) *Reparation Bulletin* p. 4, 7-8. See: *Bayley v George Eliot Hospital NHS Trust* [2017] EWHC 3398 (QB), [2017] 12 WLUK 670 [61].

1014 *AH v Greater Glasgow Health Board* [2018] CSOH 57, (2019) 169 BMLR 120 [43]. See also: *McCulloch v Forth Valley Health Board* which applied *Ah Glasgow* and rejected the necessity of disclosing alternatives not indicated in the circumstances of the case: *McCulloch v Forth Valley Health Board* [2021] CSIH 21, [2021] 3 WLUK 569 [40]. Sutherland has convincingly concluded that *AH* ‘accepted that doctors were entitled to filter information given to patients on the basis of what other doctors considered reasonable but concluded that a doctor could not withhold information about a reasonable alternative treatment and the risks associated with that on the basis of their own preference’: Sutherland, ‘*Montgomery: Myths, Misconceptions and Misunderstanding*’ [2019](3) *Journal of Personal Injury Law* p. 157, 166.

1015 *Malik v St George’s University Hospitals NHS Foundation Trust* [2021] EWHC 1913 (QB), (2021) 181 BMLR 135 [45]-[93]. Compare *Bayley*, where the court did consider alternative that were available privately and in Europe, suggesting a somewhat more patient-friendly interpretation of reasonableness: *Bayley v George Eliot Hospital NHS Trust* [2017] EWHC 3398 (QB), [2017] 12 WLUK 670 [49].

such as the patient's definition of their clinical need and the relevance of individual value judgments to the selection of alternatives, will also define which alternatives are reasonable.¹⁰¹⁶

For our two case studies of disclosable AI, based upon their pre-emption of patient choice, this reasonableness limitation is predicted to limit the informed consent obligations for AI's use in triaging. There will be forceful arguments that, in so far as alternative methods of categorising the patient's need are not immediately forthcoming, there will not be a clinically suitable (and therefore reasonable) alternative for the relevant task. Nor will it usually be possible to appeal to the autonomy principle to overcome this limitation. For, any default triaging system – e.g. one broadly aligned with medical urgency or with a first come, first served approach – may be expected to be equally unreceptive to the patient's personal objectives. Consequently, disclosure of this AI use and its, potentially uncertain, objectives will not be necessary, even where it makes certain decisions for the patient.

By comparison, a patient is currently envisaged to have available, suitable alternative options when deciding whether to undergo a diagnostic AI analysis. This includes the option of non-testing or of seeking out a narrower form of analysis. Indeed, even if other testing procedures were not immediately available, the significance of this affirmative choice for the patient's autonomy would still demand its disclosure.¹⁰¹⁷

The only challenge that could emerge for the disclosure of this aspect is arguably a lack of professional awareness: an alternative cannot be reasonable if the professional ought not to have been aware of it. As was mentioned above, unlike a traditional situation involving the non-disclosure of alternatives, ML devices do not involve the professional imposing their preference on the patient. It is the use of the tool that has the consequence of determining a choice. This choice, and the evaluative criteria underlying it, may be obscure to differing degrees, even for the user.

However, this is not anticipated to be a substantial problem for the kinds of disclosure advocated for here. As discussed in Chapter 2, a professional deploying the technology must have a broad awareness of its purposes and – in the case of AI diagnosis – they can reasonably foresee that a patient may not wish to confront the possibility that they are suffering from

1016 *Webster v Burton Hospitals NHS Foundation Trust* [2017] EWCA Civ 62, (2017) 154 BMLR 129 [31]; *Bayley v George Eliot Hospital NHS Trust* [2017] EWHC 3398 (QB), [2017] 12 WLUK 670 [96].

1017 The discussion of this autonomy challenge under the battery mechanism should also be considered relevant to this assessment.

a serious condition. Anticipating the AI's determination of this choice is possible and therefore it must be disclosed.

In sum, the existing UK approach to the disclosure of alternative options suggests that there is a concern to maintain a connection between patients' values and the decisions made in their care. Some instances where AI possess a wide discretion to determine treatment objectives must be disclosed to the patient under this head. One may say that the selection of the AI tool amounts to a pre-determination of their alternatives. Simultaneously, this is not a uniform duty and – in seeing how far it extends – one ought to consider whether the AI denied the patient a true choice that could be anticipated by the professional.

ii. AI's lesser influence on the pursuit of objectives

As the previous section has argued, the foisting of external evaluative judgments upon the patients can constitute an interference with the patient's autonomy that can trigger disclosure obligations under the law of negligence. Yet, up until now, we have considered only a particularly strong interference: where a patient is not given any choice at all whether to pursue certain goals in their care. Another related, lesser violation was seen to occur in relation to AI where it subjected the patient to external influences that were not easily accessible to rational evaluation: nudges.

English law certainly recognises the existence of external psychological pressures and that these can contribute to circumstances where it is not possible for the patient to provide their informed consent. For example, in *Thefaut v Johnston* the judge found that a conversation between surgeon and patient on the day of the surgery was not sufficient for obtaining their informed consent to the procedure because:

this is neither the place nor the occasion for a surgeon for the first time to explain to a patient undergoing elective surgery the relevant risks and benefits. At this point, on the very cusp of the procedure itself, *the surgeon is likely to be under considerable pressure of time* (to see all patients on the list and get to surgery) and *the patient is psychologically committed to going ahead*. There is a *mutual momentum towards surgery*

which is hard to halt. There is no "adequate time and space" for a sensible dialogue to occur and for free choice to be exercised.¹⁰¹⁸

There are arguably two related objections encapsulated in this quote. First, under circumstances of acute time pressure the professional is not able to facilitate the kind of comprehension that is necessary for informed consent. This is not our main concern here. Second, there is a mutually reinforcing momentum between patient and doctor, which may be described as a bias impacting them, that results in an impaired process of decision making. In short, *Thefaut* recognises that the presence of biases can defeat the patient's informed consent. This has subsequently been affirmed.¹⁰¹⁹

In *Thefaut*'s circumstances there was an available way to avoid the presence of this bias of course. It was open to the surgeon to lead the discussions with the patient at an earlier date, or under different circumstances, where time-pressure would not seriously impact either of their decision making. The extent to which the creation of these circumstances, i.e. the exclusion of such pressures, is a matter of degree has since been exemplified in *Ollosson v Lee*. While important information about a risk was only imparted to the patient on the day of a procedure, the court found that 'the information provided at the surgery was in an unpressurised situation, with time to reflect, and against a background where Mr Ollosson had arrived with some knowledge that there was a risk'.¹⁰²⁰ In consequence, the patient did give 'properly informed consent'.¹⁰²¹

These cases stand for the proposition that a doctor is obligated to create circumstances where the patient can meaningfully appreciate the information they are given for a clinical choice. This means *inter alia* that they are not influenced by biases to an unacceptable extent when giving their consent. Given the distinction of *Thefaut* in *Ollosson* it is probable, as well as understandable, that the bar for such biases will not be set too low. Biases are endemic in all decision making and they cannot be avoided. Arguably, the moral position that the doctor allegedly assumed in *Montgomery* could

1018 *Thefaut v Johnston* [2017] EWHC 497 (QB), [2017] 3 WLUK 328 [78] (my emphasis).

1019 *Hassell v Hillingdon Hospitals NHS Foundation Trust* [2018] EWHC 164 (QB), (2018) 162 BMLR 120 [53]. See also the aforementioned case of *Jones v Royal Devon and Exeter NHS Foundation Trust* [2015] Lexis Citation 3571.

1020 *Ollosson v Lee* [2019] EWHC 784 (QB), [2019] 3 WLUK 562 [157].

1021 *ibid* [158].

acceptably have influenced their advice, but they ‘must not put pressure on the patient to accept their advice’.¹⁰²²

The same can be observed in cases dealing with the legal figure of undue influence, according to which the validity of consent may be undermined by an external influence that persuades ‘the patient to depart from her own wishes, to an extent that the law regards it as undue’.¹⁰²³ It must be borne in mind that the pressures exerted to obviate consent altogether must be correspondingly higher than those that may affect a patient’s ability to give an informed consent. Nevertheless, it is instructive that in *U v Centre for Reproductive Medicine* the Court of Appeal accepted the argumentation of the first instance court that undue influence requires ‘something more than pressure (...) it does not matter how strong the persuasion was so long as it did not overbear the independence of the patient’s decision’.¹⁰²⁴ The law is clearly accustomed to distinguishing between acceptable and unacceptable external pressures on patient decision making.

It is hardly arguable that the kinds of AI-induced biases identified in our technical analysis amount to external pressures that undermine the validity of consent, overbearing the independence of a patient’s decision. By contrast, it may be arguable that they can impact the patient’s *informed* consent under the doctrine elaborated in *Thefaut*, although the biases of ML technology require a slightly different analysis. For, as outlined in Chapter 2, it is not clear that the patient’s wider position can be changed to eliminate a relevant bias, akin to the possibility in *Thefaut* of simply making the relevant disclosures earlier. Some degree of AI nudging appears inevitable and acceptable – much like the position regarding a doctor’s advice.

As NHS documentation has acknowledged, the best way to counteract biases in the use of ML technology under such conditions is to educate individuals ‘to recognise their inherent biases, and understand how these affect their use of AI-derived information’.¹⁰²⁵ In other words, where pa-

1022 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [78], citing General Medical Council, ‘Consent: Patients and Doctors Making Decisions Together’ (London 2008) para 5.

1023 *Re T (adult: refusal of treatment)* [1993] Fam 95, 121.

1024 *U v Centre for Reproductive Medicine* [2002] EWCA Civ 565, [2002] Lloyd’s Rep Med 259 [19].

1025 Nix, Onisiforou and Painter, ‘Understanding Healthcare Workers’ Confidence in AI’ (2022) <<https://digital-transformation.hee.nhs.uk/building-a-digital-workforce/dart-ed/horizon-scanning/understanding-healthcare-workers-confidence-in>

tients may be affected by AI biases, *Montgomery's* and *Thefaut's* application of the informed consent standard may translate into an obligation to explain potential AI biases to the patient to facilitate, as far as is reasonably practicable, an unbiased decision.

This still leaves the problem that ML devices would not ordinarily be in a position to influence patients to a sufficiently problematic degree. Nudging a patient into making a certain evaluative choice regarding their care is not a natural analogy to acute time pressure, which impairs their appreciation of a material risk. Yet, a more detailed analysis of problematic pressures is not forthcoming in the case law.

Drawing on the reflective dimension of autonomy it could be argued that even slight, uncertain influences (such as nudging) are sufficiently problematic where it relates to certain non-therapeutic interests of the patient. These would serve not as prompts to clarify decision-making, but as a means to subvert it. However, in the UK it is a complicating factor that these obligations have not arisen under the negligence action. Rather, similarly to the above analysis of expertise-related disclosure, successful claims have been made under the stronger battery action.¹⁰²⁶ The success of such actions under negligence must be seen as an unlikely prospect.

Moreover, in light of the physician's own, potentially limited, understanding of ML biases – a knowledge of which is not necessarily linked to responsible use – it may also be difficult to prove that a reasonable defendant has fallen short of the informed consent standard in not discussing this aspect of AI use. In consequence, although nudging may be a challenge to patients' procedural autonomy, it is not one against which the law of informed consent can provide straightforward protection.

5. Summation

In summation, the breach element of the informed consent action in UK law was redefined less than a decade ago. As such, it is understandable that uncertainties persist in the operationalisation of the standard that was laid down. Nevertheless, arguments can be advanced that the following classes of information must be imparted to the patient: that the AI possesses a

-ai> accessed 11.11.2022. This document referred specifically to practitioners but the solution to relevant biases is, as discussed in Chapter 2, not evidently effected by an individual's level of experience.

1026 *Appleton v Garrett* (1997) 34 BMLR 23.

risk-relevant status, that there have been substantial changes to the human expertise which the patient expected to be applied to the clinical decision, and that a significant choice may be pre-determined if the patient agrees to a seemingly innocuous use of AI.

D. Causation

The causation test connects the outlined elements of a relevant defendant's breach of duty to the actionable damage. Traditionally under UK law the claimant must prove that, but for the breach, they would not have suffered the damage.¹⁰²⁷ *Montgomery* applied this test to the informed consent context, considering what the particular patient would have done had they been advised of the information which the breach of duty has obscured from them.¹⁰²⁸ In other words, post-*Montgomery* the courts apply a subjective test for causation.¹⁰²⁹

Given the findings on actionable damage in UK negligence, one must distinguish two possible applications of this test. In so far as an autonomy violation itself can be counted as the injury suffered by a patient, it is sufficient to show that a requisite, serious breach of the principle has occurred. This is arguably supported by the evaluated case of *Chester v Afshar*. There, given that the non-disclosure of a material piece of information had sufficiently impacted the patient's autonomy, it was not necessary to establish a traditional form of but-for causation. As Amirthalingam has stated, this approach finds 'a causal link between the negligence and the loss of the right to informed consent, but it does not establish causation with respect to the physical injury'.¹⁰³⁰ The causation element simply asks whether these shortcomings did impact facets of decision making that are important to the individual patient. As the three arguably actionable breaches regarding ML have themselves been framed as possessing a degree of significance, they do not call for a separate analysis at this stage.

1027 *Barnett v Chelsea and Kensington Hospital Management Committee* [1969] 1 QB 428.

1028 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [96]-[105].

1029 Turton, 'Informed Consent to Medical Treatment Post-Montgomery' (2019) 27(1) Medical Law Review p. 108, 109.

1030 Amirthalingam, 'Causation and the Gist of Negligence' (2005) 64(1) The Cambridge Law Journal p. 32, 33-34.

By contrast, if a claim is made for the physical injury that is suffered attendant upon an autonomy violation, then it must be established that, but for the failure to disclose a piece of information to the patient, they would have made a decision that averted the physical injury. This was the kind of claim advanced in *Montgomery* itself, where it was found that the patient would have opted for a caesarean section had she been informed of the possibility of shoulder dystocia eventuating. On a balance of probabilities the baby would then have suffered no harm.¹⁰³¹

In making this finding the courts have had recourse to the kinds of considerations that evidence a patient's longer-held, serious commitments. In other words, those that may be taken as evidence of how they would have exercised the reflective dimension of autonomy. A vivid example is given by *Shaw v Kovac*, where the patient's hypothetical decision to refuse an operation was ascertained by reference to the fact that 'He was mentally alert (...) having the cautious and conservative nature of someone who had been born in the Scottish Highlands'.¹⁰³² By contrast, in *C v Colchester Hospital University NHS Foundation Trust* it was found that the patient would not have taken the risk of not having a surgery, given their well-documented, overwhelming desire to be cured. Causation could therefore not be established.¹⁰³³

In spite of this test's affinity with one element of autonomy, a focus on decision-making related to physical harm would nevertheless limit the protection offered to this principle. It would require a patient to demonstrate that established elements of their character would have led to their rejection of a procedure involving AI use that, in turn, would have avoided a physical injury. Whether this is possible will depend on factors incidental to autonomy.

It is also unclear in how far the courts overcome this evidentiary difficulty by appealing to an objectified patient, rather than to the individual patient's personal commitments. *Diamond v Royal Devon and Exeter NHS Foundation Trust* arguably blurred this line by endorsing an objectively rational decision and then finding that the patient would not have behaved irrationally.¹⁰³⁴ It will be seen in the American analysis that such a position

1031 *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430 [104].

1032 *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [28].

1033 *C v Colchester Hospital University NHS Foundation Trust* [2018] 2 WLUK 850 [61].

1034 *Diamond v Royal Devon and Exeter NHS Foundation Trust* [2019] EWCA Civ 585, (2019) 170 BMLR 49 [9], [19]-[22]; Austin, 'Correia, Diamond and the Chester

would restrict the protection of the individual patient's autonomy even further.

The *Chester v Afshar* autonomy-based exception to ordinary causation principles is not envisaged to provide much assistance to such claims. Subsequent courts have indicated that its analysis was exceptional – only assisting the patient when an injury is intimately connected with a failure to warn and in so far as it would have led the patient to defer or reject a relevant procedure.¹⁰³⁵ As AI's autonomy challenges have not been argued to have any distinct effect on the timing of a patient's decision, *Chester* will not meaningfully expand the scope for personal injury claims that are related to autonomy violations.

Overall, one may make two findings. To the extent that a claimant can claim successfully for a violation of their autonomy interest, the causation element should serve to emphasise that this interference must be of a sufficiently serious type to impact their process of decision making. This is entirely congruous with the outlined autonomy principle and it is not envisaged to separately restrict the identified instances where a patient can recover for undisclosed features of ML use.

Where, however, the patient brings (or must bring) a claim with respect to a personal injury they have suffered, they must go to the additional, not undemanding, lengths of showing that their decision would have been altered so as to avoid the harm.¹⁰³⁶ Although the courts still consider aspects of the patient's reflective autonomy to make out this claim, it imposes a further burden that can restrict the success of autonomy-based claims. Whether this will be the case depends very much on the kind of evidence that is available to patients regarding their wider character and foundational commitments.

Exception: Vindicating Patient Autonomy?' (2021) 29(3) Medical Law Review p. 547, 558. Note also that this understanding of rationality (what Austin refers to as an ideal desire approach) goes beyond Pugh's conception of theoretical rationality discussed in Chapter 3, which left the weighting of relevant interests to the individual.

1035 *Correia v University Hospital of North Staffordshire NHS Trust* [2017] EWCA Civ 356, [2017] 5 WLUK 285 [24], [28].

1036 For an analysis of this see: Austin, 'Correia, Diamond and the Chester Exception: Vindicating Patient Autonomy?' (2021) 29(3) Medical Law Review p. 547.

E. Awarding damages

Under UK tort law the fundamental aim of an award of damages is to put the claimant in the position that they would have been in had they not suffered the tortious injury.¹⁰³⁷ However, as Jones has eloquently stated regarding non-pecuniary loss, including not just injuries to intangible interests like autonomy, but also instances of personal injury that impair the patient's physical capacities:

The award is inevitably an arbitrary one. In practice the courts adopt a tariff or "going rate" for specific types of injury in an attempt to achieve some degree of consistency between claimants with similar injuries and to provide a basis for the settlement of claims.¹⁰³⁸

In the examined cases, where an autonomy injury has been deemed a legally cognisable loss, the arbitrariness of this award has been singled out as a particular focus of criticism.¹⁰³⁹

If the real loss suffered in these situations is to the personal autonomy of the individual, as it has been argued, then it is open to the courts to make conventional awards, akin to that awarded in *Reese*, in other non-disclosure cases.¹⁰⁴⁰ Ordinarily this will be a more appropriate response than compensating for a personal injury, which was not strictly speaking caused by the defendant, as occurred in *Chester*.¹⁰⁴¹ By comparison, if autonomy harms and other recognised categories of harm, such as personal injury, are caused simultaneously by a defendant's breach then it may be a justifiable step for the law – in order to maintain coherence and simplicity – to determine an award for autonomy protection through the rules applicable to that other category of injury.¹⁰⁴²

1037 *Livingstone v Rawyards Coal Co* (1880) 5 App Cas 25, 39.

1038 Jones, *Medical Negligence* (Sixth Edition 2021) para 12-003.

1039 Keren-Paz, 'Gender Injustice in Compensating Injury to Autonomy in English and Singaporean Negligence Law' (2019) 27(1) *Feminist Legal Studies* p. 33, 45; Priaulx, *The Harm Paradox: Tort Law and the Unwanted Child in an Era of Choice* (2014) 76.

1040 Maclean, *Autonomy, Informed Consent and Medical Law: A Relational Challenge* (2009) 189.

1041 Keren-Paz, 'Gender Injustice in Compensating Injury to Autonomy in English and Singaporean Negligence Law' (2019) 27(1) *Feminist Legal Studies* p. 33, 45.

1042 Keren-Paz, 'Compensating Injury to Autonomy in English Negligence Law' (2018) 26(4) *Medical Law Review* p. 585, 590 (fn. 37), 600-601. This is arguably what

To the extent that any special issues arise from our definition of the autonomy principle and its application to the informed disclosure of ML characteristics, it is arguable that these concern the amount to be awarded in a standalone autonomy claim. The principle does not provide a general yardstick with which to determine an appropriate amount of damages. At most, the award made for breaches of informed consent should, as Jones has stated, be consistent with the awards made in similar cases. In other words, *Rees's* £15,000 sum should serve as a reference point for particularly serious autonomy violations, akin to the denial of an individual's right to opt for or against parenthood. Here, AI's pre-determination of significant diagnostic choices provides the closest analogy.

Reductions in professional expertise or the non-disclosure of AI's risk related status arguably have lesser impacts and thus justify lesser autonomy-based awards. As the analysis in Chapter 3 has indicated, they impair but do not preclude the exercise of procedural autonomy – even where the relevant decisions have serious consequences. Ultimately one may therefore doubt whether a theoretically available award of damages, which ought to be acknowledged in such circumstances, would amount to an effective form of protection in practice, given the minimal awards that are likely to be made and the likely cost of bringing a tort claim.¹⁰⁴³

III. The UK General Data Protection Regulation

A separate in-depth examination of data protection law was rejected in Chapter 1. Nevertheless, one would be remiss not to mention that a subset of obligations under the *United Kingdom General Data Protection Regulation* (UK GDPR) and the associated *Data Protection Act* (DPA) 2018 could supplement one particular shortcoming identified under the common law causes of action.¹⁰⁴⁴ Namely, the lack of any detailed institutional disclosure

occurred in *Montgomery*. Nolan appears to accept the possibility of this position too: Nolan, 'Negligence and Autonomy' [2022](2) p. 356, 366-367.

1043 Clark and Nolan, 'A Critique of *Chester v Afshar*' (2014) 34(4) *Oxford Journal of Legal Studies* p. 659, 685. Contrast the argument in *Shaw v Kovac* that awards of relatively modest amounts may spur claims, in the hopes of forcing a settlement: *Shaw v Kovac* [2017] EWCA Civ 1028, [2017] 1 WLR 4773 [82].

1044 Following the UK's separation from the European Union, the UK adheres to the United Kingdom General Data Protection Regulation. This was implemented through: Section 3 European Union (Withdrawal) Act 2018 and the Data Pro-

obligation regarding AI use. Indeed, the provisions that will be identified as relevant to this circumstance are inapplicable to other situations that have been argued to constitute the great majority of clinical situations that involve human mediation of ML devices. It will be seen that, for UK GDPR's requirements to shape relevant disclosure obligations, there must be 'a decision based solely on automated processing'.¹⁰⁴⁵

One can therefore focus on institutional AI use. The example cited above concerned a hospital using an ML device to triage patients on the basis of the information provided by them. This would uncontroversially render the healthcare institution a data controller under article 4(7) UK GDPR with relevant obligations while they process the patient's data.¹⁰⁴⁶ One such obligation is to obtain the patient's consent for 'significant decisions based solely on automated processing'.¹⁰⁴⁷ They also have a duty to inform the patient 'at the time when personal data are obtained' of 'the existence of automated decision-making [and provide] meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject'.¹⁰⁴⁸

tection, Privacy and Electronic Communications (Amendments etc) (EU Exit) Regulations 2019.

1045 Article 22(1) UK GDPR; Section 14(1) DPA 2018. As has been previously discussed, human mediation of AI decisions must be a matter of degree and this is recognised under the GDPR too: 'To qualify as human involvement, the controller must ensure that any oversight of the decision is meaningful, rather than just a token gesture': Article 29 Data Protection Working Party, 'Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679' (2018) 20. However, so long there is more than a token acceptance of AI recommendations, it is unclear what kind of human oversight would be insufficient. In relation to our identified types of medical ML devices, this limits the practical relevance of UK GDPR and associated legislation to the purely institutional use of AI.

1046 Article 4(2) UK GDPR. See also the example of a GP surgery on the website of the Information Commissioner's Office: Information Commissioner's Office, 'What are 'controllers' and 'processors'?' (17.10.2022) <<https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/controllers-and-processors/what-are-controllers-and-processors/#1>> accessed 18.3.2023.

1047 Section 14(1), (3)(c) DPA 2018, deriving from Article 22 UK GDPR. For a detailed analysis of the configuration of this requirement in the triaging context, see: Mourby, Ó Cathaoir and Collin, 'Transparency of Machine-Learning in Healthcare: The GDPR & European Health Law' (2021) 43 Computer Law & Security Review.

1048 Article 13(2)(f) UK GDPR. See similarly Article 14(2)(g) UK GDPR, where data is obtained from a third party, rather than the patient, and Article 15(1)(h) UK GDPR, seemingly conferring an *ex post* right of access to meaningful information.

To be clear, these would be obligations outside of the common law. It would not be equivalent, or give rise, to a duty of care in negligence.¹⁰⁴⁹ Instead, medical AI uses must be brought under the statutory obligations and their requirements must be interpreted. As argued in Chapter 4, the autonomy principle has been relied upon also for these purposes in UK law. This provides an additional, forceful reason to examine this aspect of the UK GDPR: a reliance on principle is arguably appropriate and necessary given the substantial ambiguity involved in the GDPR's specification of 'meaningful information' and the lack of case law on the definition of this element.

In defining the requirement one can say, as a point of departure, that the data controller will necessarily have to disclose the fact that it is using an automated decision system. In our example, a patient must be told that a triage decision will be made by an ML device without meaningful human involvement. So much seems uncontroversial.¹⁰⁵⁰ It also represents an improvement *vis-à-vis* the common law position from the perspective of patient autonomy, given that only a limited duty to warn had been constructed for institutions under negligence.

It is also evident from the statutory text that more information must be provided: the logic and envisaged significance of the decision. However, it is hotly contested what information falls under these heads.¹⁰⁵¹ Without authoritative or persuasive legal sources on the issue, a case of uncertain

1049 Under the earlier Data Protection Act 1998 the imposition of a co-extensive duty in negligence had been forcefully rejected: *Keith Smeaton v Equifax plc* [2013] EWCA Civ 108, [2013] 2 All ER 959 [75].

1050 E.g. Hamon and others, 'Impossible Explanations?: Beyond Explainable AI in the GDPR From a COVID-19 Use Case Scenario' (FAcCT '21: 2021 ACM Conference on Fairness, Accountability, and Transparency, 03.03.2021 - 10.03.2021) 557; Mourby, Ó Cathaoir and Collin, 'Transparency of Machine-Learning in Healthcare' (2021) 43 Computer Law & Security Review, 5-6.

1051 Most of these debates are conducted regarding the General Data Protection Regulation under EU Law (EU GDPR), but they bear obvious relevance to the UK's scheme, which has taken over the relevant, outlined aspects. See: Selbst and Powles, 'Meaningful Information and the Right to Explanation' (2017) 7(4) International Data Privacy Law p. 233; Wachter, Mittelstadt and Floridi, 'Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation' (2017) 7(2) International Data Privacy Law p. 76; Custers and Heijne, 'The Right of Access in Automated Decision-Making: The Scope of Article 15(1)(h) GDPR in Theory and Practice' (2022) 46 Computer Law & Security Review p. 105727; Hoeren and Maurice Niehoff, 'Artificial Intelligence in Medical Diagnoses and the Right to Explanation' (2018) 4(3) European Data Protection Law Review p. 308; Goodman and Flaxman, 'European Union Regula-

strength can be made for a form of disclosure that responds to the AI challenges, as identified under the concept of autonomy in UK law.

In this respect, the ability of AI to independently pursue objectives falls most naturally under the significance of anticipated consequences. As outlined above, a patient ought to be made aware of the identifiable outcomes associated with the AI's purposes – particularly where they are surprising and significant. To a lesser degree, this obligation may also be framed to include the specific risks posed by AI, although these would have to be sufficiently likely to qualify as 'envisaged' and sufficiently severe to class as 'significant'. Arguably, this would be a very high bar indeed, one which is unlikely to be satisfied by many AI.

Regarding the logic of AI functioning, a case can be made that a patient should be made aware of the difference between AI expertise and human, professional expertise in a general sense. It has been argued that AI functions without necessarily relying on the same body of scientific knowledge as their human counterparts, and that this is capable of generating decision-relevant uncertainty.

Where a patient has to interact directly with an AI, the framing of UK GDPR may also require forms of disclosures that are additional to those described for negligence. Arguably these should impart that ML devices may provide explanations that are not immediately connected to the underlying functionality – a pervasive issue identified regarding explainable AI in Chapter 2. The legislation's reference to 'logic' should, at a minimum, require information about the different model types in play: one reaching a decision and one explaining that decision – designed not to replicate the decision-making process but also to be understandable. The patient is then in a better position to assess whether they are being influenced towards reaching a certain decision and they can better accommodate the AI output within their own evaluative framework. Similarly, where the kinds of features and criteria used by the AI are known, there is an argument that these must be disclosed to the patient. This can be framed as a general *ex ante* disclosure,¹⁰⁵² required by UK GDPR under the autonomy principle.

tions on Algorithmic Decision-Making and a "Right to Explanation" (2017) 38(3) AI Magazine p. 50.

1052 Framing it as an aspect of the obligation under Article 13 EU GDPR: Wachter, Mittelstadt and Floridi, 'Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation' (2017) 7(2) International Data Privacy Law p. 76, 82-83.

To conclude, the statutory requirements under DPA 2018 and UK GDPR require an institution that is relying solely on an AI/ML device for certain medical purposes, such as triaging, to give the patient notice that this is occurring. On top of this, there must be disclosure of: potential, significant outcomes; the general differences between human and AI functioning; the presence of different models with different objectives; known criteria and features relied on by the AI. If such disclosure is not made, then a claim may be possible for compensation, provided that the patient has suffered material damage or immaterial damage, including distress.¹⁰⁵³

IV. Conclusion

In summation, UK law has been argued to provide three mechanisms through which AI's autonomy challenges can be partially addressed. There is scope for argumentation that the tort of battery would, in circumstances where there is direct contact with the claimant and where an intention to use the AI is formed before such contact, obligate a professional to disclose the broad purpose of that device. This includes its ability to make certain serious choices.

The negligence mechanism also supported such disclosure. In addition, weaker arguments could be advanced that advice must be proffered regarding AI's risk-relevant status and substantial shifts in human expertise – at least where the patient had some expectation that the person undertaking their care is a specialist. For all such disclosures the negligence action imposed conditions that could severely restrict the availability of claims. For example, it was not possible to argue that institutions possess duties of care with demanding informational components.

By comparison, the force of other limiting elements was uncertain. Under the authority of the UK's highest court, it was possible to argue that (1) UK law recognised a significant autonomy violation, without more, as an actionable form of damage and (2) the significance of an autonomy violation at the breach stage would bring about a relaxation of the test at the causation stage. Subsequent case law has not sought to affirm or extend such pronouncements, however. In consequence, negligence has arguably been interpreted to provide a coherent framework for autonomy protection,

1053 Article 82 UK GDPR; Section 168 DPA 2018.

especially under the influence of the procedural autonomy principle, but how far this is accepted in practice is in serious doubt.

Finally, it was argued that a targeted legislative intervention has provided a complementary form of protection to battery and negligence, by imposing informational obligations on institutions. This would require the disclosure of unmediated AI use and potentially, albeit much less apparently, AI's more specific autonomy-related characteristics.

Chapter 7: Californian tort law

In the United States of America, as in England, it is not the abstracted principle of patient autonomy, but its implementation through specific common law mechanisms that will determine the informational duties owed to patients regarding medical uses of artificial intelligence (AI). Both battery and negligence were identified as the modalities with the most potential in this respect. It is the task of this chapter to evaluate the capacity of these torts to meet the outlined, novel autonomy challenges introduced by machine learning (ML).

In order to conduct this analysis in the requisite depth, and with a sufficiently nuanced understanding of the applicable rules-based framework, the tort law of California has been selected to serve as a case study. California provides an instructive example partly because the pronouncements of its courts have had a significant impact on the development of informed consent law in the other states – setting persuasive precedent and inspiring innovation – and partly because its rules typify wider American trends.¹⁰⁵⁴ These trends have led to a more stringent adherence to the doctrinal underpinnings of tort law, representing a notable contrast with the analysis of the previous chapter.

As outlined in Part II, the autonomy principle will nevertheless have an integral role to play in our assessment. To gauge the common law's response to novel problems one must not limit oneself to the established rules, but one must also be capable of anticipating their potential for development. In the following, I therefore analyse how California's torts of battery and negligence can be applied to the challenges of medical AI's implementation, while allowing for realistic developments in line with the principle of patient of autonomy.

Generally, these torts will have a *post facto* role: awarding compensation for violations that have already occurred. In so far as battery is concerned,

1054 These will be pointed out in the course of the analysis. Examples include: tight restrictions on the individuals on whom informed consent duties can be imposed, a narrow approach to legally compensable injuries in negligence and the (determinative) recourse to the reasonable person in shaping negligence's breach and causation requirements.

a close connection to the patient's right to privacy will, for a restricted subset of medical decisions, also enable a constitutional claim to be brought on the permissibility of certain actions. Although AI's autonomy problems are expected to remain outside of the scope of this right (as discussed in Chapter 5), it should be mentioned that even here the outlined actions and elements will provide a central reference point and, as such, deserve close attention.¹⁰⁵⁵

I. Battery

A battery constitutes both a tort and a crime. Although our focus is on the former, the Californian courts have generally operated under the assumption that the statutory, criminal definition can also be applied to the common law tort.¹⁰⁵⁶ A preliminary understanding can therefore be gleaned from California Penal Code section 242, which states: 'A battery is any willful and unlawful use of force or violence upon the person of another'.¹⁰⁵⁷ An individual's consent to a tortious battery will generally provide a defence to the one perpetrating the wilful use of force or violence.¹⁰⁵⁸

With a preliminary definition in hand, one can understand why Californian courts, representing a theme in the United States more generally,¹⁰⁵⁹ utilised the battery cause of action early on to establish the legal requirement that a patient's consent is necessary for medical treatment. Such claims tended to involve invasive surgeries that were performed without authorisation. For instance, in *Valdez v. Percy* the plaintiff authorised the surgeon to remove an enlarged axilla gland but alleged that she had not consented to the procedure which was then performed during the same surgery: the removal of her right breast. The Court of Appeal held that,

1055 See for example: *Barber v. Superior Court* (1983) 147 Cal.App.3d 1006.

1056 *McChristian v. Popkin* (1946) 75 Cal.App.2d 249, 260; *Fraguglia v. Sala* (1936) 17 Cal.App.2d 738, 742. However: 'the torts of assault and battery are not defined by statute, and the court is afforded the opportunity to extend the concept of the tort beyond the limits placed on the corresponding crime by its statutory definition': *California Jurisprudence* (Third Edition 2022), Assault and Other Willful Torts § 27.

1057 California Penal Code section 242.

1058 Witkin, *Summary of California Law* (Eleventh Edition 2022), Torts § 457.

1059 Faden, King and Beauchamp, *A History and Theory of Informed Consent* (1986) 120-125; Schultz, 'From Informed Consent to Patient Choice: A New Protected Interest' (1985) 95(2) *The Yale Law Journal* p. 219, 224-226.

without obtaining the plaintiff's consent, these actions would constitute an actionable battery.¹⁰⁶⁰

On the basis of *Valdez* it was subsequently asserted in *Estrada v. Orwitz* that '[t]here can be no doubt that an action based on the theory that an operation is performed without consent of the patient charges an assault and battery, and that negligence has nothing to do with such an action'.¹⁰⁶¹ Hence, when the plaintiff argued that certain teeth had been extracted without consent, the court determined that they had charged a battery.¹⁰⁶² Many other cases could be adduced to support this proposition.¹⁰⁶³ Collectively they establish that, to avoid a liability in battery, a medical treatment must ordinarily be based on the patient's consent.

This has also connected the battery cause of action to the protection of patient autonomy. Both courts and commentators have recognised that battery serves a distinct role in this endeavour. In particular, it recognises the patient's ability to exercise control over their body as a significant instantiation of the right to make autonomous decisions.¹⁰⁶⁴ In *Thor v. Superior Court* the Californian Supreme Court adduced the outlined battery case law as the first piece of evidence for the common law's long standing tradition of protecting the principle of patient autonomy.¹⁰⁶⁵ The principle was not only seen as intertwined with the informed consent doctrine ('a corollary'),¹⁰⁶⁶ but it was also seen to further the patient's ability to undertake their own balancing of relevant interests.¹⁰⁶⁷ The recent case of *Stewart v. Superior Court* similarly drew on federal and state law to define 'the right to personal autonomy' that was denied a patient by the doctors' decision to sign the consent form for them.¹⁰⁶⁸

1060 *Valdez v. Percy* (1939) 35 Cal.App.2d 485, 491-492.

1061 *Estrada v. Orwitz* (1946) 75 Cal.App.2d 54, 57. Note that the usage of the term 'assault' in this statement represents a common ambiguity.

1062 *ibid* 57.

1063 *Weinstock v. Eissler* (1964) 224 Cal.App.2d 212; *Hundley v. St. Francis Hospital* (1958) 161 Cal.App.2d 800; *Keister v. O'Neil* (1943) 59 Cal.App.2d 428.

1064 'Patient autonomy was initially identified with and subsumed under an interest in physical security, protected by rules proscribing unconsented touch': Schultz, 'From Informed Consent to Patient Choice: A New Protected Interest' (1985) 95(2) *The Yale Law Journal* p. 219, 224.

1065 *Thor v. Superior Court* (1993) 5 Cal.4th 725, 735.

1066 *ibid* 735.

1067 *ibid* 734-736.

1068 *Stewart v. Superior Court* (2017) 16 Cal.App.5th 87, 104-107.

Moreover, battery's role in the protection of this interest has also been recognised to have a constitutional dimension in California. Specifically, this stems from the constitutional guarantee of Privacy under Article I, section 1 of the California Constitution. This provides a right to 'give or withhold informed consent with respect to a proposed medical treatment'.¹⁰⁶⁹ As was already addressed in Chapter 5, analyses of both the California and United States Constitution place heavy reliance on the common law reasoning and development in the shaping of this right. As such, it does not require separate analysis, but it does highlight the preeminent part that the battery cause of action plays in the protection of one of the patient's most fundamental interests. This strengthens the legitimacy of its principled application and development in a relevant healthcare context.

With this fundamental framework in place, two dimensions of the battery cause of action should constitute our focus. First, while battery is associated with these autonomy objectives, the Californian interpretation is subject to rule specific considerations that must be addressed. Second, one must consider the information that must be disclosed to obtain a valid consent to battery.

A. Limitations flowing from the battery doctrine

Within the context of medical treatment, *Ashcraft v. King* provides a summation of battery's requirements that builds upon the above definition and which is frequently utilised by the courts: 'A battery is any intentional, unlawful and harmful contact by one person with the person of another'.¹⁰⁷⁰ This states the three elements that must be considered before engaging with the nature of the patient's consent: the requisite contact; the unlawful or harmful nature of that contact (in spite of the 'and' formulation, it will be seen that fulfilling one of the conditions suffices); and the intent of the person touching the other.

1069 *Foy v. Greenblott* (1983) 141 Cal.App.3d 1, 11.

1070 *Ashcraft v. King* (1991) 228 Cal.App.3d 604, 611. Although there is a general problem of courts using varying, partially conflicting, definitions of these elements, this represents a widely accepted one in the clinical sphere.

1. Contact

Regarding the element of contact, it is clear that absent any touching, as may occur when a physician decides not to treat the patient, there can be no battery and an action lies, if at all, only in negligence.¹⁰⁷¹ This was stated in *Scalere v. Stenson*: ‘Under a battery theory the doctor's failure to disclose the risks and benefits of non-treatment would not be actionable because there was no unconsented touching’.¹⁰⁷² Thus, where an AI leads to a decision not to treat, as where a mere analysis of data provides a false negative diagnosis, a battery cause of action will be unavailable. This is manifestly a restriction unrelated to the patient's autonomy.

Beyond this, one must ask what type of contact between doctor and patient fulfils the contact requirement. Californian case law has set a low threshold in this respect. It objects to ‘touching of any kind’.¹⁰⁷³ For instance, the touching of the plaintiff's hands, arms and shoulder post-surgery were enough to sustain a relevant claim in *So v. Shin*.¹⁰⁷⁴ Requiring such direct, albeit minimal, physical contact is one natural interpretation of the statement that a ‘mere touching’ is required by, as well as sufficient for, a battery.¹⁰⁷⁵

As has already been suggested in the British analysis, this would further limit the circumstances in which a battery claim corresponds to AI's autonomy challenges. It would mean that battery could not occur in a class of actions that may be characterised as indirect physical contacts.¹⁰⁷⁶ In particular, it encompasses the scenario where the doctor gets the patient to

1071 Regarding U.S. tort law generally see: Schultz, ‘From Informed Consent to Patient Choice: A New Protected Interest’ (1985) 95(2) *The Yale Law Journal* p. 219, 229-232; Dobbs, Hayden and Bublick, *Dobbs' Law of Torts: Practitioner Treatise Series* (Second Edition 2022) § 33.

1072 *Scalere v. Stenson* (1989) 211 Cal.App.3d 1446, 1455.

1073 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260. The California Supreme Court also endorsed the statement that “‘It has long been established that ‘the least touching’ may constitute battery. In other words, force against the person is enough; it need not be violent or severe, it need not cause bodily harm or even pain, and it need not leave a mark” in: *People v. Shockley* (2013) 58 Cal.4th 400, 404-405.

1074 *So v. Shin* (2013) 212 Cal.App.4th 652, 671-672.

1075 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 240.

1076 Simons, ‘A Restatement (Third) of Intentional Torts’ (2006) 48 *Arizona Law Review* p. 1061, 1077.

ingest forms of medication,¹⁰⁷⁷ which is not obviously a less serious decision than many forms of medical treatment involving physical touch.

In California no decision directly confronts this point. *Obiter* statements were made in *Freedman v. Superior Court* that an action in battery should not be barred by the manner in which contact was made. Specifically: ‘we would not suggest distinctions in terms of the physical nature of the touching (punctures or cuttings being batteries and ingestion of medication not, for instance)’.¹⁰⁷⁸ This was intermingled with a discussion of the informed consent necessary for treatment, suggesting that responding to grave deviations from the patient’s authorisation to treatment was the primary concern.¹⁰⁷⁹

Some support for this position can also be gleaned from out-of-state case law, which purported to follow the principles underlying a leading Californian case on informed consent. In *Mink v. University of Chicago* the District Court perceived that the mechanics by which a treatment was administered were indistinguishable according to the principle underlying *Cobbs v. Grant*.¹⁰⁸⁰ It was held that a battery claim should be permitted for such indirect contact: ‘[h]ad the drug been administered by means of a hypodermic needle, the element of physical contact would clearly be sufficient. We believe that causing the patient to physically ingest a pill is indistinguishable in principle’.¹⁰⁸¹

Since the question at issue, whether or not there was contact sufficient for battery, goes directly towards the question of whether the patient’s bodily integrity was threatened, *Mink*’s extension of ‘touching’ to indirect contacts must be supported by an alternative justification. Namely, the principle of autonomy. This would explain why the doctrinal validity of the court’s approach remains highly contested.¹⁰⁸² It is because, by accepting the relevance of the principle of patient autonomy, the court interpreted the

1077 *ibid* 1077.

1078 *Freedman v. Superior Court* (1989) 214 Cal.App.3d 734, 740, fn. 2.

1079 *ibid* 740, fn. 2.

1080 *Mink v. University of Chicago* (N.D.Ill. 1978) 460 F.Supp. 713, 716-718. *Cobbs* constitutes a seminal Californian informed consent case, as was discussed in Chapter 5 and will be discussed further below.

1081 *ibid* 718.

1082 See: Simons, ‘A Restatement (Third) of Intentional Torts’ (2006) 48 Arizona Law Review p. 1061, 1077.

battery cause of action in a way that departed from the more traditional interpretation, shaped by narrower value of bodily integrity.¹⁰⁸³

Ultimately, whether Californian courts would follow the same reasoning remains an open question, but there is *obiter* support for this position and it is arguably in line with the autonomy principle underlying the Californian case law that inspired *Mink* itself. As such, we will proceed on the assumption that only a very narrow band of medical cases (those involving non-treatment) is excluded by the contact requirement. All manner of positive treatment plans, which include AI assistance, would fall to be considered under this element of battery.

2. Unlawful nature

The second condition for a successful battery claim is that the touching was unlawful, harmful or – to use a term that has frequently been treated as synonymous – offensive.¹⁰⁸⁴ This restriction bears a strong connection to the patient's autonomy. Conduct must be offensive as defined by law and in the clinical sphere California's courts have consistently equated an offensive touching with one that was without, or exceeded, the patient's consent.¹⁰⁸⁵

Thus, the court in *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* stated that 'Although typically a battery is a violation of a person's wishes to avoid bodily contact that is hostile, aggressive or harmful, the tort is committed if there is unwanted intentional touching of any kind. (...) For example, a person is entitled to refuse well-intentioned medical treatment.'¹⁰⁸⁶ *Rains v. Superior Court*, to which *Conte* referred, found that an intentional deviation from consent was sufficiently objectionable – as, by implication, was 'a physician's good faith effort to effect a cure by

1083 This is something that one alternative explanation for the application of battery to medication, is not able to account for as neatly. Namely, that it constitutes an extension of the doctrine that certain non-bodily contacts (e.g. the touching of some object connected to the body) is sufficient for battery. This would be much less controversial, see: Ezra, 'Smoker Battery: An Antidote to Second-Hand Smoke' (1989) 63(4) Southern California Law Review p. 1061, 1092-1093.

1084 *Rains v. Superior Court* (1984) 150 Cal.App.3d 933, 938; *Barbara A. v. John G.* (1983) 145 Cal.App.3d 369, 375.

1085 *People v. Miranda* (2021) 62 Cal.App.5th 162, 175.

1086 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1266.

exceeding the precise treatment to which consent was given'.¹⁰⁸⁷ Likewise, a contact has repeatedly been described as unlawful for the purposes of clinical battery where it has not been consented to.¹⁰⁸⁸

This means that this limitation, including its application to medical AI, will be co-extensive with the nature of the consent that the law requires. If the consent provided to a contact is invalid or non-existent, then the contact is also offensive. The conditions under which consent is deemed insufficient will be dealt with in Section I.B.

3. Intention

The final limitation on the battery cause of action, which cannot be related to the objective of protecting the patient's ability to make a medical decision for themselves, is the requirement that the doctor must have had a requisite intention. This is a strong rule-specific limitation that flows from battery's status as a paradigmatic intentional tort, a characteristic that distinguishes it from the tort of negligence.¹⁰⁸⁹ Continuing uncertainty persists, however, regarding the intention that is in fact required by U.S. tort law doctrine. Is it merely to do the act in question, such as inserting a scalpel? Is it to inflict some kind of offense or harm? Or is it an intention to act without, or in excess of, the consent of the patient?¹⁰⁹⁰

Californian case law appears to represent a microcosm of this uncertainty, referring to a mixture of these intentions. For example, in *Freedman v. Superior Court* it was held that 'battery requires no showing of "scienter" or any intent to do wrong—only an intent to cause the harmful unconsented touching'.¹⁰⁹¹ In other words, the first position is supported: the

1087 *Rains v. Superior Court* (1984) 150 Cal.App.3d 933, 941.

1088 *Ashcraft v. King* (1991) 228 Cal.App.3d 604, 611; *Fluharty v. Fluharty* (1997) 59 Cal.App.4th 484, 497; *Daley v. Regents of University of California* (2019) 39 Cal.App.5th 595, 602; *Piedra v. Dugan* (2004) 123 Cal.App.4th 1483, 1495.

1089 Dobbs, Hayden and Bublick, *Dobbs' Law of Torts: Practitioner Treatise Series* (Second Edition 2022) § 28.

1090 Moore provides an overview of the confusion in the case law, as well as the lack of clarity in the Second Restatement of Torts: Moore, 'Intent and Consent in the Tort of Battery: Confusion and Controversy' (2012) 61(6) *American University Law Review* p. 1585, 1597-1606.

1091 For this proposition, see also *People v. Miranda* (2021) 62 Cal.App.5th 162, 175-176: 'Battery is a general intent crime (...), which means it requires the intent to do the act involved, not an intent to cause a resulting harm'.

physician must only intend to carry out the act that is classified as offensive. By stark contrast, in *Austin B. v. Escondido Union School District* it was found that if the touching is lawful 'it is appropriate and, indeed required, that the jury be instructed that to be liable for battery, a defendant must intend to harm or offend the victim'.¹⁰⁹²

In medical cases there is generally no intent to cause harm, given the aim of the therapeutic endeavour to confer a benefit on the patient. This means that there must be a direct intention to act unlawfully, i.e. without consent or other legal authorisation.¹⁰⁹³ Merely the intent to do the procedure is insufficient.

Indeed, this was the position adopted by the Court of Appeal in the medical case of *Piedra v. Dugan*. In response to the plaintiff's claim that the doctor had violated a condition that they had placed on their consent, the court asked whether it could be maintained that the doctor had intentionally violated that condition.¹⁰⁹⁴ Here, given that the treating doctor did not know of this condition, and that knowledge of it could not be imputed to him,¹⁰⁹⁵ no battery was found. Agreeing with the trial court, it was ultimately held that the doctor had not 'intentionally rendered a treatment that had not been consented to'.¹⁰⁹⁶ Conversely, one may consider the aforementioned case of *Ashcraft*, where the jury could 'infer an intent to wilfully disregard plaintiff's conditional consent'.¹⁰⁹⁷ This was borne out by the fact that the defendant, having acknowledged the plaintiff's condition to use only family donated blood in their operation, proceeded to use blood from the hospital's general supply.

Taken together, these cases stand for the proposition that there are instances where the patient's consent has been exceeded, and a grave violation of their ability to make their own medical decisions may have occurred, but a battery claim will not succeed because the requisite intention on the part of the defendant is missing. Yet such instances appear limited to

1092 *Austin B. v. Escondido Union School Dist.* (2007) 149 Cal.App.4th 860, 872.

1093 *ibid* 872-873; *Dennis v. Southard* (2009) 174 Cal.App.4th 540, 554; *Barouh v. Haberman* (1994) 26 Cal.App.4th 40, 46-47.

1094 *Piedra v. Dugan* (2004) 123 Cal.App.4th 1483, 1497-1499. The ability of patients to impose conditions on treatment will be explored in the next section.

1095 The condition had only been expressed orally to other employees of the institution and the court found that '[t]here is no authority for imputing knowledge of Fountain Valley employees to Dr. Dugan on the claim for medical battery': *Piedra v. Dugan* (2004) 123 Cal.App.4th 1483, 1498.

1096 *ibid* 1494.

1097 *Ashcraft v. King* (1991) 228 Cal.App.3d 604, 613.

the giving of conditional consent – the purported violation of which was at stake in *Piedra* and *Ashcraft*.

In the great preponderance of cases, dealing with actions that exceed consent, it is suggested that the Californian courts reason primarily on the basis of the nature of the consent obtained and the gravity with which a deviation impacts the patient's ability to make their decision. For the most part, the professional's intent to deviate is then inferred, so that an argument that the scope of the consent had been inadvertently forgotten or deviated from by the defendant would gain little traction.¹⁰⁹⁸ The law presumes in such cases that the doctor could not have had a reasonable belief that they were acting within the consent of the patient and then they will also be found to have had the requisite intention for battery.

This will encompass cases of non-consent, where the patient has not consented to the procedure at all and it will also be relevant to cases where the doctor carries out a substantially different procedure to the one that the patient consented to. Thus, it was stated in *Cobbs v. Grant* that '[w]hen the patient gives permission to perform one type of treatment and the doctor performs another, the requisite element of deliberate intent to deviate from the consent given is present'.¹⁰⁹⁹ The Supreme Court reached a determination on the issue of intent by reference to the nature of the consent and also by reference to the gravity of the violation. In contrast, where 'an undisclosed inherent complication with a low probability occurs, no intentional deviation from the consent given appears' and (partially for this reason) battery is inapplicable.¹¹⁰⁰

It was subsequently made clear in *Dennis v. Southard* that it is indeed a *legal presumption* that the professional has the deliberate intent to deviate from consent in such cases – a presumption that was rightly summarised in the Judicial Council of California Civil Jury Instructions (CACI).¹¹⁰¹ Whereas, for conditional consent cases it was determined that these do not allow for an equivalent legal inference from the doctor's actions.¹¹⁰² In

1098 Moore, 'Intent and Consent in the Tort of Battery: Confusion and Controversy' (2012) 61(6) American University Law Review p. 1585, 1547-1548.

1099 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 240-241.

1100 *ibid* 240-241. See also *Burchell*, where the belief that there was an emergency did not defeat the requisite intent: *Burchell v. Faculty Physicians & Surgeons of Loma Linda University School of Medicine* (2020) 54 Cal.App.5th 515, 526.

1101 *Dennis v. Southard* (2009) 174 Cal.App.4th 540, 544.

1102 *ibid* 544. Arguably, there is uncertainty in how far intention must be proven even in the conditional consent scenario, however. The aforementioned case of *Conte*

People v. Miranda the Court of Appeal also took care to summarise the clinical case law as follows:

The law in the medical context likewise defines the circumstances when a patient's advance consent to a procedure while unconscious means that a doctor has a reasonable belief that the procedure is lawful and thus does not commit a battery by performing it. A doctor commits a battery when deviating from the consent given to perform a substantially different procedure than the one for which consent was given¹¹⁰³

Therefore, Californian law typically only requires that the doctor has an intention to do the offensive act in question. If this act is then performed in circumstances where the law will not attribute to them a reasonable belief that they are acting within the patient's consent, then the element of battery is fulfilled. This includes the performance of procedures that are substantially different to the one consented to and it includes instances where the patient has not consented at all.

4. Summation

In sum, one can see that the courts' stretching of traditional doctrine permits the nature of the patient's consent to shape the wider requirements of the battery action and bears testament to the significance attributed to the principle of patient autonomy. These actions highlight that, when considering failures to provide information on AI use, the key issue will be the legal validity of patient consent.

B. Battery and the nature of valid consent

Due to the aforementioned role of consent in establishing the offensiveness of the act, it is incumbent upon the patient to prove that their consent was absent or tainted in such a way as to justify their claim.¹¹⁰⁴ If they can,

refers primarily to the specificity of the condition on consent, rather than the doctor's intent to deviate from it: *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1269.

1103 *People v. Miranda* (2021) 62 Cal.App.5th 162, 176.

1104 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1266.

then a battery action provides one mechanism for realising the autonomy principle's demand for the patient to be informed in their medical decision making. As will be seen, battery requires at a minimum that a patient is informed as to the basic kind of procedure they are being subjected to, the identity of the professional providing their care and certain motivations that may be involved in that professional's decision making.

Negligence will not be able to provide a full replacement for these claims, given the greater restrictions that legal doctrine imposes on this tort.¹¹⁰⁵ Moreover, it must be added that the intervention of the Californian legislature, through the *Medical Injury Compensation Reform Act* (MICRA), substantially limits plaintiffs' prospects for recovery under a professional negligence claim, but generally not under battery.¹¹⁰⁶ This currently makes the latter action all the more attractive to plaintiffs.¹¹⁰⁷

As was highlighted in the overview, battery establishes the requirement that a patient's consent is required for a medical procedure. Conversely, 'one who consents to a touching cannot recover in an action for battery'.¹¹⁰⁸ In relation to AI, with the largely assistive functions highlighted in Chapter 2, one should however assume that some consent to a wider procedure has been given by the patient. The crucial question thereby becomes not one of consent and non-consent, but rather a question of the validity conditions that the battery mechanism imposes on a given consent. One must ask whether AI's unique characterises, with their defined impact on patient autonomy, are cognisable as challenges to the validity of consent, so as to fulfil this requirement of the battery action.

The way to frame these problems under Californian law is to ask whether AI use *per se*, or certain characteristics thereof, can (when left undisclosed) render a procedure substantially different to the one that was *prima facie* consented to. If so, and the contact requirement is fulfilled, then liability in battery will ensue. This is the first aspect to be addressed here. A second

1105 Moore, 'Intent and Consent in the Tort of Battery: Confusion and Controversy' (2012) 61(6) *American University Law Review* p. 1585, 1646.

1106 *Perry v. Shaw* (2001) 88 Cal.App.4th 658, 668; *Larson v. UHS of Rancho Springs, Inc.* (2014) 230 Cal.App.4th 336, 348-349; *Saxena v. Goffney* (2008) 159 Cal.App.4th 316, 324-325.

1107 'A problem that sometimes arises is when a plaintiff hoping to evade the restrictions of MICRA, will choose to assert intentional torts': *Unruh-Haxton v. Regents of University of California* (2008) 162 Cal.App.4th 343, 352-353.

1108 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1266.

question would arise if a patient were to make their consent conditional on non-AI use.

1. Substantially different procedure

Cobbs v. Grant expressly limited the relevance of the battery action in circumstances where the patient has given some form of consent, but alleges that this was invalidated by a lack of information.¹¹⁰⁹ California's highest court held that the mechanism should be invoked only in situations where this deficit led to the procedure being 'substantially different' to the one consented to.¹¹¹⁰ A legal obligation to disclose more nuanced types of information would flow, if at all, from professional negligence.¹¹¹¹ Henceforth battery could be claimed only for particularly grievous interferences with medical decision making. Namely, in situations where the consent procedure does not determine the essential character of the procedure that the patient is subjected to.¹¹¹²

The cases relate various types of shortcomings to this criterion of a 'substantially different' procedure and the AI challenges identified in Chapter 3 must be evaluated in relation to these. In the course of this analysis, some ambiguity must be tolerated as a result of the jury's prominent role in U.S. tort cases. As the jury is responsible for applying the substantially different criterion to the facts, there is often no need for a specific legal determination of the matter.¹¹¹³ Furthermore, as was stated in *Kaplan v. Mamelak*,

1109 This can be contrasted with earlier cases, where a wider role was envisioned for battery: *Dow v. Kaiser Foundation* (1970) 12 Cal.App.3d 488.

1110 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 239.

1111 *ibid* 240-241. This has been clearly affirmed in subsequent cases. For instance, in *Saxena v. Goffney* it was stated: '[o]ur high court has made it clear that battery and lack of informed consent are separate causes of action', and later '[p]erforming a medical procedure without informed consent is not the same as performing a procedure without any consent': *Saxena v. Goffney* (2008) 159 Cal.App.4th 316, 324-327.

1112 *Rains v. Superior Court* (1984) 150 Cal.App.3d 933, 939-941.

1113 For example, in *Kaplan v. Mamelak* it was held: '[i]n the absence of any definitive case law establishing whether operating on the wrong disk within inches of the correct disk is a "substantially different procedure," we conclude the matter is a factual question for a finder of fact to decide and, at least in this instance, not one capable of being decided on demurrer': *Kaplan v. Mamelak* (2008) 162 Cal.App.4th 637, 647. This illustrates the more general claim of Gardner that: 'US legal systems, unlike other common law legal systems, tend to use jury trials for the

the distinction laid down in *Cobbs* does not lend itself to the formulation of some overarching test.¹¹¹⁴ In consequence, in order to understand the application of the test, one must draw analogies to three identifiable classes of cases that have been defined by the courts and which can be analysed to provide guidance.

i. Physical nature of the procedure

The first class of cases are the ones where the physical manifestation and effects of a procedure are deemed substantially different. For instance, in *Perry v. Shaw* the plaintiff consented to a removal of excess skin, but the physician also performed a breast enlargement procedure. Following *Cobbs*, this could only be characterised as an operation to which the plaintiff had not consented.¹¹¹⁵ In *Burchell v. Faculty Physicians & Surgeons of Loma Linda University School of Medicine* surgery on the plaintiff's penis, during an operation where the plaintiff's consent had been given 'to have a small mass removed from his scrotum', was found to be substantially different to the consented to procedure.¹¹¹⁶ More generally, it has been stated: 'Consent to an operation carries with it a consent to remove an organ or body tissue which is a normal incident to the operation (...), but such a consent does not carry with it permission to remove a different organ or one which is not normally excised as an incident to the operation consented to'.¹¹¹⁷

These framings already exhibit the cases' narrow focus on interventions with a markedly different physical manifestation than the one anticipated.¹¹¹⁸ This holds true in the more contentious, borderline, decisions. So that in *Daley v. Regents of University of California* the heart of the contro-

bulk of their tort litigation. It may be that trial judges tend to pass the buck to the jury on the point in question, and appellate judges then stay clear of it, with the result that the legal position is indeterminate': Gardner, *Torts and Other Wrongs* (2019) 1-2.

1114 *Kaplan v. Mamelak* (2008) 162 Cal.App.4th 637, 646-647.

1115 *Perry v. Shaw* (2001) 88 Cal.App.4th 658, 664.

1116 *Burchell v. Faculty Physicians & Surgeons of Loma Linda University School of Medicine* (2020) 54 Cal.App.5th 515, 524-525.

1117 *Rainer v. Buena Community Memorial Hosp.* (1971) 18 Cal.App.3d 240, 256-257.

1118 Sometimes reference is made to the nature of the procedure in general: 'Cobbs implies that the failure to discuss the nature of the treatment sounds in battery': *Nelson v. Gaunt* (1981) 125 Cal.App.3d 623, 634. But this does not appear to be reflected in the development of the case law.

versy turned on the significance of physical differences – ‘a percutaneous surgery (with access to the organs established by a needle puncture)’, as opposed to ‘an open laparotomy and open hysterotomy’ – rather than an assessment of the wider significance for the patient of such differences.¹¹¹⁹

In *Osborn v. Irwin Memorial Blood Bank* the court was faced with the question of whether the nature of a procedure could be changed by a circumstance that was not immediately connected to the physical nature of the procedure. Namely, the plaintiffs consented to an operation after they had been led to believe that they could not give direct blood donations to their child. In fact, they could have directed such donations.¹¹²⁰ This was a significant circumstance that could (and ought to) have shaped their decision making and impacted their consent. Yet the court curtly dismissed any argument on this point: for their consent to be valid it was enough for the parents to understand the mechanics of the procedure, which involved their child receiving blood transfusions from general supplies.¹¹²¹

Consequently, the potential relevance of this category to medical AI could only emerge from the fact that reliance on the technology may itself constitute a physical change in the procedure. For instance, leading to the use of a different type of device. Moreover, one characteristic of AI use is its impact on an intervention’s risk profile, which could also be argued to impact the physical effects of the procedure. However, neither of these arguments is particularly persuasive.

In the first respect, it is highly improbable that AI use (largely comprised of assisting human actors in the performance of their tasks) would be characterised as a separate, substantially different physical dimension. Californian courts have proved extremely generous in subsuming different aspects of clinical care into one overall procedure to which consent has been given. For example, in *Piedra v. Dugan*, the administration of a drug to a child was seen to be a necessary and nonelective component of a procedure that came under a general consent provided by the parents.¹¹²² In *Kaplan v. Mamelak* it was held to be a question for the jury whether consent to an operation

1119 *Daley v. Regents of University of California* (2019) 39 Cal.App.5th 595, 600.

1120 *Osborn v. Irwin Memorial Blood Bank* (1992) 5 Cal.App.4th 234, 246.

1121 *ibid* 287. See also the case of *Richmond v. Patel* (Dec. 17, 2021, No. B310903) [non-published opinion]: The nature of the procedure was determined by the physical processes taking place when a biopsy was performed, rather than the subsequent use of the plaintiff’s tissue for research.

1122 *Piedra v. Dugan* (2004) 123 Cal.App.4th 1483,1491.

on one spinal disk included a consent to an operation on another.¹¹²³ The comparatively small changes in physical procedures that are anticipated to be brought about by AI hardly exceed these, much more substantial, variations.

Regarding the risks of an intervention, there are some indications that the gravity of the procedure and its actual side effects are relevant to the determination of substantial difference.¹¹²⁴ However, since *Cobbs* the courts have stopped short of invalidating consent on the basis of a procedure's risks alone (irrespective of their manifestation).¹¹²⁵ For our purposes this is well illustrated by *Daum v. SpineCare Medical Group, Inc.* Here, the investigatory status of a device that was implanted into the plaintiff could not invalidate the given consent.¹¹²⁶ More starkly still, in *Stone v. Foster* it was held that under *Cobbs*' categorisation a patient's complete lack of awareness that the procedure involved *any* risks, fell properly to be considered in negligence.¹¹²⁷

Although we have discussed how machine learning technologies may alter the risk profile of an intervention, akin in many ways to forms of innovative treatment, it is thus not anticipated that they will alter its physical characteristics of an intervention in a way that is demanded by the substantial difference criterion. The primarily risk-related matters that flow from the employment of this additional device and the technologies' wider

1123 *Kaplan v. Mamelak* (2008) 162 Cal.App.4th 637, 647.

1124 In *Burchell* the court appears to have considered the substantial difference criterion together with the eventuation of serious risks: 'Barker removed the mass from both the scrotum and the penis, a different and substantially more invasive procedure than had been contemplated. Burchell suffered serious side effects, some of which are permanent and irreversible': *Burchell v. Faculty Physicians & Surgeons of Loma Linda University School of Medicine* (2020) 54 Cal.App.5th 515, 518. *Berkey v. Anderson* even drew a direct connection: 'The procedure as outlined by the doctors obviously entailed much more, both as to comfort and risk. Appellant asked Dr. Anderson, "What is a myelogram; is it like the electromyograms that I have been having?" The jury could have found that this called for more than a few mollifying words which grossly understated the seriousness of the procedure': *Berkey v. Anderson* (1969) 1 Cal.App.3d 790, 804.

1125 In *Kerins v. Hartley* the mere increase in risk to the patient, generated by their surgeon's HIV-positive status, was insufficient to ground a battery claim: *Kerins v. Hartley* (1994) 27 Cal.App.4th 1062, 1077-1078. Although it must be noted that there were special considerations at play in this case: the plaintiff had allegedly imposed a condition that the surgeon should be in good health on the procedure and they were claiming for emotional distress: *ibid* 1066-1067.

1126 *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285, 1313.

1127 *Stone v. Foster* (1980) 106 Cal.App.3d 334, 346-347.

impacts on the non-physical nature of a procedure fall to be assessed under negligence.

ii. Identity of the professional

The next category to be examined is the identity and expertise of the treating doctor, factors external to the medical procedure.¹¹²⁸ This may be affected by AI's impact on the expertise that is brought to bear on the medical decision, as well as by the influence of its determinations on the decision making of human professionals, which was argued to lead to a partial substitution of human decision making.

California provides relatively little case law on this point. From *obiter* statements in *Newhouse v. Board of Osteopathic Examiners* and *Clarke v. Hoek* it can be deduced that an unanticipated third party's intervention in a procedure, whether experienced or not, is liable to be judged a battery.¹¹²⁹ Similarly, it was stated *obiter* in *Rains v. Superior Court* that fraud as to identity would be capable of giving rise to a battery if it affected the essential character of the act.¹¹³⁰ Regarding a more specific characteristic of a physician (in this case their HIV positive status), *Kerins v. Hartley* indicated that a consent may be invalidated. Yet this was seemingly tied to an express condition imposed by the patient regarding their surgeon's health, rather than a matter of substantial difference.¹¹³¹

All in all, one must say that these discussions do not provide purchase for an analogy to be drawn with AI's characteristics. An argument has not been advanced that AI constitutes a third party in itself, nor that it alters the personal identity of a treating physician, but only that it may affect the physician's level of expertise. Under U.S. tort law doctrine, and

1128 Ihekweumere, 'Doctor, Are You Experienced? The Relevance of Disclosure of Physician Experience to a Valid Informed Consent' (2002) 18(2) The Journal of Contemporary Health Law and Policy p. 373, 396.

1129 *Newhouse v. Board of Osteopathic Examiners* (1958) 159 Cal.App.2d 728, 732-733; *Clarke v. Hoek* (1985) 174 Cal.App.3d 208, 218.

1130 *Rains v. Superior Court* (1984) 150 Cal.App.3d 933, 938-940. Albeit the court attributes more weight to the (non-therapeutic) purpose of such actors, than their identity. This will be the focus of the next section.

1131 *Kerins v. Hartley* (1994) 27 Cal.App.4th 1062, 1066-1067.

consistent with *Cobbs*, this discussion fits most naturally within the realm of professional negligence.¹¹³²

Furthermore, any objectionable reliance on AI expertise is something that will occur in the course of a wider procedure that calls for the professional's judgment. *Conte* intimates that such a balancing exercise by the professional will support the finding that the issue should be left to be regulated under negligence. There the surgeon, in the course of a surgery, chose not to fixate a broken shoulder. This was an aspect of medical judgment that should be adjudged under professional negligence standards, so as to not unnecessarily fetter the exercise of clinical discretion.¹¹³³ This supports our conclusion that battery does not provide a mechanism through which the patient can assert a right to be informed about the professional's relative expertise or a supplementation of it *via* AI.

iii. Non-therapeutic motivations

A medical professional who acts with a non-therapeutic intent may be found liable in battery. This is in spite of the physical mechanics of an action remaining the same. Several Californian cases have addressed this matter.

In *Rains v. Superior Court* the California Court of Appeal conceded that the 'alleged violent touching of which plaintiffs now complain are physically identical to the contact consented to',¹¹³⁴ but nevertheless found that the 'nontherapeutic purpose of touching by a psychiatrist goes to the "essential character of the act itself" and thus vitiates consent obtained'.¹¹³⁵ Here, the alleged non-therapeutic purpose was for the defendant psychiatrists to control the patients,¹¹³⁶ but it was also considered that a desire for personal

1132 Iheukwumere, 'Doctor, Are You Experienced? The Relevance of Disclosure of Physician Experience to a Valid Informed Consent' (2002) 18(2) *The Journal of Contemporary Health Law and Policy* p. 373. This view is also supported by the framing of problems of 'ghost surgery' as issues for professional regulation: McDonald, *California Medical Malpractice: Law & Practice* (Revised Edition 2022) § 2:10.

1133 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1269-1270.

1134 *Rains v. Superior Court* (1984) 150 Cal.App.3d 933, 937.

1135 *ibid* 941.

1136 *ibid* 938.

gratification or the mere pursuit of a fee could assume this role.¹¹³⁷ So *v. Shin* is also instructive. The actions of an anaesthetist, in seeking to forestall an embarrassing report rather than seeking to provide patient care, were classified as a battery.¹¹³⁸

These instances are to be contrasted with *Freedman v. Superior Court*, where a patient was deceived as to the purpose of a medication administered while she was giving birth. She was made to believe that the drug was necessary to prevent infection, whereas it was in fact used to induce labour.¹¹³⁹ In spite of this divergence, which can be expected to have impacted the patient's decision making significantly, the Court of Appeal found that, as the defendant's purpose was not alleged to be non-therapeutic, the relevant consent remained valid. Since the mutual end of physician and patient remained the treatment of the latter, the essential character of the treatment remained unchanged.¹¹⁴⁰

This approach again highlights that the factors capable of affecting consent under California's battery action, and thus requiring disclosure, have been very narrowly delimited and without reference to patient autonomy. The principle seeks to foster the patient's decision making and to allow them to weigh their interests for themselves. Allowing physicians to pursue an abstract end of beneficial treatment clearly defeats this goal, endorsing a form of medical paternalism.

More to the point for the present investigation, this broad-brush approach effectively precludes the possibility of battery being used to regulate clinical AI, especially their propensity to pursue, or to serve, purposes that may diverge from the patient's own. The law would almost certainly sanction purposes that are framed primarily in therapeutic terms – although, as will be discussed below in relation to the tort of negligence, there may be related non-therapeutic dimensions. Patient consent would not be invalidated by the pursuit of such purposes. Moreover, even if a non-therapeutic goal could be attributed to an AI or AI developer, it does not seem feasible to impute this purpose to a physician treating the patient with its help. Consequently, battery would not require the disclosure of AI's relatively independent pursuit of objectives.

1137 *ibid* 940-941.

1138 *So v. Shin* (2013) 212 Cal.App.4th 652, 667.

1139 *Freedman v. Superior Court* (1989) 214 Cal.App.3d 734, 737.

1140 *ibid* 738-739.

2. Conditional consent

In addition to the above criterion of substantial difference, the Californian battery doctrine further instantiates a relatively strong right on behalf of the patient to impose individual conditions on their consent. If such a condition is imposed and the intentionality requirement is fulfilled, then a treating physician must normally respect it or be subject to liability.¹¹⁴¹

In theory this would provide patients with the power to determine that an AI should not be used in their care or, if it is, for what purposes. This may be one way for patients who are *ex ante* well informed about AI and its associated difficulties to assert control over the use of the technology.

However, this solution is hardly ideal, given the need for patient awareness and proactive behaviour regarding a modality that largely remains hidden from view. The already examined case of *Conte* further clarifies that any condition would have to be specific. The plaintiff's purported condition on his consent to shoulder surgery, that it be '*with repair*', was deemed too intangible to fulfil this requirement.¹¹⁴²

Our preceding analysis has demonstrated that clinical AI use is already incredibly varied and challenging to define. Demanding of a patient that they delimit the bounds of permissible and impermissible applications of the technology in their care appears unrealistic and ineffective. Rather, the autonomy challenges posed by AI must be solved primarily through the clinician's proactive role as a facilitator of their patient's decision-making.

C. Summation

It was seen that the structural doctrinal limitations flowing from the tort of battery were substantially reshaped to favour the principle of autonomy in the medical sphere. However, when one considered the requirements of valid patient consent, it emerged that the courts operated within narrow categories that did not bear an obvious relation to procedural autonomy considerations.

Focusing the analysis on the similarities and differences of the physical nature of the procedure and ignoring wider factors, which may be of sub-

1141 *Grieves v. Superior Court* (1984) 157 Cal.App.3d 159; *Ashcraft v. King* (1991) 228 Cal.App.3d 604.

1142 *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1269.

stantial significance to a patient's decision making, almost precludes this mechanism from performing a useful role in the regulation of AI's novel, more nuanced autonomy challenges. This impression is further supported by the fact that no reliable analogies could be drawn to cases dealing either with the identity of the professional or their extraneous motivations.

In addition, while there is a well-documented possibility for an individual to impose conditions on the use of AI in their care, it seems unlikely that many patients would be in a position to exercise this discretion effectively (i.e. in a legally recognised way). To maintain the patient's positive freedom in interactions with clinical AI, providing them with an opt-out or conditional opt-in is clearly insufficient.

Ironically, it may be precisely because Californian (and other U.S.) courts have been prepared to loosen battery's wider doctrinal restrictions in the medical context, that it became all the more important to impose rule-adjacent limitations through the case law. Limitations that delineate clearly between the deficiencies in consent that allow for a plaintiff to bring a claim in battery and those that allow them to bring a claim in negligence. Furthermore, the practical relevance of this distinction was undoubtedly heightened by California's unique legislative background, which sought to limit professional negligence actions in the healthcare sector in a number of ways.

On the basis of such limiting factors, Californian common law has opted to offer battery's strong protection only to a very narrow subset of medical cases. For our purposes this means that the burden of protecting the patient against AI/ML's violations of procedural autonomy (strong and weak) is placed squarely on the negligence action.

II. Negligence

The tort of negligence imposes obligations on individuals to act with due care in their interactions with others. Due care is usually measured by reference to the ordinarily prudent person, who serves as a standard both for the actions and omissions of individuals.¹¹⁴³ We will see that, for a successful claim, a plaintiff must show that a defendant owed a relevant duty, fell short of their obligation(s) and caused some legally cognisable injury.

1143 *Fouch v. Werner* (1929) 99 Cal.App. 557, 564.

If these elements are made out, then a plaintiff is entitled to monetary compensation: ‘the obligations will primarily be enforced by the traditional judicial remedy of an action for damages for their breach’, although it is often additionally envisaged that these will have a ‘prophylactic effect’.¹¹⁴⁴

As already outlined, the role of the negligence action in the enforcement of the principle of patient autonomy arises from the California Supreme Court’s decision to tether the emerging doctrine of informed consent to it.¹¹⁴⁵ Although the tort of battery retains a minimal role in this endeavour, securing the validity of a patient’s consent, it is negligence that imposes the most demanding informational duties on healthcare professionals under their overarching duty to take care. More concretely, it requires one to ask whether ‘the doctor in obtaining consent may have failed to meet his due care duty to disclose pertinent information’.¹¹⁴⁶

It should also be noted that, although Californian courts have been prepared to find the basis for a physician-patient relationship in an implied contract,¹¹⁴⁷ the cause of action considered here remains the tortious one based in ordinary negligence. This is primarily because, while it is possible for patient and physician to enter into a contractual relationship and to specify the nature of their relationship and the obligations arising under it,¹¹⁴⁸ the courts ‘have shown an historical distaste (...) to allow verbal artfulness to obscure the fact that a tort cause of action truly underpins the plaintiff’s grievance’.¹¹⁴⁹ Therefore, in the absence of clear evidence of express agreements, the applicable standard is that to be found in general negligence.¹¹⁵⁰

1144 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120, 178-179.

1145 *Cobbs v. Grant* (1972) 8 Cal.3d 229.

1146 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 240-241.

1147 *McNamara v. Emmons* (1939) 36 Cal.App.2d 199, 204.

1148 *Custodio v. Bauer* (1967) 251 Cal.App.2d 303, 315.

1149 McDonald, *California Medical Malpractice: Law & Practice* (Revised Edition 2022) § 1:3. See also *Tunkl v. Regents of University of Cal.*, where the court found that a hospital’s contractual waiver could not stand and could not bring about a deviation from negligence’s general standard of care: *Tunkl v. Regents of University of Cal.* (1963) 60 Cal.2d 92, 102-104.

1150 ‘It is thoroughly settled in California that “In the absence of an express contract, the physician or surgeon does not warrant cures. By taking a case he represents that he possesses the ordinary training and skill possessed by physicians and surgeons practicing in the same or similar communities, and that he will employ such training, care, and skill in the treatment of his patients (...)”’: *McNamara v. Emmons* (1939) 36 Cal.App.2d 199, 205. “To recover for breach of warranty or contract in a medical malpractice case, there must be proof of an express contract

Similarly, while the courts have been prepared to categorise the relationship between physician and patient as fiduciary in nature, this has not defined the form of the plaintiff's action for recovery. To be sure, it has helped to establish the kinds of information that must be disclosed and thus influenced the analysis under the action. This will be discussed further below, regarding the Supreme Court's decision on the disclosure of a physician's personal interest in *Moore v. Regents of University of California*.¹¹⁵¹ Nevertheless, the elements of the claim (including a demanding causation requirement)¹¹⁵² remain anchored in negligence.¹¹⁵³

Ultimately the claim of a patient that they were entitled to some form of information in their medical care – and that this was not provided – is to be brought in negligence, under ordinary negligence principles. Resulting obligations may sometimes be modified or influenced by the contractual and fiduciary nature of the relationship.

The elements that must be proved under these principles in California are summarised in the following formula: 'a legal duty, breach thereof, proximate causation and resulting damage'.¹¹⁵⁴ A slight amendment to this classification is necessary, since it subsumes the requirement that there must be some legally cognisable injury under the causation requirement. Fusing these elements would not do justice to the analysis of whether patient autonomy is directly or indirectly protected through an informed consent claim. Consequently, and consistently with the approach undertaken in Chapter 6, this is the first element to be considered.

by which the physician clearly promises a particular result and the patient consents to treatment in reliance on that promise': *McKinney v. Nash* (1981) 120 Cal.App.3d 428, 442.

1151 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120.

1152 This was highlighted by the dissenting Justice Mosk in: *Moore v. Regents of University of California* (1990) 51 Cal.3d 120, 179-180. See also: Giuffrida, 'Moore v. Regents of the University of California: Doctor, Tell Me Moore' (1991) 23(1) Pacific Law Journal p. 267, 307-309.

1153 *Jameson v. Desta* (2013) 215 Cal.App.4th 1144, 1164-1165.

1154 *Stafford v. Shultz* (1954) 42 Cal.2d 767, 774. Note that this also applies to informed consent cases and that, consistent with these elements, the burden of proof for lack of informed consent is placed on the plaintiff: 'she bears the burden of proving the elements of every legal theory she proffers in support of that negligence claim—including her informed consent theory': *Flores v. Liu* (2021) 60 Cal.App.5th 278, 297-298.

A. Damage

In the absence of a legally cognisable injury, it will not be possible for a plaintiff to bring a claim for negligence.¹¹⁵⁵ This is not always made clear under U.S. and Californian case law and, ordinarily, the eventuation of some form of (uncontroversially accepted) personal injury will be a practically significant factor in many patients' decision to bring a legal action in the first place.¹¹⁵⁶ However, in cases that deal with new classes of damage, or in claims where the timing of an injury must be determined to ascertain the corresponding action's limitation period, this factor has been emphasised.

In the context of professional clinical negligence, the necessity of damage is also well documented. For example, in *Hills v. Aronsohn* it was held that some legally compensable injury or 'damages' is an element of the cause of action.¹¹⁵⁷ In *Turpin v. Sortini* the California Supreme Court aptly described the plaintiff's suffering of an injury as a 'threshold question' for the relevant negligence claim.¹¹⁵⁸ Without passing over this hurdle an action cannot succeed.¹¹⁵⁹

1155 'If the allegedly negligent conduct does not cause damage, it generates no cause of action in tort': *Budd v. Nixen* (1971) 6 Cal.3d 195; 'If the allegedly negligent professional conduct does not cause damage, it generates no cause of action in tort. The mere breach of a professional duty, causing only nominal damages, speculative harm, or the threat of future harm—not yet realized—does not suffice to create a cause of action for negligence': *Van Dyke v. Dunker & Aced* (1996) 46 Cal.App.4th 446, 452.

1156 Weisbard, 'Informed Consent: The Law's Uneasy Compromise with Ethical Theory' (1986) 65(4) *Nebraska Law Review* p. 749, 753-754. *Looney v. Moore* is further instructive. This federal case applied Alabaman law, but also made reference to the negligence-based approach expounded in *Cobbs v. Grant*: 'Alabama common law requires an actual injury to maintain a negligence cause of action—and in the specific context of informed consent claims, so do the majority of other courts. Finally, although it is true that in cases discussing informed consent claims, the Alabama Supreme Court did not list actual injury as a required element for those claims, there was no dispute in those cases that an actual injury existed, and the court was focusing on what constitutes informed consent. Thus, we do not consider the omission of "injury" from the list as dispositive': *Looney v. Moore* (11th Cir. 2018) 886 F.3d 1058, 1069.

1157 *Hills v. Aronsohn* (1984) 152 Cal.App.3d 753, 762, fn. 7.

1158 *Turpin v. Sortini* (1982) 31 Cal.3d 220, 235-236.

1159 '[T]he obstacle to recovery in a wrongful life case is that it is impossible to determine that the plaintiff has been harmed': Kearn, 'Turpin v. Sortini: Recognizing the Unsupportable Cause of Action for Wrongful Life' (1983) 71(4) *California Law Review* p. 1278, 1286-1287.

The damage that is capable of giving rise to a self-standing negligence action must then be of a certain, legally determined, type.¹¹⁶⁰ As has been intimated, personal injury is one kind of damage that triggers the threshold for liability and provides a readily available basis for most clinical negligence claims. In *Lantis v. Condon* it was further recognised that a spouse's loss of consortium constitutes a separate injury in response to a separate right.¹¹⁶¹ By contrast, a type of injury that has been expressly rejected in the medical malpractice suits is the illegitimacy of a child.¹¹⁶² This begs the question: 'which types of injury provide a sufficient basis for an informed consent claim?' This will be central to determining the extent to which patient autonomy can be protected against medical AI's challenges through this action. Three candidates appear most relevant in this respect: personal injury, a setback to the autonomy interest itself and emotional distress.

1. Personal injury

In *Townsend v. Turk* the Court of Appeal cited with approval the statement that: '[s]tandard negligence analysis protects an interest in physical well-being. The doctrine of informed consent injects into the established framework of negligence a concern with patient choice that would otherwise be absent'.¹¹⁶³ This makes clear that the orthodox approach, both in California and the United States, is to understand informed consent as protecting the patient's right to be free from personal injury.

The case of *Warren v. Schechter* illustrates the nuances of this understanding. As is common in informed consent actions, the claim here concerned a procedure for which a surgeon had disclosed a variety of dangers but had failed to disclose one particular risk (metabolic bone disease) from

1160 I use 'damage' to refer to the type of setback of interests that gives rise to a negligence claim; distinct from 'damages' which refers to types of setback that are compensated once an action has been established. The two are sometimes treated synonymously, as in *Hills v. Aronsohn* (1984) 152 Cal.App.3d 753, 762, fn. 7. Nevertheless, the distinction is recognised as in England: Donovan, 'Is the Injury Requirement Obsolete in a Claim for Fear of Future Consequences' (1993) 41(5) UCLA Law Review p. 1337, 1342-1343.

1161 *Lantis v. Condon* (1979) 95 Cal.App.3d 152.

1162 *Alexandria S. v. Pacific Fertility Medical Center, Inc.* (1997) 55 Cal.App.4th 110.

1163 *Townsend v. Turk* (1990) 218 Cal.App.3d 278, 284, citing Schultz, 'From Informed Consent to Patient Choice: A New Protected Interest' (1985) 95(2) The Yale Law Journal p. 219, 232.

which the plaintiff then, subsequently, suffered. It was held that the action only accrued when this risk materialised (when the plaintiff broke her back years later).¹¹⁶⁴ *Warren* provides a concrete illustration of the claim that the damage which provides the basis for an informed consent negligence action is ordinarily the manifestation of a physical risk, rather than the violation of a more intangible interest in autonomy. For, in the latter case, the injury would have occurred immediately.

Where claims concern information pertaining to alternative treatments or non-treatment, this link with personal injury is more indirect and this element will consequently be more difficult to establish.¹¹⁶⁵ A plaintiff will hardly succeed in these instances, unless it is also the case that their physical well-being was impaired by the selection of an alternative or by the pursuit of treatment over non-treatment.

This illustrates why the outlined challenges associated with ML technologies are not automatically subsumed under this classification. On occasion it will be possible to point to a specific risk that has been altered by their use. For example, the harm that results from a failure to discover a progressing disease from a false negative diagnosis. But the more abstract alteration in the understanding of risks, the alteration of expertise, or the pursuit of differing treatment goals will not necessarily be linked to an identifiable setback of an individual's interest in physical well-being.

2. Autonomy interest

If the protection of the patient from bodily harm is necessarily of limited assistance to a patient who feels that their decision-making process is affected by non-disclosure of AI features, then it is natural to enquire whether it may not be open to them to bring a negligence action directly on the grounds that their autonomy interest has been violated. Although judgments are sparse on this issue, several commentators in the U.S. have advocated for the adoption of this position.¹¹⁶⁶

1164 *Warren v. Schecter* (1997) 57 Cal.App.4th 1189, 1204.

1165 Meisel, 'A Dignitary Tort as a Bridge between the Idea of Informed Consent and the Law of Informed Consent' (1988) 16(3-4) *Law, Medicine and Health Care* p. 210, 214-215.

1166 Meisel, favours basing informed consent on a somewhat more elusive dignitarian interest, and usefully outlines the lack of precedent, positive or negative, on the issue: 'few courts have rejected such protection outright, but that is probably

Some indirect, tentative support can be gleaned from various aspects of Californian medical malpractice case law. Early *dicta* for example emphasised the need for the significance of the patient's familiarity with treatment alternatives, a need that appeared to be only indirectly connected to the eventuation of harm.¹¹⁶⁷ Moreover, the seminal decision in *Truman v. Thomas* held that a physician could be liable for a failure to warn a patient of the consequences of opting to not undertake a diagnostic test.¹¹⁶⁸ This arguably opened the door for a more general recognition of the patient's protected interest in informed decision making, beyond a warning of the risks to bodily integrity that would accompany a patient's consent.¹¹⁶⁹

There is also some evidence that causes of action for medical malpractice have received separate recognition in Californian courts due to their impact on personal autonomy. For instance, in *Zambrano v. Dorough* the Court of Appeal held that an injury that resulted in the patient's loss of reproductive capacity was of a different type than the injuries she suffered from the misdiagnosis of a tubal pregnancy (including the rupture of a fallopian tube and unnecessary surgery).¹¹⁷⁰ The basis for this distinction was the significance of the former injury for the patient's reproductive autonomy, rather than the gravity for the physical well-being of the patient, which had already been substantially impacted.

In spite of these indicators and arguments, a standalone negligence claim for the violation of a plaintiff's autonomy has not been recognised in the courts. Indeed, the impact of *Truman* has proven relatively narrow: mandating disclosure regarding the refusal of a procedure only once it has been

because so few have been called on to recognize such a right': *ibid* 211. See also: Weisbard, 'Informed Consent: The Law's Uneasy Compromise with Ethical Theory' (1986) 65(4) *Nebraska Law Review* p. 749, 764 (*inter alia* suggesting 'viewing patient self-determination as a goal independent of the avoidance of physical injury'); Twerski and Cohen, 'Informed Decision Making and the Law of Torts: The Myth of Justiciable Causation' [1988](3) *University of Illinois Law Review* p. 607, 609 ('courts should identify and value the decision rights of the plaintiff which the defendant destroyed by withholding adequate information').

1167 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243.

1168 *Truman v. Thomas* (1980) 27 Cal.3d 285.

1169 Meisel, 'A Dignitary Tort as a Bridge between the Idea of Informed Consent and the Law of Informed Consent' (1988) 16(3-4) *Law, Medicine and Health Care* p. 210, 215-216.

1170 *Zambrano v. Dorough* (1986) 179 Cal.App.3d 169, 172-174.

recommended by a physician.¹¹⁷¹ The common law continues to require the eventuation of a specific physical harm to trigger informed consent claims. *Zambrano's* holding must similarly be classed as of little consequence for the present discussion. Another Court of Appeal decision has declined to follow it altogether, even regarding the narrow purpose of ascertaining whether a plaintiff's claim is part of the same cause of action or constitutes a separate one.¹¹⁷² Therefore, there is little prospect that *Zambrano* will alter the wider view on the threshold issue of damage.

It may be thought that another possibility is to construct a form of autonomy damage by reference to the constitutional right to privacy. One commentator has argued that, in the context of Californian wrongful birth claims, the courts assume that the injury consists in a violation of the plaintiff parents' privacy interest.¹¹⁷³ However, it appears that such a move is of limited assistance. Not least, because the above claim is followed almost immediately by the undoubtedly correct insight that, *vis-à-vis* private actors like doctors, 'tort liability does not arise from the mere existence of constitutional rights'.¹¹⁷⁴

Much more problematic for this line of argument though, is the fact that the class of wrongful birth case law that is alluded to, simply does not bear out the claim. In the leading Californian case on wrongful birth, *Custodio v. Bauer*, the legally cognisable injuries are stated in terms of alternative and comparatively orthodox categories: emotional suffering (of either parent), physical injury and consequential economic expenses.¹¹⁷⁵ If a constitutional right has imbued the common law with a novel autonomy interest, then it appears that this has yet to be discovered.

1171 *Scalere v. Stenson* (1989) 211 Cal.App.3d 1446, 1450; *Munro v. Regents of University of California* (1989) 215 Cal.App.3d 977, 986-988.

1172 *DeRose v. Carswell* (1987) 196 Cal.App.3d 1011, 1023-1026. It was also held in *Massey v. Mercy Medical Center Redding* that a morphine injection without informed consent was not an independent cause of action from the fall that it was allegedly administered to cover up: *Massey v. Mercy Medical Center Redding* (2009) 180 Cal.App.4th 690, 699.

1173 Goebelsmann, 'Putting Ethics and Traditional Legal Principles Back into California Tort Law: Barring Wrongful-Birth Liability in Preimplantation Genetic Testing Cases' (2010) 43(2) *Loyola of Los Angeles Law Review* p. 667, 686.

1174 *ibid* 687.

1175 *Custodio v. Bauer* (1967) 251 Cal.App.2d 303, 322-324. Note that this Californian case avoids the issues that were seen to arise around wrongful birth claims in England. Since there is no attempt to limit recovery under some heads of damage, there is no need to create a separate injury to allow for fair compensation.

The next subsection will illustrate that the categories of injury that negligence can remedy are malleable, not fixed, under American tort law. Nevertheless, without direct, supportive authority on this point – and in light of the decidedly negative trend in recognising autonomy violations as distinct actions – an argument from principle cannot find sufficient purchase here. The doctrinal limitations currently appear too strong to grant it independent protection.

3. Emotional distress

Even if there is no recovery for a violation of the patient's interest in informed decision making *per se*, Californian law has recognised that a negligence action may be founded on another type of intangible injury. Namely, the emotional distress of the patient.

The availability of this action is, by itself, not uncontroversial. Historically, American courts and commentators were content to follow the inherited English position that '[t]he mere temporary emotion of fright not resulting in physical injury is, in contemplation of law, no injury at all, and hence no foundation of an action'.¹¹⁷⁶ This original position is represented in California by *Sloane v. Southern California Railway*. It was stated that 'mental suffering alone will not support an action, yet it constitutes an aggravation of damages when it naturally ensues from the act complained of'.¹¹⁷⁷

Against this background, California became the first U.S. jurisdiction to challenge the *status quo* and to recognise that a physician may be liable for the negligent infliction of emotional distress to a direct victim.¹¹⁷⁸ In *Molien v. Kaiser Foundation Hosp.* a physician was held to be liable to the plaintiff for negligently diagnosing his wife with syphilis and advising her

1176 Throckmorton, 'Damages for Fright' (1923) 57(6) American Law Review p. 828, 835-836.

1177 *Sloane v. Southern Cal. Ry. Co.* (1896) 111 Cal. 668, 679-680. See also *Espinosa v. Beverly Hospital* for a case of medical malpractice, where the trial court was found to have properly instructed the jury: '[u]nless you find that plaintiffs suffered actual physical injury as the proximate result of defendants' negligence, they cannot recover in this case because fright or mental suffering alone will not sustain a recovery for plaintiffs': *Espinosa v. Beverly Hospital* (1952) 114 Cal.App.2d 232, 234.

1178 Meisel, 'A Dignitary Tort as a Bridge between the Idea of Informed Consent and the Law of Informed Consent' (1988) 16(3-4) Law, Medicine and Health Care p. 210, 212-213.

to communicate this to him.¹¹⁷⁹ The plaintiff did not suffer a physical injury as a result, but merely an extreme emotional reaction. Without more, this was found to be a compensable injury for the purposes of negligence.¹¹⁸⁰

Viewed in the abstract, it is imaginable that such a claim could form one means of expanding the application of informed consent requirements to medical AI. The crux of Part I's theoretical argument maintains that the patient's process of decision making is impacted by an uninformed use of ML technologies. It is conceivable that some of these uses will provoke an emotional reaction in the patient. Covert manipulation of choices, or the subjection of a medical decision to a bias regarding sensitive or protected characteristics, may serve as good examples.

However, *Molien* and subsequent decisions have imposed demanding limitations on recovery for emotional distress which make an action for this injury ill-suited to informed consent claims in general,¹¹⁸¹ and claims concerning medical AI in particular. The precise nature of the test(s) to be applied remains confused.¹¹⁸² Depending on the circumstances, the Californian courts have variously required: a reasonable foreseeability of emotional harm,¹¹⁸³ a guarantee of genuineness of that harm,¹¹⁸⁴ or proof that a specific injury is more likely than not to occur.¹¹⁸⁵

In spite of the confusion, it appears that any of these restrictions would exclude most forms of AI use. The outlined aspects of the technology may no doubt appear unnerving to patients, and more so to some than to others. Yet, in all but the most serious cases, it is not reasonably foreseeable that emotional distress will result from the technology's application.¹¹⁸⁶ Neither

1179 *Molien v. Kaiser Foundation Hospitals* (1980) 27 Cal.3d 916.

1180 *ibid* 930-931.

1181 Meisel, 'A Dignitary Tort as a Bridge between the Idea of Informed Consent and the Law of Informed Consent' (1988) 16(3-4) *Law, Medicine and Health Care* p. 210, 212-213.

1182 Heidenreich, 'Clarifying California's Approach to Claims of Negligent Infliction of Emotional Distress' (1995) 30(1) *University of San Francisco Law Review* p. 277, 298-302.

1183 *Molien v. Kaiser Foundation Hospitals* (1980) 27 Cal.3d 916, 923.

1184 *Burgess v. Superior Court* (1992) 2 Cal.4th 1064, 1079.

1185 *Potter v. Firestone Tire & Rubber Co.* (1993) 6 Cal.4th 965, 997. This was applied in *Kerins v. Hartley* to a plaintiff's fear of catching HIV from her surgeon: *Kerins v. Hartley* (1994) 27 Cal.App.4th 1062, 1073-1074.

1186 For a critical analysis of the difficulties in applying the reasonableness requirement, especially regarding plaintiff's with an existing predisposition see: Chillag, 'Negligent Infliction of Emotional Distress as an Independent Cause of Action in

can a ‘genuine’ emotional grievance be anticipated, nor a link to a future eventuation of a disease that is more probable than not.

4. Summation

The Californian negligence action, in line with wider trends of U.S. law, is realistically expected to require the eventuation of physical injury to the patient before they can bring a successful informed consent claim. This is not fatal to our envisaged claim for lack of informed consent because, in the high-stakes medical arena, there will be circumstances where a patient suffers physical injury as a result of another’s use of AI. Depending, of course, on the fulfilment of negligence’s other elements, a claim could then be made out.

Still, this limitation means that even serious intrusions into a patient’s decision-making procedures will only be actionable upon the incidental occurrence of physical harm. This is far removed from a coherent protection of the patient from AI’s autonomy challenges.

B. Duty of care

To succeed in an informed consent claim, a patient will further have to establish that a relevant party owed them a duty to disclose such information. To fulfil this element it would be sufficient in most circumstances to point to the general duty of care that the Californian legislature has created.¹¹⁸⁷ However, this duty applies only to situations where a defendant actively creates a risk for another; so that, for example, not every individual has a duty to take affirmative action to assist another.¹¹⁸⁸ It will be seen that the duty of informed consent has been similarly limited to certain individuals. In the following, I consider two classes of persons who may be argued to owe a duty to obtain the informed consent of a patient: medical professionals and healthcare institutions.

California: Do Defendants Face Unlimited Liability’ (1982) 22(1) Santa Clara Law Review p. 181, 195-198.

1187 ‘Everyone is responsible, not only for the result of his or her willful acts, but also for an injury occasioned to another by his or her want of ordinary care or skill in the management of his or her property or person’: California Civil Code section 1714, subdivision (a).

1188 *Brown v. USA Taekwondo* (2021) 11 Cal.5th 204, 213-214.

The potential duty of AI developers and manufacturers to convey information to the patient will not be explored further. In line with the law of all other states,¹¹⁸⁹ California has adopted the learned intermediary doctrine.¹¹⁹⁰ Where a doctor utilises a medical device in a medical environment and under medical supervision, this doctrine applies and holds unquestionably that developers' informational obligations are directed towards the doctor.¹¹⁹¹ As has been stated at length, such a professional utilisation is the normal situation envisaged for AI use. Here, similarly to many other medical products 'the physician stands in the shoes of the product's ordinary user'.¹¹⁹²

Moreover, even in the small subset of cases where the patient is placed in a position to make use of AI directly, in a non-clinical environment, any envisaged duty (arising either from negligence or strict product liability) is a very narrow one: to 'warn of a particular risk'.¹¹⁹³ This duty does not provide a suitable vehicle for requiring information on AI's unique features and therefore need not be explored further.¹¹⁹⁴

1. Medical professionals

It is not disputed that a physician has a duty to obtain their patient's informed consent. Already in *Salgo v. Leland Stanford Jr. University Bd.*

1189 'Every state in the country, along with the District of Columbia and Puerto Rico, has adopted the learned intermediary doctrine in some iteration': *Dearinger v. Eli Lilly and Company* (Wash. 2022) 199 Wash.2d 569, 574-575.

1190 This was recently reaffirmed in: *Amiodarone Cases* (2022) 84 Cal.App.5th 1091.

1191 *Bigler-Engler v. Breg, Inc.* (2017) 7 Cal.App.5th 276, 319.

1192 *Gall v. Smith & Nephew, Inc.* (2021) 71 Cal.App.5th 117, 122.

1193 *Bigler-Engler v. Breg, Inc.* (2017) 7 Cal.App.5th 276, 312-317.

1194 As will be explored in the breach element below, a professional can only be held liable to disclose information that they know or (under ordinary negligence standards) ought to know. Chapter 2 considered that a reasonable use of AI is predicated on a degree of knowledge regarding the characteristics of AI that were taken to render it challenging for autonomy: its general relationship to risk characteristics, its alteration of expertise, the ability to pursue goals relatively independently and a user's propensity to be influenced by it. This knowledge does not stem from a manufacturer's duty to warn. Simultaneously, developers will realistically provide descriptions of their device. Liability will more appropriately be imposed in light of this, where they actively mislead a physician about the nature of these characteristics. As this presumes positive malfeasance by the actors, however, it will not be considered as a general problem here.

of *Trustees* the court offered its seminal judgment on the duty to disclose in these terms: 'A physician violates his duty to his patient and subjects himself to liability if he withholds any facts which are necessary to form the basis of an intelligent consent by the patient to the proposed treatment'.¹¹⁹⁵ Similarly, when *Cobbs* solidified the doctrine's association with negligence, it framed the duty 'as an integral part of the physician's overall obligation to the patient'.¹¹⁹⁶ Many subsequent judgments have placed a similar emphasis on the person of the physician.¹¹⁹⁷

Although there may be specific circumstances where the finding of such a relationship is controversial, AI use does not appear to add to them. In *Hale v. Superior Court* it was made clear that the patient-physician relationship can be initiated by a doctor's examination, diagnosis or furnishing of treatment.¹¹⁹⁸ Where an ML device is used as an assistive technology, the determination of a relationship can still proceed according to such, relatively straightforward, involvement.

A further factor that must be considered here is that, in spite of relatively broad statements regarding physicians, several cases have found that not all of those who owe the patient a general duty of care, also owe a duty to obtain their informed consent. Specifically, in *Mahannah v. Hirsch* and *Townsend v. Turk* it was held that a pathologist and a radiologist respectively did not owe a duty to obtain the patient's informed consent.¹¹⁹⁹ Information does not need to be conveyed at every treatment level and it is 'the therapist',¹²⁰⁰ the one actively managing the plaintiff's care,¹²⁰¹ or the professional in direct contact with the patient that owes this obligation.¹²⁰²

This will restrict a patient's entitlement to information about AI use. Where the technology is employed by a specialist, to which the patient

1195 *Salgo v. Leland Stanford Jr. University Bd. of Trustees* (1957) 154 Cal.App.2d 560, 578.

1196 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243.

1197 For example: *Conte v. Girard Orthopaedic Surgeons Medical Group, Inc.* (2003) 107 Cal.App.4th 1260, 1266-1267.

1198 *Hale v. Superior Court* (1994) 28 Cal.App.4th 1421, 1423-1424.

1199 *Mahannah v. Hirsch* (1987) 191 Cal.App.3d 1520; *Townsend v. Turk* (1990) 218 Cal.App.3d 278.

1200 *Jamison v. Lindsay* (1980) 108 Cal.App.3d 223, 230.

1201 *Townsend v. Turk* (1990) 218 Cal.App.3d 278, 286.

1202 *Mahannah v. Hirsch* (1987) 191 Cal.App.3d 1520, 1527; *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125, 1135. See also *Quintanilla v. Dunkelman*, which highlighted the importance that 'the role of Dr. Dunkelman was more than merely that of a referring physician': *Quintanilla v. Dunkelman* (2005) 133 Cal.App.4th 95, 118-119.

or their data has been referred, there is no readily cognisable duty of disclosure to the patient. Even the obligation to disclose this information to the doctor will only arise if it has some bearing on their fulfilment of the ordinary, beneficence-oriented, standard of care.¹²⁰³

Furthermore, one must investigate the situation of clinical personnel, other than physicians, who can assume a primary caregiver role, such as nurses. That these owe a duty with a professional standard is well established.¹²⁰⁴ The case law and statute additionally suggest that this encompasses disclosure obligations in appropriate circumstances.¹²⁰⁵ With respect to the former, one can point to *Massey v. Mercy Medical Center Redding* in which it appears to have been assumed that a nurse's decision to administer a morphine injection could have constituted a breach of her informed consent obligation (even though the relevant action was in fact barred by the statute of limitations).¹²⁰⁶ Regarding the latter, it is worth noting that the legislature has, through various forms of legislation, recognised specific situations in which a nurse will be responsible for obtaining the patient's informed consent. For example, a recent amendment to the California Business and Professions Code, namely section 2746.54, lays down a number of factors with regard to which a certified nurse-midwife must obtain informed consent.¹²⁰⁷

A good argument can therefore be made that a duty encompassing informed consent requirements has been, or can consistently be, imposed on non-physician direct caregivers. This is a welcome addition in those

1203 *Townsend v. Turk* (1990) 218 Cal.App.3d 278, 287.

1204 *Fraijo v. Hartland Hospital* (1979) 99 Cal.App.3d 331, 341-342.

1205 For an example of the situation where informed consent obligations were not properly imposed on caregivers other than the physician, see: *Ermoian v. Desert Hospital* (2007) 152 Cal.App.4th 475, 510-513.

1206 *Massey v. Mercy Medical Center Redding* (2009) 180 Cal.App.4th 690, 698-699. See also *Anderson v. PIH Health Hospital-Whittier*, which implies that a duty to obtain informed consent could be imposed on nurses and nonphysician personnel where they were charged with the task of obtaining the informed consent of the patient: *Anderson v. PIH Health Hospital-Whittier* (Cal. Ct. App., Apr. 8, 2022, No. B308407) [unpublished opinion].

1207 Under California Business and Professions Code section 2746.54, subdivision (a)(1) and (5), informed consent must be obtained regarding the fact that 'the certified nurse-midwife is not supervised by a physician and surgeon' and that '[t]here are conditions that are outside of the scope of practice of a certified nurse-midwife that will result in a referral for a consultation from, or transfer of care to, a physician and surgeon'.

situations involving AI where the status or expertise of primary caregivers has been lessened or altered.

2. Healthcare institutions

There are different mechanisms for holding an institutional healthcare provider liable for negligence under United States and Californian tort law. This encompasses both a theory of vicarious liability, for the negligent actions of parties suitably associated with an institution, and corporate liability where the corporation owes a direct duty of care to the patient.

Vicarious liability attributes liability to an institution for another's breach of their duty. As such, it does not generate a new obligation, but merely identifies an additional party that a plaintiff can hold responsible for their injury in medical malpractice.¹²⁰⁸ There is ample precedent that a hospital can be liable for the negligence of their agent or their ostensible agent in this manner.¹²⁰⁹ Given that an informed consent action is merely a species of medical malpractice, a healthcare institution may also be vicariously liable for shortcomings related to informed consent.¹²¹⁰ Yet, as this duty does not provide additional protection, but only reinforces the existent one, it will not be examined further here.

The other basis for liability is provided by the concept of corporate institutional liability. This refers to an institution's 'violation of a duty—as a corporation—owed directly to the patient which resulted in injury'.¹²¹¹ Various duties have been imposed on healthcare institutions (specifically hospitals) under this head, whereby a plaintiff need not refer to the specific actions of any agent to establish their claim.¹²¹² A hospital has, for instance, been held to owe a duty to exercise reasonable care in 'screening the competency of its medical staff to insure the adequacy of medical care rendered to patients at its facility'.¹²¹³ However, the nature of this duty is clearly restricted and, also in light of the previous section's finding that informed

1208 *Ermoian v. Desert Hospital* (2007) 152 Cal.App.4th 475, 501-502.

1209 *Elam v. College Park Hospital* (1982) 132 Cal.App.3d 332, 337; *Ermoian v. Desert Hospital* (2007) 152 Cal.App.4th 475, 502-503.

1210 See the extensive analysis in *Ermoian* on this basis (even though the Court of Appeal ultimately upheld a finding that there had been no negligence): *Ermoian v. Desert Hospital* (2007) 152 Cal.App.4th 475, 514-516.

1211 *Elam v. College Park Hospital* (1982) 132 Cal.App.3d 332, 338, fn. 5.

1212 *Murillo v. Good Samaritan Hospital* (1979) 99 Cal.App.3d 50, 55.

1213 *Elam v. College Park Hospital* (1982) 132 Cal.App.3d 332, 338-347.

consent duties are applicable only to a subset of medical professionals, one must ask whether a hospital can owe such a duty to the patient directly.

Referring back to the analysis in Chapter 2, it must be recalled that such a direct form of liability is expected to assume added relevance to our analysis of AI. Applications of the technology will allow for circumstances where it is only an institution (rather than any individual agent) who can provide the user with information on the technology's use and nature. Specifically, instances were discussed where individuals engage with AI that triage, and partially determine, a clinical process before they enter a specific professional-patient relationship.

Having legally mandated informational duties in these, relatively unorthodox, scenarios is arguably necessary to address AI's unique autonomy challenges (especially its relative independence) and constitutes an important guarantor of patient self-determination. Unfortunately, Californian courts have declined to impose such a duty and have thereby situated themselves well within a wider trend running through the American case law on this matter.

The first step in this analysis, is to reiterate the distinction drawn above between mere duties to warn of products and informational obligations that can secure the patient's informed consent. Although a healthcare provider may owe the former duties,¹²¹⁴ they are narrowly delineated and are of little assistance in addressing the problems associated with AI. What is of concern to us, is whether a wider, more malleable informational obligation can be imposed on these institutions.

The leading precedent in California on this aspect is *Walker v. Sonora Regional Medical Center*. This case reiterated that a hospital was under a corporate duty to protect the patient from harm, but it declined to hold that this also encompassed an obligation to provide the patient with information (laboratory test results) relevant to her clinical decision.¹²¹⁵ Admittedly, this case was subject to special circumstances, given that federal regulation and state legislation expressly limited the persons to whom this information could be disclosed.¹²¹⁶ Nevertheless, in reaching its conclusion, the *Walker* court referred to policy-based argumentation that is in

1214 *Bigler-Engler v. Breg, Inc.* (2017) 7 Cal.App.5th 276, 317-318.

1215 *Walker v. Sonora Regional Medical Center* (2012) 202 Cal.App.4th 948, 959-963. See also: *California Jurisprudence* (Third Edition 2022) Healing Arts and Institutions § 416.

1216 *ibid* 960-962.

line with the approach of other U.S. states, which have similarly rejected informed consent claims against hospitals and healthcare institutions.¹²¹⁷ In particular, the court referred to the primary importance of the physician-patient relationship, where disclosure of the relevant information could be presumed, and to the undesirability of the hospital interposing itself into this relationship.¹²¹⁸

Furthermore, it stated ‘that it is physicians or other licensed medical practitioners, not hospitals as corporate entities, who actually practice medicine’ and render advice to the patient on the basis of their individual circumstances.¹²¹⁹ In other words, the court hinted that one of the bases for imposing informed consent obligations – a relationship of personal dependency with an imbalance of knowledge – was absent.¹²²⁰

Lastly, it appears that the court further characterised the hospital’s role as a supplementary, advisory one. It was analogous to the role of a consultant, as in the case of *Mahannah* that was discussed above.¹²²¹ As Gatter has outlined, this rejection of hospital liability for deficiencies in informed consent is almost universal and the kinds of justifications advanced in *Walker* could, by now, be said to represent a traditional policy within U.S. common law.¹²²²

At the same time, it may be said that the ruling provided scope for development. The Court of Appeal was keen to limit its findings to the specific circumstances before it.¹²²³ Moreover, it paid lip service to the

1217 For an overview of these arguments and their widespread acceptance, see: Gatter, ‘The Mysterious Survival of the Policy against Informed Consent Liability for Hospitals’ (2006) 81(4) *Notre Dame Law Review* p. 1203.

1218 *Walker v. Sonora Regional Medical Center* (2012) 202 Cal.App.4th 948, 961-963.

1219 *ibid* 965, fn. 19.

1220 This basis will be discussed further below in relation to the informed consent standard.

1221 *Walker v. Sonora Regional Medical Center* (2012) 202 Cal.App.4th 948, 961, fn. 13.

1222 Note that the exceptions to this *almost* universal approach relate to vicarious liability (which has been addressed above) and a hospital’s housing of clinical research, which is inapplicable to the therapeutic use of AI: Gatter, ‘The Mysterious Survival of the Policy against Informed Consent Liability for Hospitals’ (2006) 81(4) *Notre Dame Law Review* p. 1203, 1218-1223.

1223 ‘Whatever may be the scope of a hospital’s duty under other circumstances, we must focus our inquiry on the particular facts before us’: *Walker v. Sonora Regional Medical Center* (2012) 202 Cal.App.4th 948, 963.

hospital's role in facilitating a patient's autonomy.¹²²⁴ If the relevance of the autonomy principle were embraced, then it is arguable that policy factors would favour the imposition of a duty on healthcare institutions where they directly rely on AI. In this eventuality: the hospital's role would become that of primary caregiver and thus cease to be supplementary; the relevant expertise would not be supplied by a physician, but by the selected device; and, accordingly, there would be no danger of interfering with the professional-patient relationship. The *Walker* court's emphasis on policy considerations for the purpose of identifying duties of care ought to facilitate such innovation.¹²²⁵

Consequently, a development of the law in this direction could certainly find support. However, there is no clear evidence that the courts, in California or beyond, have thus far drawn a strong association between the corporate liability of hospitals and patient autonomy. More significantly, the outlined policy justifications for refusing to impose informed consent duties on hospitals have been found wanting for some time – even before one considers the introduction of AI use.¹²²⁶ Therefore, judging by the present state of Californian law, and the firm entrenchment of the rule against allowing informed consent claims against hospitals within almost all states, it cannot be assumed that a relevant duty of care would be imposed on healthcare institutions.

3. Summation

Concluding this element of the negligence action, it is possible to focus the analysis going forward on the duties of a restricted group of medical professionals that constitute a patient's primary caregivers or points of contact. Although this group has been argued to be relatively broad, it still leaves one considerable restriction that is unconnected to autonomy. Namely, professionals who utilise AI in a secondary capacity will have no direct obligations to disclose this use. As a consequence, there will be

1224 '[T]he Hospital arguably took the single most effective measure toward achieving the desired result of having Amber receive information and counseling regarding the laboratory test': *ibid* 966-967.

1225 'The determination that a legal duty is owed in a particular set of circumstances is "only an expression of the sum total of those considerations of policy which lead the law to say that the particular plaintiff is entitled to protection."': *ibid* 958.

1226 Gatter, 'The Mysterious Survival of the Policy against Informed Consent Liability for Hospitals' (2006) 81(4) *Notre Dame Law Review* p. 1203, 1232.

situations where a qualitatively similar, or even identical, violation of an individual's autonomy occurs, constituting a breach of duty in the sense outlined below, but the patient will only be informed of one of them.

Institutions may be vicariously liable when the outlined class of individuals constitute their agents or ostensible agents, but they are not envisaged to owe informational duties that are relevant to AI's autonomy challenges at the corporate level. This generates one further potential lacuna in the common law's response to AI's autonomy challenges.

C. Breach

Under American tort law, if a duty of care arises, then a defendant must be found to have breached that duty to be liable.¹²²⁷ For the purposes of California, the California Civil Code section 1714, subdivision (a) refers to an individual evincing a 'want of ordinary care or skill in the management of his or her property or person' and case law has interpreted this as a continuation of the common law 'standard of an ordinarily prudent man under normal circumstances'.¹²²⁸

Specifically in the medical malpractice context, in *Huffman v. Lindquist*, the California Supreme Court cited previous case law for the proposition that:

The "law has never held a physician or surgeon liable for every untoward result which may occur in medical practice" (...) but it "demands only that a physician or surgeon have the degree of learning and skill ordinarily possessed by practitioners of the medical profession in the same locality and that he exercise ordinary care in applying such learning and skill to the treatment of his patient"¹²²⁹

This reference to the learning and skill ordinarily possessed by members of the profession serves to specify the standard of ordinary prudence for the circumstances of clinical practitioners, circumstances which include their specialised education and training.¹²³⁰ In other words, the standard of

1227 Dobbs, Hayden and Bublick, *Dobbs' Law of Torts: Practitioner Treatise Series* (Second Edition 2022) § 124.

1228 *Fouch v. Werner* (1929) 99 Cal.App. 557, 564-565.

1229 *Huffman v. Lindquist* (1951) 37 Cal.2d 465, 473.

1230 *Flowers v. Torrance Memorial Hospital Medical Center* (1994) 8 Cal.4th 992, 997-998.

care remains constant, but it is applied in a form that is relevant to their situation.¹²³¹

For the purposes of establishing a breach of this specified standard, it is understandable why it will usually be necessary to rely on expert testimony.¹²³² An exception is normally only appropriate where the alleged breach can be adjudged by reference to the layman's common experience.¹²³³

It should further be noted that the judgment cited above states an innovation of American common law, the so-called 'locality rule'. Hereby, the relevant assessment must be restricted to the practice within a particular region.¹²³⁴ However, the implications of this rule for the standard of care have been contested in California.¹²³⁵ It now appears that the locality of the practitioner does not definitively delimit the practice of the professional community, but constitutes only one relevant factor to be considered in the overall assessment of the standard.¹²³⁶

With this general framework in place, one can turn to the standard that is applicable under the professional's duty to disclose information, which is a

1231 See the further elaboration of the standard in *Meier v. Ross General Hospital*. Here the court approved the rule that 'when a physician chooses one of alternative accepted methods of treatment, with which other physicians disagree, and which is in fact unsuccessful, the jury may not automatically deem him negligent': *Meier v. Ross General Hospital* (1968) 69 Cal.2d 420, 434. See also: *Clemens v. Regents of University of California* (1970) 8 Cal.App.3d 1, 12-13.

1232 *Huffman v. Lindquist* (1951) 37 Cal.2d 465, 474-475.

1233 *ibid* 474-475; 'Some questions concerning medical negligence require no expertise. Technical knowledge is not requisite to conclude that complications from a simple injection (...), a surgical clamp left in the patient's body (...), or a shoulder injury from an appendectomy (...) indicate negligence. Common sense is enough to make that evaluation': *Franz v. Board of Medical Quality Assurance* (1982) 31 Cal.3d 124, 141.

1234 McDonald, *California Medical Malpractice: Law & Practice* (Revised Edition 2022) § 2:6.

1235 'California decisions have wafted considerably on the rule's meaning and significance. Even now the topic remains in a nether realm of inconsistency, confusion and double-talk': *ibid* § 2:6.

1236 'It is now well established, however, that the locality of the practitioner is merely one of the "circumstances" to be considered in evaluating the physician's adherence to the standard of care': *California Jurisprudence* (Third Edition 2022) Healing Arts and Institutions § 442. For an overview of the development of the matter see: *Rainer v. Buena Community Memorial Hosp.* (1971) 18 Cal.App.3d 240, 259-260.

form of professional negligence.¹²³⁷ It will be seen that, although the courts have been prepared to shift from the above position somewhat, they have done so only to a very limited extent. As a result, Californian negligence law maintains a relatively narrow focus on the defendant's obligations, as defined by their profession and as evidenced by expert testimony.

In the following, I first subject this standard of review to a closer analysis, ascertaining its requirements, particularly with a view to questioning the extent to which it can protect patient autonomy. On this basis I then examine classes of cases that provide analogies to AI's autonomy challenges.

1. The informed consent standard

The terminology of informed consent was established for the first time in *Salgo v. Leland Stanford Jr. University Bd. of Trustees*. Without expressly committing to a negligence analysis, it was said that liability arose with a failure to disclose 'any facts which are necessary to form the basis of an intelligent consent by the patient to the proposed treatment'.¹²³⁸ This understanding of the doctrine was to inspire the recognition of informed consent claims within the United States more generally,¹²³⁹ but it left unanswered crucial questions regarding the specific categories of information whose disclosure the law would require and how the test, allegedly deriving from the patient's needs, was to be reconciled with negligence law's orthodox orientation towards professional practice.¹²⁴⁰

i. The meaning of reasonable disclosure

Cobbs v. Grant was the first Californian case to directly confront 'the yardstick to be applied in determining the reasonableness of disclosure'

1237 *Bigler-Engler v. Breg, Inc.* (2017) 7 Cal.App.5th 276, 321-322.

1238 *Salgo v. Leland Stanford Jr. University Bd. of Trustees* (1957) 154 Cal.App.2d 560, 578.

1239 Faden, King and Beauchamp, *A History and Theory of Informed Consent* (1986) 125-129.

1240 The court did mention the element of risk, highlighting 'that in discussing the element of risk a certain amount of discretion must be employed consistent with the full disclosure of facts necessary to an informed consent': *Salgo v. Leland Stanford Jr. University Bd. of Trustees* (1957) 154 Cal.App.2d 560, 578.

and sought to provide the requisite responses.¹²⁴¹ It remains the governing precedent for the standard of care in Californian informed consent cases.

In a first significant step, the court departed from the majority of U.S. states and the orthodox approach outlined above, by finding that the determination of the information to be disclosed could not be left to medical custom. Leaning into the focus on the patient previously exhibited in *Salgo* and the federal case of *Canterbury v. Spence*,¹²⁴² it was held that vesting '[u]nlimited discretion in the physician is irreconcilable with the basic right of the patient to make the ultimate informed decision regarding the course of treatment to which he knowledgeably consents to be subjected'.¹²⁴³ As such, the patient was viewed as the primary decision-maker in the medical encounter: an important prerequisite for establishing the relevance of their practical autonomy.

In giving the opinion of the Supreme Court, Justice Monk went to commendable lengths to outline the nature of this right to informed decision making and to contextualise it for the purposes of the physician-patient relationship. As Chapter 5 argued, many of these statements are consistent with the procedural conception of autonomy that underlies the present analysis.

Not only was it recognised that a patient lacks parity of medical knowledge with their physician, but also that they possess 'an abject dependence upon and trust in [their] physician for information upon which [they rely] during the decisional process'.¹²⁴⁴ From this 'emerges a necessity, and a resultant requirement, for divulgence by the physician to his patient of all information relevant to a meaningful decisional process (...) it is the prerogative of the patient, not the physician, to determine for himself the direction in which he believes his interests lie'.¹²⁴⁵ Conclusively it was stated: '[t]he scope of the physician's communications to the patient, then, must be measured by the patient's need, and that need is whatever information is material to the decision'.¹²⁴⁶ The patient's decision-making process was to be centre stage in the determination of reasonable disclosure.

1241 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243.

1242 *Canterbury v. Spence* (D.C. Cir. 1972) 464 F.2d 772.

1243 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243.

1244 *ibid* 242.

1245 *ibid* 242-243.

1246 *ibid* 245.

ii. The operationalisation of reasonable disclosure

Given these astute observations on the patient's position during the medical encounter and given their recognised interest in the process of decision making, it would initially appear that the disclosure required under *Cobbs* should be capable of addressing some of the challenges posed by clinical ML devices. The patients are the ones who are assessing what is in their interest and they are the ones who place their trust in the medical professional on the basis of their expertise.

However, the realisation of the stated principles through operational legal standards has proven difficult and has tended to be restrictive.¹²⁴⁷ As Weisbard observed in his analysis of the American approach to informed consent (including *Cobbs*): 'the law has fallen short of its own rhetorical promises, to say nothing of its underlying philosophical premises'.¹²⁴⁸ To understand why this is so, one can begin by looking at the specific formula that Justice Mosk provided.

Alongside its more abstract elaborations, *Cobbs* proposed a two-stage test that effectively restricted the scope for a patient-based assessment. This assessment would be relied upon only to establish 'minimal disclosure' levels.¹²⁴⁹ Hereby 'a medical doctor has a duty to disclose to his patient the potential of death or serious harm, and to explain in lay terms the complications that might possibly occur'.¹²⁵⁰ A subset of available alternatives must also be disclosed (variously described as therapeutic,¹²⁵¹ feasible,¹²⁵² recommended or reasonable)¹²⁵³ together with their comparative advantages and disadvantages.¹²⁵⁴ This is the first limb of the legally defined standard that purports to satisfy the materiality test, conveying the information that has been deemed so decisionally important for the patient.

1247 The *Cobbs* court itself recognised that 'scope of the disclosure required of physicians defies simple definition': *Cobbs v. Grant* (1972) 8 Cal.3d 229, 244.

1248 Weisbard, 'Informed Consent: The Law's Uneasy Compromise with Ethical Theory' (1986) 65(4) Nebraska Law Review p. 749, 751.

1249 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 244-245.

1250 *ibid* 244-245. See also *Arato v. Avedon* (1993) 5 Cal.4th 1172, 1190-1191; *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285, 1301-1302.

1251 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243.

1252 *Contreras v. St. Luke's Hospital* (1978) 78 Cal.App.3d 919, 927-928.

1253 *Scalere v. Stenson* (1989) 211 Cal.App.3d 1446, 1450-1453.

1254 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 242; *Jamison v. Lindsay* (1980) 108 Cal.App.3d 223, 230; *Thor v. Superior Court* (1993) 5 Cal.4th 725, 735.

A second standard, *prima facie* applying to all other classes of information,¹²⁵⁵ was based on the established practice of the relevant medical community: ‘a doctor must also reveal to his patient such additional information as a skilled practitioner of good standing would provide under similar circumstances’.¹²⁵⁶ Successfully establishing this aspect of the claim can be dependent upon expert testimony.¹²⁵⁷ Through this element of the test, one can see how the Supreme Court reintroduced the traditional community standard and struck a balance between the doctrinal requirements of negligence and the imperative of the common law principle of patient autonomy.¹²⁵⁸

It must be added that, regardless of which category one is dealing with, determining the information required is generally a question for the trier of fact (i.e. for the jury)¹²⁵⁹ and the courts have further emphasised how significant the individual circumstances of a plaintiff are to informed consent claims.¹²⁶⁰ As Applebaum and others have stated: ‘[a]fter a particular kind of problem is litigated, more concrete standards may begin to evolve (...) But within very broad limits the jury is free to write on a clean slate’.¹²⁶¹ Such an approach limits the relevance of an abstracted legal assessment of

1255 ‘When a physician recommends one or more courses of treatment, the information that is “material” (and, hence, that must be disclosed in order to obtain the patient’s informed consent) falls into two categories—namely, (1) “minimal” disclosures that are always material, and (2) “additional” disclosures that might be material if “skilled practitioner[s] of good standing” would “provide” those disclosures “under similar circumstances.”’: *Flores v. Liu* (2021) 60 Cal.App.5th 278, 293.

1256 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 244-245.

1257 *Betterton v. Leichtling* (2002) 101 Cal.App.4th 749, 754-755.

1258 A statement from *Arato v. Avedon* is telling in this respect: ‘In *Cobbs v. Grant*, we not only anchored much of the doctrine of informed consent in a theory of negligence liability, but also laid down four “postulates” as the foundation on which the physician’s duty of disclosure rests’: *Arato v. Avedon* (1993) 5 Cal.4th 1172, 1182-1183. For an examination of such a hybrid analysis more generally, see: Appelbaum, Lidz and Meisel, *Informed consent: Legal theory and clinical practice* (Second Edition 2001) 51-52.

1259 *Arato v. Avedon* (1993) 5 Cal.4th 1172, 1184; *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125, 1135.

1260 ‘The extent of a physician’s duty to disclose is directly controlled by the unique situation of each patient’: *Brown v. Regents of University of California* (1984) 151 Cal.App.3d 982, 990-991.

1261 Appelbaum, Lidz and Meisel, *Informed consent: Legal theory and clinical practice* (Second Edition 2001) 49.

the forms of disclosure that a patient would require regarding the use of ML technologies.

That being said, a number of judgments have been prepared to determine more concrete legal parameters that help ascertain the disclosable features under the first limb of *Cobbs*. They have done so under an interpretation of the materiality test that represents one of two strands of jurisprudence in the U.S.¹²⁶² Namely, they have found that disclosable, material information is that which would be regarded as significant by a reasonable person in the patient's position.

This parameter can be gleaned from the decision of *Truman v. Thomas*. In considering whether certain disclosure was legally necessary, the court relied on the aforementioned quote from *Cobbs* that 'all information material to the patient's decision' must be given and then went on: '[m]aterial information is that which the physician knows or should know would be regarded as significant by a reasonable person in the patient's position when deciding to accept or reject the recommended medical procedure'.¹²⁶³ On the basis of this standard, the California Supreme Court established for the first time in the United States that a serious risk of *not* undergoing a procedure was disclosable.¹²⁶⁴ This figure of the reasonable patient has been drawn upon by subsequent courts to establish the necessity of various other forms of disclosure.¹²⁶⁵

In *Moore v. Regents of the University of California*, for example, it was determined that a physician must 'disclose personal interests unrelated to the patient's health'.¹²⁶⁶ In so holding, the court expressly referenced that, although the failure to disclose risks was often the focus of medical malpractice actions, the requirement of materiality was capable of covering more diverse situations: 'the concept of informed consent (...) is broad enough to encompass [the situation where the physician has a personal interest]'.¹²⁶⁷ Because a reasonable patient would want to know that there

1262 Weisbard, 'Informed Consent: The Law's Uneasy Compromise with Ethical Theory' (1986) 65(4) Nebraska Law Review p. 749, 759-760; King, 'The Reasonable Patient and the Healer' (2015) 50(2) Wake Forest Law Review p. 343, 350-351.

1263 *Truman v. Thomas* (1980) 27 Cal.3d 285, 291.

1264 *ibid* 290-295.

1265 *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125, 1133-1134; *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285, 1302-1305; *Jambazian v. Borden* (1994) 25 Cal.App.4th 836, 844-845; *Flores v. Liu* (2021) 60 Cal.App.5th 278, 292-293.

1266 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120, 131-132.

1267 *ibid* 129.

is a possibility of extraneous interests affecting the physician's judgment – something that is established *inter alia* by reference to previous judicial pronouncements on non-disclosure related issues – such information is material.¹²⁶⁸ One can therefore see that the materiality test was applied in a way that both focussed on analogy (previous categories of disclosure), as well as appealing to an underlying legal determination of the needs of the reasonable patient.

Californian courts have also been prepared – indeed, somewhat more prepared – to identify classes of information that are not material to a medical decision, thereby restricting the scope for permissible (legal) argumentation. *Cobbs* already referred to 'relatively minor risks' in this manner, and highlighted that a patient must not be given '[a] mini-course in medical science'.¹²⁶⁹ Information that is commonly appreciated has also been excluded from the standard of materiality.¹²⁷⁰

In addition, the seminal California Supreme Court case on this issue, *Arato v. Avedon*, highlighted the judiciary's general scepticism towards recognising specific categories of material information that would have to be disclosed as a matter of law – such a move was described as unwise.¹²⁷¹ Accordingly, the court rejected a claim that statistical life expectancy data should have been disclosed.¹²⁷² Whether such information was material was a situational judgment for the jury to make and, more generally, the non-therapeutic interests of a patient were not a suitable basis on which to determine the appropriateness of disclosure.¹²⁷³

Lastly, a restriction regarding information that must be disclosed also emerges from the aforementioned nature of the negligence action: that it does not demand perfect behaviour, but only conduct that meets a reasonable standard. In consequence, if a patient is unable to establish that a physician ought to have known a particular piece of information, then

1268 *ibid* 129-130. The court referred specifically to a case not dealing with informed consent – *Magan Medical Clinic v. California State Bd. of MedicalExaminers* – which stated that: 'a sick patient deserves to be free of any reasonable suspicion that his doctor's judgment is influenced by a profit motive': *Magan Medical Clinic v. California State Bd. of MedicalExaminers* (1967) 249 Cal.App.2d 124, 132.

1269 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 242, 244.

1270 *Truman v. Thomas* (1980) 27 Cal.3d 285, 291.

1271 *Arato v. Avedon* (1993) 5 Cal.4th 1172, 1185.

1272 *ibid* 1186-1187.

1273 *ibid* 1186-1189.

they cannot be obligated to disclose it.¹²⁷⁴ It was concluded in Chapter 2 that a reasonable practice involving ML technologies must be predicated on some knowledge of the general characteristics that have been deemed relevant to the patient's exercise of their autonomy. Without a knowledge of AI's general risk-related status, their provision of expertise, their ability to influence the professional's decision making and the kinds of purposes it is suited for, a professional cannot responsibly employ the technology. At the same time, the possession of information about these factors will be a matter of degree. Where relevant in the subsequent analysis, a potential lack of knowledge and corresponding duty will therefore be pointed out.

iii. Summation

The current standard set by Californian common law for a professional's disclosure obligation to a patient has been conceived as a balance between the demands of the negligence doctrine and the imperatives of informed patient consent. These imperatives can be interpreted as largely consistent with the procedural conception of autonomy adopted in this thesis, under which an understanding of AI's challenges has been developed.

Deeming specific information disclosable is not done through recourse to an untethered normative analysis, however. Rather, it also requires reference to the standard of an objective, reasonable patient and, in practice, one must proceed primarily *via* an intentionally cautious development of the established categories of required disclosure.

Consequently, the following should be read as an argument by analogy that explores the relevance of classes of legally recognised information to AI's novel interferences with patient autonomy. In so far as AI-relevant information can be subsumed directly under established categories, this provides the strongest basis for finding a breach of informed consent obligations. Complementing this, an argument can also be constructed by identifying the needs of the prudent person. This will further reflect, albeit more indirectly, the law's consideration of the demands of patient autonomy.

1274 *Townsend v. Turk* (1990) 218 Cal.App.3d 278, 284-285, citing: *Moore v. Preventive Medicine Medical Group, Inc.* (1986) 178 Cal.App.3d 728,739.

2. The risks of medical AI

In Chapter 3 an analysis was provided of how clinical AI can impact the risk profile of an intervention. Two arguments were advanced: an analysis of certain AI that will impact the specific type of risks a patient is exposed to in their treatment and a claim that understanding certain general characteristics of AI will allow the patient to undertake an appropriate risk calculus, without being overloaded with information. The professional would interfere with the positive dimension of the patient's practical autonomy if they did not disclose such information and foster the requisite understanding.

In fitting these factors into the outlined Californian framework, one can begin by noting that the courts have linked inadequate risk disclosure with a diminishment of positive, practical autonomy. As was argued in Chapter 5, *Cobbs* and its progeny themselves have recognised this connection. It was also seen above that they have emphasised its importance by mandating the disclosure of material, serious risks. Yet, as this indicates, informed consent obligations are not breached by just any failure to disclose information about risks. The courts have carved out a relevant subset, the conditions for which will be applied to AI in this section. In line with our perception of AI's challenges, we first consider AI's specific risks and then their more abstract risk-related status.

i. Specific risks

The first requirement is that a given risk must be established by reference to medical expertise.¹²⁷⁵ In and of itself this appears uncontroversial. Without proving the existence of a risk, a patient cannot complain that the risk materialised or that their decision making was impacted by a failure to disclose it. With regard to AI, the existence of specific risks will have to be proven and this can be problematic due to the uncertainty surrounding, and difficulty testing, even approved AI. As referred to above, if a patient is unable to establish that a physician ought to have known of a risk, then they

1275 *Jambazian v. Borden* (1994) 25 Cal.App.4th 836, 849-850; *Betterton v. Leichtling* (2002) 101 Cal.App.4th 749, 756.

also need not disclose it.¹²⁷⁶ This constitutes a real problem for demanding the disclosure of specific AI-risks.

Assuming that the plaintiff provides proof of risk, however, it must then also be shown that the requisite risk information is material. That is, it would be considered significant by the reasonable person in the patient's position. For specific risks, this can be established according to relatively well-defined criteria: it is measured both by the gravity of the harm and the probability of its occurrence. For instance, in *Truman*, the undisclosed risk of cervical cancer was of a very low probability, but the potential harm (death after a failure of early detection) was of the gravest kind.¹²⁷⁷

Certain AI risks will be significant in this sense. As previous examples demonstrate, AI will be used in critical functions where there is a potential for very serious harm, such as for the identification of a brain haemorrhages. Moreover, initial high values for accuracy (taken as a rough stand-in for the probability of harm eventuating) were also shown to often require downward revision for various reasons. Depending on the surrounding circumstances and the specific uses, there will thus sometimes be a meaningful probability of this harm eventuating. In such cases, AI do not pose minor, low probability risks, but rather material risks.

However, since an ML device will usually impact the existing material risks of a procedure or decision in which it is integrated, one must further enquire whether its contribution is disclosable independently from the information provided about risks more generally. To be a suitable subject of the disclosure doctrine in these circumstances, a material risk must be of a type that distinguishes it from the overall risk assessment.

Morgenroth v. Pacific Medical Center, Inc. illustrates the issue. The court held that the risk of a stroke during a necessary diagnostic procedure was not different to the risks of death and serious harm that had already been disclosed to the patient.¹²⁷⁸ This meant that mentioning only the general and not the specific risk was not a breach of the defendant's obligation, such disclosure met *Cobbs'* materiality test.¹²⁷⁹ In comparison, *Warren v. Schecter* reached the opposite conclusion upon its facts. The doctor did

1276 *Townsend v. Turk* (1990) 218 Cal.App.3d 278, 284-285, citing: *Moore v. Preventive Medicine Medical Group, Inc.* (1986) 178 Cal.App.3d 728, 739.

1277 'Although the probability that Mrs. Truman had cervical cancer was low, Dr. Thomas knew that the potential harm of failing to detect the disease at an early stage was death': *Truman v. Thomas* (1980) 27 Cal.3d 285, 293.

1278 *Morgenroth v. Pacific Medical Center, Inc.* (1976) 54 Cal.App.3d 521, 534.

1279 *ibid* 534.

advise the patient of serious risks (the eventuation of which called for a second surgery) and of a risk of death from gastric surgery, but not of the risk of severe osteoporosis.¹²⁸⁰ The latter was regarded as a distinct type, warranting separate disclosure.¹²⁸¹

This presents a particular problem for medical AI because, unless they themselves represent a separate procedure (an issue that is considered below), AI are unlikely to generate a new kind of serious harm – as the osteoporosis in *Warren* unquestionably did. For example, while the use of the Acumen Hypotension Prediction Index Software may lead to different kinds of errors and varying degrees of uncertainty in the prediction of a hypotensive event, it does not itself generate the risk of such an event or alter its gravity. Under this approach then, the specific risk changes associated with AI that are integrated into a wider process (i.e. the great majority of cases) will not be material.

Nor is it easy to argue for an alteration of this approach from the perspective of procedural autonomy. Focussing on material risks, as well as classifying risks of a similar kind and magnitude together, serve a useful purpose in the fortification of the patient's positive autonomy: it prevents the patient from being overwhelmed by information in their decision making. As was stated in *Cobbs*: 'the patient's interest in information does not extend to a lengthy polysyllabic discourse on all possible complications'.¹²⁸² Listing the specific effects of AI technology on the risks and benefits of a complex procedure would arguably run afoul of these objectives.

Ultimately, the limited disclosure obligation under Californian law reinforces the view expressed in Chapter 3, that requiring information about the particular risks involved in the use of an AI device is of limited utility in addressing the underlying, more nuanced challenges to the patient's decisional and practical autonomy. Such an approach would not even mandate the disclosure of AI use, its separate characteristics or its dangers. Instead, such matters would be obscured under the head of an overall calculus of risks.

1280 *Warren v. Schecter* (1997) 57 Cal.App.4th 1189, 1195-1196.

1281 *ibid* 1202.

1282 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 244.

ii. Risk-relevant status

A wholly different approach is conceivable if AI/ML itself constitutes a material consideration because of its abstracted relation to the risk analysis. The distinct nature of the technology would be moved into the foreground, without overloading the patient with information, and their decision making would be facilitated. Such an approach is in line with *Cobbs*' central tenets and, as will be explored here, it can be further supported by Californian courts' treatment of the innovative nature of a procedure as a material, risk-related piece of information.

One can begin by noting that the burden of proving a status seems considerably lighter than that of proving the existence of a specific risk. If one focuses on an abstract classification of AI for disclosure purposes, then it is easier for the plaintiff to demonstrate (1) that the technology that was applied to them was of a certain kind, and (2) that this kind is associated with features that are relevant to the patient's risk assessment.

An analogy for this analysis can be found in *Clemens v. Regents of University of California*. Here the plaintiff claimed that they had been subjected to a procedure with an experimental status and that this bore special significance for disclosure purposes.¹²⁸³ Given that this was the first claim of its kind, the plaintiff may be forgiven for failing to show that the procedure they underwent was in fact experimental. While not denying the significance of the status, and admitting that the procedure was new, the court held that the testimony and practice of the defendants (who had been employing it for over two years) did not bear out the classification.¹²⁸⁴

In the subsequent case of *Trantafello v. Medical Center of Tarzana* it was then possible to build upon this limited recognition. The Court of Appeal considered that the plaintiff's evidence had successfully raised a triable issue of fact regarding the claim that the use of a certain substance in their surgery was innovative.¹²⁸⁵ As will be discussed below, both courts appear to have been satisfied that, if the plaintiffs can make good their claim as to the experimental or innovative status of an intervention, this information ought to have been disclosed to the patient.

For AI it is envisioned that a similar approach can be consistently adopted. A patient should be able to prove that a technology utilised in their

1283 *Clemens v. Regents of University of California* (1970) 8 Cal.App.3d 1, 9.

1284 *ibid* 7.

1285 *Trantafello v. Medical Center of Tarzana* (1986) 182 Cal.App.3d 315, 320.

care possessed the relevant capabilities, most likely by employing a form of ML (as per the definition of Chapter 2). This will require expert evidence. Establishing the materiality of these AI/ML capabilities requires the second, separate evaluation – drawing on their relation to the risk analysis. That the possession of a risk-related status is an indicator of materiality, mandating disclosure, can be derived most straightforwardly from a deeper look at the two examined cases.

In *Clemens* the court was prepared to accept the validity of a jury instruction that ‘a physician seeking consent to a “new or experimental” procedure should inform the patient that it is new or experimental when seeking to consent to it’.¹²⁸⁶ This disclosure was deemed necessary, alongside the disclosure of the more specific risk that was in fact disclosed and which eventuated.¹²⁸⁷

Furthermore, although *Clemens* preceded *Cobbs*, its continued relevance can be gleaned from *Trantafello*, which stated expressly that the professional’s ‘duty to inform plaintiff of the alleged innovative nature of the treatment [arises] from the general rules stated in *Cobbs v. Grant* (...) and *Clemens v. Regents of University of California*’.¹²⁸⁸ As in *Clemens*, the *Trantafello* court determined that this disclosure was required separately from information regarding the available options and dangers involved.¹²⁸⁹ And, in line with the above analysis, the significance of this information appears to have been deduced primarily from its suspected impact on the probability of causing the patient physical harm (i.e. its risk-related status).¹²⁹⁰

As such, I do not agree with McDonald, who argues that the innovative nature of a procedure shapes disclosure requirements by heightening existing risk disclosure obligations associated with a procedure.¹²⁹¹ Both *Clemens* and *Trantafello* indicate that other obligations are not so much

1286 *Clemens v. Regents of University of California* (1970) 8 Cal.App.3d 1, 9.

1287 *ibid* 8-9.

1288 *Trantafello v. Medical Center of Tarzana* (1986) 182 Cal.App.3d 315, 320, fn. 2.

1289 ‘Dr. Richland was required to advise plaintiff prior to obtaining plaintiff’s consent, of the innovative nature of the operation and the available options and dangers involved’: *ibid* 320.

1290 ‘The theory of plaintiff’s case is that the generally accepted practice in disk surgery is to implant a bone graft for this purpose; that the use of methyl methacrylate was an innovative procedure not generally accepted in the United States because of a high probability it will not properly fuse or heal to the bone and which has a high incidence of pseudo arthrosis’: *ibid* 319.

1291 McDonald, *California Medical Malpractice: Law & Practice* (Revised Edition 2022) § 2:11.

heightened, as an additional category of material information is established. Once a plaintiff proves that a device or procedure with general risk-related qualities was involved in their care, they have a right to this information, even if specific risk disclosure obligations are not applicable or have been fulfilled. This is what the patient's right to self-determination demands, both as outlined by *Cobbs* in the abstract and in these two cases in particular.

Another case that deserves attention in this respect, is *Daum v. Spinecare Medical Group, Inc.* It was already seen in our earlier analysis that the use of an experimental device was not deemed capable of constituting a battery. Yet this did not prevent the Court of Appeal from stating that a breach of statutory informed consent requirements – which were said to represent ‘procedural additions to the general common law requirements’ – could be found where an individual was not provided with the written consent forms for the use of an investigational device.¹²⁹²

Although an explicit materiality analysis was rendered unnecessary by the statutory background of the claim,¹²⁹³ the court made some useful comments during its causation analysis that, as will be seen below, similarly incorporates reference to the reasonable plaintiff.¹²⁹⁴ Namely, it was found that an investigational or experimental device possessed several risk-related characteristics that can be significant to the prudent patient in Mr. Daum's position.¹²⁹⁵ These factors included: a lack of evidence regarding a device, its unapproved regulatory status, the availability of alternatives, and the plaintiff's preference for a conservative approach (literally: ‘Mr. Daum claimed he would have been unwilling to be the subject of an experiment with an unproved device (...) he knew from this experience that “the first things that you will build always have some problems with them.”’).¹²⁹⁶ Note that none of these factors relate directly to specific complications, but

1292 *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285, 1307-1309. Elsewhere the device was referred to as experimental: *ibid* 1296.

1293 ‘The disclosure at issue in *Daum* was required by statute and regulation, and in those circumstances we held it was reversible error to require the jury to consider only expert testimony in deciding whether the defendants had complied with their duty of disclosure’: *Betterton v. Leichtling* (2002) 101 Cal.App.4th 749, 755.

1294 For an analysis of the two roles of the reasonable patient see: King, ‘The Reasonable Patient and the Healer’ (2015) 50(2) Wake Forest Law Review p. 343, 349-350.

1295 Generally the device was described as investigational, but experimental was treated as a synonym: *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285, 1296.

1296 *ibid* 1311-1312.

that they nevertheless raised a question of fact whether a prudent patient would have declined the relevant surgery.¹²⁹⁷ Transferring such insights to the materiality analysis, this finding arguably suggests that a reasonably prudent patient would also consider the investigational status of the device (with all the associated characteristics) material.

The reasoning of the innovative and experimental cases provides several grounds for comparison to AI. In Chapter 3, a covert employment of ML methods was found to be problematic for a patient's theoretical autonomy because: they have a more tenuous connection to the state of medical knowledge, they are subject to an irreducible risk of relying on confounding variables, there are still unsolved difficulties in their scientific validation according to established methods and there is an inherent opacity in their operation that may obscure a more specific assessment of risks. Furthermore, it is true that, unlike the device in *Daum*, it is not suggested that use of an AI automatically constitutes an investigation or experimentation in and of itself (either under the state or federal statutes). Nevertheless, it was commented in Chapter 3 that one objective of online ML functioning will ordinarily be to improve its own performance. In such cases the analogy with experimental or investigational procedures is heightened by the introduction of this research-related purpose.

On these bases it is argued that Californian courts should be prepared to extend the reasoning of *Cobbs*, *Clemens*, *Trantafello* and *Daum* to AI technology and find that the AI/ML status is in itself material information. This would give due deference to existing doctrine, in the form of *Cobbs*' two-stage analysis, while also appropriately recognising the autonomy principle's role in shaping this area of the law. As far as the breach element of the negligence action is concerned, such a limited extension is capable of addressing one of AI's unique autonomy challenges.

3. Alterations to expertise

As argued in Chapter 3, unknown changes from human to AI expertise arguably present a problem for patient autonomy, hindering or altogether preventing them from exercising their theoretical rationality appropriately. This raises the question whether, in certain scenarios, it is the obligation of the physician to affirmatively disclose basic information about their status

1297 *ibid* 1312.

and identity, about their relative expertise and about their relation to other professionals involved in the patient's care.

Different strands of Californian case law suggest that the non-disclosure of the (lack of) medical expertise of a caregiving professional constitutes a breach of their duty to obtain the patient's informed consent. This section offers an interpretation of these strands that is consistent with the principle of procedural autonomy, providing it with coherence, and argues that it should be affirmed on this basis. It is further maintained that a logical extension to medical AI will require professionals to disclose those situations in which the technology takes over a significant aspect of their expert judgment.

The first authority that is directly in point, is *Davis v. Physician Assistant Bd.* A physician assistant sought relief against a disciplinary action that revoked their licence.¹²⁹⁸ *Inter alia* this decision had been based on the assistant's negligent failures to obtain patients' informed consent. The patients had not been advised that only the assistant – rather than a medical doctor – would be the one performing their procedures.¹²⁹⁹

Seeking a basis on which to make this novel determination, the Court of Appeal began by highlighting that legally mandated disclosure was not limited to information concerning risks and alternatives.¹³⁰⁰ Echoing *Moore v. Regents of the University of California*, the court then stated that the concept of informed consent is 'broad enough to include information about whether the person who is going to perform a patient's surgery is a doctor or not'.¹³⁰¹ By applying the reasonable patient test of materiality, the *Davis* court determined that: '[c]learly identifying the practitioner who would perform surgery, making clear whether the person performing the procedure is a physician assistant and not a doctor, and making clear whether or not a physician would be involved at all are matters relevant to informed consent'.¹³⁰²

In this manner, the court expressly noted the relevance of status and, more concretely, the non-doctor status of the person providing primary care. *Davis'* injunction to inform the patient of (any) physician involvement opens the door for argumentation that there is a broader right to know

1298 *Davis v. Physician Assistant Bd.* (2021) 66 Cal.App.5th 227, 231.

1299 *ibid* 275.

1300 *ibid* 277.

1301 *ibid* 278.

1302 *ibid* 278.

about the actors involved in one's care, their levels of expertise and their relationship to one another.

However, to be contrasted with this finding is *dicta* from *Quintanilla v. Dunkelman*. Here a breach of professionals' informed consent obligations was found, but primarily on the basis that the nature of the patient's choices was not adequately explained to her.¹³⁰³ The court did not judge on several issues surrounding the characteristics of the treating surgeon: the plaintiff claimed that she did not know that the surgeons treating her were not gynaecologists;¹³⁰⁴ the court accepted that she was not aware of the identity of the person who performed the procedure, since she believed that it would be the physician with whom she enjoyed a pre-existing relationship;¹³⁰⁵ and the court commented upon a disparity in the expertise between these two individuals. Whereas the person who she believed to be operating was 'a board-certified general surgeon who performs gynecological surgery', the actual surgeon was 'a general surgeon who spent four to six months in residency in gynecology'.¹³⁰⁶

One way to interpret the court's failure to problematise these aspects, is as an implicit judgment that the non-disclosure of the physician's personal characteristics does not violate informed consent obligations. This finding could be reconciled with *Davis* on the basis that there was no fundamental issue concerning the surgeon's status in *Quintanilla*: the individuals were clearly qualified and licensed to perform the actions that they did.¹³⁰⁷ But this would be an expansive understanding of a limited discussion. A more straightforward explanation can be found in the court's preference for limiting its finding to orthodox grounds, which were readily available. Namely, the aforementioned failure to disclose the nature of the procedure.

In addition, it is illustrated by the facts of *Quintanilla* that it would be anomalous from the perspective of patient autonomy to regard status, but

1303 *Quintanilla v. Dunkelman* (2005) 133 Cal.App.4th 95, 113-115. See also the discussion *infra*.

1304 *ibid* 101, 107.

1305 *ibid* 118-119, 102-103.

1306 *ibid* 101-102.

1307 What was left unclear in *Davis*, however, is in how far the professional status and the professional regulation of the defendant – that certain acts constituted unlicensed, unauthorised, unsupervised practice, etc. – were relevant. That the licensed or unlicensed nature of conduct was not mentioned in the informed consent discussion arguably indicates that this was not understood as a defining characteristic of the material disclosure: *Davis v. Physician Assistant Bd.* (2021) 66 Cal.App.5th 227, 276-279.

not expertise and identity, as actionable non-disclosures. The patient had very clearly placed her trust in one particular individual with whom she had several interactions and whom she understood to possess a requisite degree of skill to carry out a sensitive procedure on her. Surely this was material information from the perspective of a reasonable patient making their medical decision. Under an approach that focuses on procedural autonomy, a patient should simply be required to prove that the human expertise applied to their care was lessened or morphed to such an extent that it would have assumed practical significance in their decision making.

That *Quintanilla* cannot be taken as instructive on the professional's duty to advise a patient of their relative expertise, is also supported by *Moore v. Preventive Medicine Medical Group, Inc.* The Court of Appeal recognised that a patient must be told the risk of being examined only by a non-expert: 'material information included the risk to Moore if he was not examined by the specialist'.¹³⁰⁸ The lesser capability of the referring physician required this risk to be impressed upon the patient at the point of their interaction.¹³⁰⁹ Disparities in the expertise of one's professional is therefore something that can be material to the reasonable patient's decision making in a specific context.¹³¹⁰

However, it is true that the applicability of *Moore* has been delineated much more narrowly than suggested by *Davis*. Subsequent case law has affirmed that a duty to advise is only triggered where a less qualified professional actively refers the patient for expert assessment.¹³¹¹ A duty to discuss the disparities of expertise and the associated risks will therefore not arise without more. For example, where the physician could recommend an expert opinion, although they are not required to do so by ordinary negligence principles,¹³¹² but they do not think it necessary on the facts of the case.¹³¹³

1308 *Moore v. Preventive Medicine Medical Group, Inc.* (1986) 178 Cal.App.3d 728, 738-739.

1309 *ibid* 738-739.

1310 For a similar interpretation, see: Terrion, 'Informed Choice: Physicians' Duty to Disclose Nonreadily Available Alternatives' (1993) 43(2) Case Western Reserve Law Review p. 491, 511-512.

1311 *Scalere v. Stenson* (1989) 211 Cal.App.3d 1446, 1450.

1312 *ibid* 1453.

1313 See the next section for an analysis of the requirements for the disclosure of alternatives.

Moore's duty is also construed narrowly in that the difference in expertise must be relevant to a specific procedure and pose a corresponding risk. What matters are not the general characteristics of the professional and the expert, but rather the specific danger that is associated with a failure to follow through on a particular course of action that has been recommended. In *Moore* this course of action was the physician's statement that one could not know a mole's nature for sure, unless the patient saw a specialist and 'got it removed or studied microscopically'.¹³¹⁴

Bringing AI-induced changes directly under *Moore* would prove difficult as a result. One aim of leveraging ML technologies is to supplement a less qualified professional's expertise with specialist knowledge. As such, after AI analysis, there will not often be a duty under ordinary negligence to refer a patient to an 'actual' human expert. Nor will it always be easy to frame the information that is relevant to ML-generated expertise as a risk that is associated with the non-acceptance of a procedure. ML assistance may constitute a broader form of assistance, which is not linked to one specific recommendation, as in the case of AI-Pathway Companion Prostate Cancer.

Consequently, the line of case law originating with *Moore v. Preventive Medicine Medical Group, Inc.* is one important judicial recognition of the significance that differing levels of professional expertise can have on the patient's understanding of their care pathway. But it does not provide obligations to disclose such information in its own right.

In light of these conflicting signals, an interpretation of *Davis* that imposes an obligation to disclose the identity of relevant actors in the patient's care, their status and their levels of expertise arguably offers the most coherent interpretation of these different strands of case law. It would vindicate the importance that has been attached to differing levels of professional expertise and it would reflect the significance that a patient is likely to attach to the identity and the capabilities of their carer.

Such an interpretation could then require the disclosure of shifts of expertise precipitated by AI. Recall in this respect an ML device like IDx-DR. Its use allowed a less qualified professional to assume a task that was previously reserved for a more skilled practitioner. *Davis* supports the proposition that a patient has an interest in knowing the expertise-related differences in the actors involved in their care. This is especially true where one actor, such as the physician, is expected to be the central figure in

1314 *Moore v. Preventive Medicine Medical Group, Inc.* (1986) 178 Cal.App.3d 728, 734.

their care. If it is not obvious that an individual with a lesser status and/or expertise is using such a device, then it is arguable that discrepancies in expertise, and the role of the device, constitute material information.

In addition, a patient could strengthen their argument by demonstrating that the discrepancy has the requisite degree of significance. This could emerge for example – as in *Moore* and *Quintanilla* – from the high-stakes nature of the decision. It could also be inferred where the patient can demonstrate that there was a decrease in the overall level of expertise brought to bear on the decision – as in *Davis*. For example, that an AI did not fully substitute the capabilities of the expected human actor.

Things are arguably different where the primary carer's expertise remains unchanged. Where an AI like Mia (Mammography Intelligent Assessment) is involved, whereby the AI merely provides a second opinion to a human radiologist, there would arguably be no disclosable change in the expertise of relevant actors. Any deficiencies in the AI would be compensated for by the involvement of the human professional.

Such a result constitutes a legitimate, plausible development of existing case law. Nevertheless, the incomplete nature of this case law, as well as the conflicting indications that are provided by it, limit the strength of the arguments that can be made in this respect. Requiring the disclosure of shifts in AI-human expertise must therefore be seen as an available argument of uncertain strength under Californian common law.

4. Information concerning the choice of goals

AI's independent pursuit of objectives was argued to be problematic on two accounts in Chapter 3. A very strong challenge emerged from AI that could partially determine an aspect of a clinical decision. A related interference stemmed from AI's ability to impact human decision making through a lesser form of influence, through nudging. Patients are thereby subjected to latent pressures to pursue non-personalised objectives.

The manner in which the Californian informed consent doctrine secures a patient's control over their goals of treatment is by making available to them information concerning: the choices available in their care, including the purposes of recommended options, and by requiring the relevant professionals to disclose any extraneous interest that they may have in their care. Recourse to the autonomy principle is necessary to subsume the requisite information about ML devices under these heads.

i. Understanding choices

Under California's informed consent doctrine, different options must be disclosed to the patient where they constitute material information.¹³¹⁵ Ordinarily this will involve a choice between courses of action with distinct risk profiles. For example in *Jamison v. Lindsay* the focus was on dangerous therapeutic interventions, for which 'the therapist must inform the patient of the available alternatives and the hazards involved so that the patient is able to give effective consent to the proposed treatment.'¹³¹⁶ Similarly, in the aforementioned case of *Quintanilla v. Dunkelman* it was found that a failure to discuss available choices and their dangers, including the surgical procedures actually performed on the patient, constituted a breach of duty.¹³¹⁷

ML devices will be subsumable under the existing paradigm in so far as they are deemed to possess risk-related characteristics. Opting for a diagnosis or treatment method involving such AI over another procedure, or over non-treatment, with a different risk-benefit balance, would require a discussion of both options (subject to the caveat discussed below that it is a reasonable alternative). Yet, disclosing the risk-related characteristics of AI does not provide the patient with information regarding AI's objectives or their influence on the non-risk related goals that are pursued in patient care. To consider whether the disclosure of such factors is required, one must analyse them under the wider rationale which informs a professional's duty to advise a patient of the available options.

Namely, allowing a patient to choose amongst alternatives is a recognition of the fact that they are entitled to choose in light of their own objectives. As Applebaum and others have stated in their analysis of U.S. informed consent law: 'the choice among options cannot be made on the basis of objective criteria (...) Informing the patient about alternatives permits patient to make a decision in light of his values, preferences, goals and needs'.¹³¹⁸ Where a professional makes a choice amongst options for the patient, they exert a very strong influence on the patient's decision making

1315 'A duty of reasonable disclosure of the available choices with respect to proposed therapy': *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243. See also: *Traxler v. Varady* (1993) 12 Cal.App.4th 1321, 1330-1331.

1316 *Jamison v. Lindsay* (1980) 108 Cal.App.3d 223, 230.

1317 *Quintanilla v. Dunkelman* (2005) 133 Cal.App.4th 95, 114-115.

1318 Appelbaum, Lidz and Meisel, *Informed consent: Legal theory and clinical practice* (Second Edition 2001) 59.

– a decision making that is (as it was remarked in *Cobbs*) dependent on being facilitated by the physician's greater expertise. Without the physician's disclosure of alternatives, there is a very real danger that the patient's care is disconnected from the pursuit of their medical and non-medical interests.

Multiple *dicta* testify to the fact that Californian courts have recognised this rationale. In *Mathis v. Morrissey* the court gave the example of a professional disagreement on the optimal treatment for breast cancer.¹³¹⁹ It was held that, even if the pursuit of either procedure can be supported within the relevant medical community, it was for the patient to make the final, personal decision, so that the professional recommending one procedure had to disclose the other, recognised schools of thought.¹³²⁰ In *Wilson v. Merritt* the court was also prepared to recognise the importance that a paraplegic, in selecting among treatment options, would attach to maintaining the use of their arms and shoulders in light of their dependence on them for their mobility.¹³²¹ The danger of disconnecting the patient from the pursuit of their own, individual objectives therefore clearly shapes the information that must be provided in scenarios involving the discussion of alternatives.

From the expositions in Chapter 3 it can be deduced that AI's control over objectives will not normally lead to a similar disconnect. Even if an AI possesses advanced capabilities and nudges the patient towards making a particular kind of choice among their options, *inter alia* by withholding or framing information, the oversight and assistance of a human professional is envisaged to maintain the patient's ability to direct their care. Selecting treatments involving AI will, for the most part, not amount to a pre-determined choice of one alternative (i.e. the AI-favoured one) above another.¹³²²

However, in one identified subset of cases an analogous separation between the AI's goals and a patient's personal preferences was identified. This was where an AI is used to determine an aspect of clinical decision making and possesses a relatively wide discretion to select non-personalised objectives. The examples given in this regard were: (1) an AI that may be used as a general-purpose diagnostic tool, which is capable of generating surprising insights into serious underlying conditions (2) an AI that triages the patient according to its own criteria. Such applications of

1319 *Mathis v. Morrissey* (1992) 11 Cal.App.4th 332, 343.

1320 *ibid* 343-344. Relying on *Cobbs v. Grant* (1972) 8 Cal.3d 229, 243-244.

1321 *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125, 1139.

1322 Although it must also be remembered that this is a matter of degree, partially dependent on the abilities of the intervening human professional.

ML are comparable to situations where a physician acts upon their clinical judgment to select a procedure without discussing the available options with the patient. In both cases, the pursuit of goals by an external entity binds a patient. Facilitating reasoning on these aspects requires the patient to be informed of the types of commitments that they subscribe to through the decision to use AI, commitments which they could avoid without such a subscription. In other words, the choice whether or not to use AI, and the purposes implicit in this, move to the foreground.

To substantiate the obligation to advise a patient of this information, one must turn to the more specific requirements that the courts have imposed on the disclosure of choices. In particular, a breach is generally only identified where a patient was denied a true choice, one between reasonable, medically indicated options.¹³²³ As stated in *Vandi v. Permanente Medical Group, Inc.*: ‘there is no general duty of disclosure with respect to nonrecommended procedures’.¹³²⁴

Beyond this, however, it refers also to availability. A concrete illustration of this requirement can be provided through *Spann v. Irwin Memorial Blood Centers*. Here, the Court of Appeal rejected the plaintiff’s claim that a blood centre was liable for failing to disclose the possibility of a donor reduction programme, which was desirable to mitigate the risk of disease transmission. Since such a programme did not physically exist, and there was no duty under ordinary negligence principles to offer it, there could be no duty to disclose it either.¹³²⁵

This line of authority is most relevant to AI’s role in triaging patients, because here the patient is arguably denied no choice for aligning the therapeutic procedure with their preferences at all. Where an AI is integrated into a workflow, automatically prioritising images or callers within a given system for instance, another route may not be in existence. Indeed, the pre-existing option (the default) may be a first come, first served basis, which would be equally, if not more, unreceptive to the patient’s values. As a result, if a patient is seeking medical assistance at a particular institution or from a particular specialist, then their automatic subjection to AI triaging could be understood as the only available choice.

1323 *Schiff v. Prados* provides a detailed analysis: *Schiff v. Prados* (2001) 92 Cal.App.4th 692, 701-703.

1324 *Vandi v. Permanente Medical Group, Inc.* (1992) 7 Cal.App.4th 1064, 1071.

1325 *Spann v. Irwin Memorial Blood Centers* (1995) 34 Cal.App.4th 644, 658.

This is reinforced by the use of institutional, financial and geographic constraints to shape the legal definition of availability in the U.S. and California.¹³²⁶ As discussed in Chapter 3, an ML triage tool's attraction often lies in it being the most accessible and financially sustainable option that can be offered to patients. If an institution or area uses the tool, and has good financial reasons for doing so, then this reinforces a view of the tool as non-optional. In consequence, there will not be a general requirement to make a patient aware of the possible objectives involved in such non-optional AI use.

By contrast, the information associated with AI's ability to diagnose particular conditions is of a different quality. The patient does have a true choice to make, which is most akin to the scenario considered in *Mathis*. There, disclosure of options was demanded to enable patient determination of a care pathway for a very serious decision: the treatment of breast cancer. Where an AI has the capability to diagnose severe conditions, it is expected that the doctor would have knowledge of this purpose and it represents a meaningful choice for the patient whether they wish to run the risk of confronting this information. Declining diagnostic AI use must ordinarily be seen as a feasible, realistic option that can be exercised once the patient is informed of the relevant objectives.¹³²⁷

In the latter scenario there is an even stronger basis for disclosure, because the device's purpose constitutes part of the nature of the recommended and selected procedure, rather than just being a contextualising piece of information concerning choices.¹³²⁸ By selecting an AI that can pursue certain objectives, a professional is already making a certain determination. In this regard, the appropriate analogy is to a case like *Quintanilla*, where the choices at issue were the ones actually imposed on the patient by the physician. Discussing the choices involved means discussing the nature of the procedure that is subsequently carried out by subjecting the patient to an ML device. It is not an abstract option that can be limited by reference to the practice of the medical community. In this case, patient autonomy

1326 Terrion, 'Informed Choice' (1993) 43(2) Case Western Reserve Law Review p. 491, 493-496.

1327 *Truman v. Thomas* was premised on such an approach. The patient had to be informed of the purpose of the pap smear, which was capable of detecting cervical cancer, but was then in a position to make an informed refusal: *Truman v. Thomas* (1980) 27 Cal.3d 285, 293-294.

1328 *ibid* 293-294.

demands that a meaningful indication of AI objectives and capabilities be provided to the patient.

ii. AI's lesser influence on the pursuit of objectives

The last facet of AI use that was considered for its effect on patient autonomy, was its uncertain potential to influence patients to act in accordance with pre-determined objectives. It was noted that, although it is expected that AI assistance will systematically direct patient choice, it is difficult to actually assess this influence in a given medical decision. Particularly as there will normally be a degree of mediation by an expert physician. A range of factors could determine whether such an expert could help the patient overcome AI-induced biases. The most that the informed consent doctrine can do to respond to this issue, is to convey these concerns to the patient and allow them to adapt their decision making as they see fit.

That this is a type of information which is deemed material under the law of informed consent, emerges from the Supreme Court of California's decision in *Moore v. Regents of the University of California*. Through an analysis of the nature, scope and purpose of informed consent obligations the court determined:

(1) a physician must disclose personal interests unrelated to the patient's health, whether research or economic, that may affect the physician's professional judgment; and (2) a physician's failure to disclose such interests may give rise to a cause of action for performing medical procedures without informed consent or breach of fiduciary duty.¹³²⁹

The materiality of these interests stemmed from the fact that they had a potential, relatively tenuous influence on the decision maker: 'a physician who does have a preexisting research interest might, *consciously or unconsciously*, take that into consideration in recommending the procedure (...) the physician's extraneous motivation *may affect* his judgment and is, thus, material to the patient's consent'.¹³³⁰ In a similar vein, California's highest court referenced *Magan Medical Clinic v. Cal. State Bd. of Medical Examiners*, a case where it had been found that 'a sick patient deserves to be free

1329 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120, 129.

1330 *ibid* 131 (emphasis added).

of any reasonable suspicion that his doctor's judgment is influenced by a profit motive'.¹³³¹

Consequently, a professional need not in fact deviate from their patient's desires or impose extraneous goals upon them. It suffices if these goals are present and there is a readily comprehensible manner in which they *may* influence professional and patient decision making.

Moore did not require the disclosure of all such influences, though. It focussed on the presence of non-therapeutic motivating factors, those extraneous to the patient's health. Thus, the patient had stated a cause of action for a breach of informed consent specifically where their physician had an undisclosed economic interest and an undisclosed research interest in their care.¹³³² These interests were singled out because they were in opposition to the therapeutic goal of benefiting the patient.¹³³³ This opposition gave rise to the prospect of 'potentially conflicting loyalties', which the court found so odious to the reasonable patient.¹³³⁴ Moreover, as the framing of conflicting or opposing interests suggests, it was not necessary to show that the physician did not have any therapeutic purpose for their actions at all.¹³³⁵

In sum, *Moore* allows for an argument that influences capable of systematically causing a physician to prioritise the non-medical over the medical interests of a specific patient must be disclosed under the informed consent doctrine. This is so even where there is no direct indication that there was no therapeutic purpose or that non-therapeutic objectives were determinative.

For the purposes of medical AI, the precedent laid down in *Moore* identifies a kind of breach that could partially protect against its lesser influences on patient decision making. The first condition, that a readily comprehensible influence on the professional's and patient's process of decision making must be present, but not necessarily active, appears to be fulfilled for many uses of AI. As discussed in Chapters 2 and 3, although it is ultimately dependent upon the design of ML systems, it is expected that this technology will have an increased propensity to influence decision makers and nudge them towards a desired kind of decision. A patient

1331 *ibid* 129-130; *Magan Medical Clinic v. California State Bd. of Medical Examiners* (1967) 249 Cal.App.2d 124, 132.

1332 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120, 132-133.

1333 *ibid* 130-131.

1334 *ibid* 130-131.

1335 *ibid* 133.

should therefore often be in a position to demonstrate that this potential is present.

Regarding the requirement that this influence must be directed towards certain non-therapeutic objectives, it has also been argued that there will be extraneous factors that flow into a device's reasoning. Two in particular were highlighted in Chapters 2 and 3. First, there will almost inevitably be some financial considerations, at least in the sense that an AI must account for local limitations flowing from the availability of resources and from reimbursement conditions. However, it also stands to reasons that one goal of some devices will go beyond this: to make recommendations in a way that saves an institution costs.

Second, for the subset of AI that operate in an online manner, the objective of improving or maintaining their performance is expected to provide one influence directing the AI's functioning. If not research itself, this activity to improve the general experience of other users (the wider population), is straightforwardly comparable and disconnected from benefitting the individual patient. In this sense, AI may be said to introduce the kind of objectionable, non-therapeutic goals that are likely to generate conflicting interests, which shape the doctor's and patient's decision making.

It could be objected that these nudges are less likely to consciously influence the professional than their own extraneous objectives, bypassing their rationality. However, the above reasoning explicitly recognised that finding a conscious reliance on the relevant objectives was not necessary.

Similarly, while it has been argued that AI are likely to maintain an overriding therapeutic objective, it was explicitly stated in *Moore* that the absence of any therapeutic purpose for the relevant actions is not a requirement. So, where ML devices introduce identifiable financial- or improvement-related nudges to clinical decision-making, it will be possible to argue by analogy to *Moore*.

How far such identification will be possible in practice is the only remaining, and yet most substantial, issue. It is an additional hurdle for a patient dealing with an AI that they must prove the existence of certain objectives and demonstrate how they were (abstractly) capable of directing reasoning. Without easy access to the technology or insights into the processes by which a particular decision, or kind of decision, was reached – including information on how human decisionmakers tend to interact with the ML device – the patient will experience difficulties in showing that objectionable non-therapeutic objectives were present in their care and ought to have been disclosed. Likewise, it may be a considerable challenge

to demonstrate that the doctor had or ought to have had specific knowledge of certain non-therapeutic objectives pursued by the AI. A negligence complaint that AI brought extraneous interests to bear on a patient's care is therefore likely to remain a largely theoretical one. The autonomy-based argumentation would not counteract this underlying problem.

5. Summation

The Californian informed consent standard has been shaped by a careful balancing between the demands of the negligence doctrine and the imperative of protecting the patient's autonomy. In the course of this, it has delineated a number of relatively clear categories that serve as focal points for argumentation. Fitting AI into these categories, while still not straightforward, has allowed one to argue that a patient must be informed about: the risk-relevant status of AI; non-obvious substitutions of the human expertise brought to bear on their care, and; the pre-determination of certain serious choices.

D. Causation

For the purposes of causation, a plaintiff must show that the defendant's breach of duty caused their injury.¹³³⁶ Hereby it should be remembered that we are concerned with connecting the breach of informed consent obligations with the eventuation of physical injuries – as has been discussed under the damage element above.

Cobbs v. Grant also serves as the seminal authority on applying this element to cases of informed consent. California's Supreme Court was not content with a subjective test, requiring only that the individual patient would have avoided the damage (e.g. refused their consent for the operation).¹³³⁷ Instead, conscious of the evidentiary difficulties that flow from the self-serving testimony that plaintiffs may provide after the fact, the court settled on an objective test: 'what would a prudent person in the patient's position have decided if adequately informed of all significant perils'.¹³³⁸ A

1336 *Vasquez v. Residential Investments, Inc.* (2004) 118 Cal.App.4th 269, 288.

1337 *Cobbs v. Grant* (1972) 8 Cal.3d 229, 245.

1338 *ibid* 245.

subsequent addition has made this test yet more favourable to the defendant. Namely, even if the plaintiff can prove that the reasonable person would not have suffered the damage, then it is still open to the defendant to prove that the individual, particular plaintiff would have consented to the procedure and suffered the damage.¹³³⁹

At times the courts have not entirely adhered to this test, adjusting the reasonable person standard in order to zero in on the hypothetical decision of the individual patient or, at least, on the hypothetical decision of a reasonable patient with some of the same commitments as the individual patient.¹³⁴⁰ As anticipated in Chapter 5, these are arguably cases where the judiciary has been confronted with the reflective dimension of autonomy.

In *Hernandez ex rel. Telles-Hernandez v. U.S.* (a federal decision applying California law) the District Court considered the duty of a professional to disclose to a woman giving birth the relative risks and benefits of continuing with a vaginal delivery, which constituted the plaintiff's preferred and chosen option, and those of a caesarean section.¹³⁴¹ Rather than simply asking what information the reasonable person would have required, the court was prepared to consider 'Mrs. Telles-Hernandez' emphasis on prenatal care and her desire to deliver her baby without the use of medication' as evidence relevant to causation.¹³⁴² The commitments implicit in these past clinical choices were indicative of the decision she would have made, had she been given adequate disclosure.

A similar approach was implicit in the above-analysed case of *Wilson v. Merritt*. The court accommodated the patient's heightened desire for maintaining mobility in its assessment. It further acknowledged that this desire was demonstrated by the moderately successful therapy that the patient had previously undertaken, and which presented no threat to this goal.¹³⁴³

These cases hardly had the ability – nor did they purport – to overrule *Cobbs*. And the reasonable person standard of causation is arguably too entrenched to provide an opportunity for principle-based developments.

1339 *Truman v. Thomas* (1980) 27 Cal.3d 285, 294, fn. 5; *Warren v. Schecter* (1997) 57 Cal.App.4th 118, 1206; *Flores v. Liu* (2021) 60 Cal.App.5th 278, 297-298.

1340 In addition to the cases outlined below, see: *Morgenroth v. Pacific Medical Center, Inc.* (1976) 54 Cal.App.3d 521, 534-535.

1341 *Hernandez ex rel. Telles-Hernandez v. U.S.* (N.D. Cal. 2009) 665 F.Supp.2d 1064, 1078-1079.

1342 *ibid* 1078-1079.

1343 *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125, 1128-1129, 1139.

However, *Hernandez* and *Wilson* do suggest that the courts are prepared to utilise a limited discretion to specify the reasonable person standard in a way that is more amenable to the reflective dimension of autonomy and, consequently, to the protection of the deeply held commitments of the individual.

For AI's autonomy challenges this means that, on top of establishing the materiality of the outlined information to the decision making of the prudent patient, a plaintiff will have to demonstrate that this information would also have been of such significance to an abstracted individual that they would have altered their decision and avoided the damage. Ordinarily this would mean showing that there are weighty, rational grounds for refusing a procedure involving an AI or for preferring an alternative. If a patient can additionally demonstrate that they have coherent, long held or robust commitments relevant to AI (for instance with regard to an objective that the AI has overruled), then this may lead to an interpretation of the reasonable person standard that accommodates this commitment.

Much of the above discussion, in this chapter and beyond, has been dedicated to demonstrating that there are compelling grounds for a patient to be concerned about ML technology's influence on their decision making. Yet, it is undeniable that this will not always be sufficient to outweigh other demands and become the dominant concern in a high-stakes clinical situation.

The problem is well illustrated by *Spann v. Irwin Memorial Blood Centers*. Here the plaintiff was aware of the relevant risk and argued only that alternatives and other information related to the plaintiff's condition ought to have been disclosed.¹³⁴⁴ Even supposing that it was a breach of duty not to provide this information, however, the court was satisfied that, in light of the life-threatening seriousness of the patient's condition, a reasonable person in their position would not have declined the intervention.¹³⁴⁵ Implicit in such an approach is the dominant importance that the U.S. courts attach to risk disclosure above all other factors.

In light of this, there will be many situations where a concern for AI-specific factors does not weigh particularly heavily with a reasonable person requiring more-or-less acute medical care. From the adduced examples, we may point to circumstances where a professional fails to disclose alterations in human-AI expertise (particularly where there is no loss in the cumulat-

1344 *Spann v. Irwin Memorial Blood Centers* (1995) 34 Cal.App.4th 644, 658.

1345 *ibid* 657.

ive level of expertise) and to situations where an AI indirectly influences a patient's choices. Would the knowledge that an AI incorporates some objectives to improve its own performance, and systematically nudges a user, really sway a patient to reject its recommendation? The strongest cases for causation could, by contrast, be advanced with regard to risk-related information and with regard to a failure to advise the patient of their pre-determined commitments by subscribing to AI use. Here, the relation to a specific decision is clear and the importance of this kind of information for clinical choices is well-established.

To a certain extent this approach may be justifiable as a rough means of addressing only significant autonomy violations. But it has already been discussed that autonomy is not only impacted by physical injury. Moreover, the objective causation test laid down in *Cobbs* fails to address autonomy interferences that are recognised as particularly severe under a procedural theory; a focus on the reasonable patient will sometimes allow an external perspective to override the individual's deeply held subjective commitments. Therefore, causation provides another element that significantly broadens the class of cases where the patient's decisional autonomy is impaired without remedy.

E. Awarding damages

When we first began the exposition of tort law's role in the protection of a plaintiff's informed consent, it was stated that the primary means by which its obligations would be enforced is through a *post facto* award of damages – albeit such awards were frequently associated with a prophylactic effect. During our analysis of the negligence action, it was further seen that the first major hurdle for any plaintiff to overcome is demonstrating that they had suffered some form of legally cognisable injury or damage. Under Californian precedent, the only realistic category of injury that could ground an action for the analysed types of AI autonomy interferences, was argued to be physical injury. Merging these two strands, this section examines the kinds of damages that a plaintiff can hope to recover if the requirements of negligence have been made out.

Regarding the first aspect, the type of damage to be recovered, one can begin with the premise that the 'remedy in tort is compensatory in nature and damages are generally intended not to punish a negligent defendant

but to restore an injured person as nearly as possible to the position he or she would have been in had the wrong not been done'.¹³⁴⁶ Compensatory damages may be categorised either as economic or non-economic under Californian statute.¹³⁴⁷ Our focus will be on the latter, given our starting premise that the primary damage suffered to ground a successful claim is physical and that any consequent autonomy violations are associated with 'subjective, non-monetary losses including, but not limited to, pain, suffering, inconvenience, mental suffering, emotional distress, loss of society and companionship, loss of consortium, injury to reputation and humiliation'.¹³⁴⁸

In particular, in awarding damages for these losses, the law seeks to remedy a host of values deriving from the patient's autonomy, including an individual's 'capacity to participate in the physical world' and to use their 'cognitive and expressive faculty' to synchronise their 'individual existence culturally with the lives of a host of others'.¹³⁴⁹ In principle, a patient who suffers physical damage will be able to claim for all of the detriment that flows proximately therefrom.¹³⁵⁰

In the informed consent context this includes damages for the injuries and losses the plaintiff would have avoided had they been adequately informed. This was established in *Warren v. Schechter*, where the patient suffered both from complications that the professional had disclosed and from one complication that they had, in breach of their duty, failed to inform her about.¹³⁵¹ The court held that the patient could recover 'not only for the undisclosed complications, but also for the disclosed complications, because she would not have consented to either surgery had the true risk been disclosed, and therefore would not have suffered either category of complications'.¹³⁵²

However, it must also be noted that, in response to a crisis involving a dwindling availability of medical insurance in California, the legislature enacted MICRA.¹³⁵³ This statute limits the non-economic damages that an

1346 *Turpin v. Sortini* (1982) 31 Cal.3d 220, 232.

1347 California Civil Code section 1431.2, subdivision (b).

1348 *ibid* section 1431.2, subdivision (b)(2).

1349 McDonald, *California Medical Malpractice: Law & Practice* (Revised Edition 2022) § 16:7.

1350 California Civil Code section 3333.

1351 *Warren v. Schechter* (1997) 57 Cal.App.4th 1189, 1195.

1352 *ibid* 1195.

1353 *Reigelsperger v. Siller* (2007) 40 Cal.4th 574, 577-578.

individual can recover against a healthcare provider or healthcare institution based on professional negligence.¹³⁵⁴ Under amendments coming into effect January 1st 2023, the plaintiff is limited to a total recovery of either \$1,050,000 for actions not involving wrongful death,¹³⁵⁵ or \$1,500,000 for actions that do involve wrongful death.¹³⁵⁶ It has been established that these capitations apply to claims involving damage caused by breach of informed consent, as these actions are properly characterised as forms of professional negligence.¹³⁵⁷

Given the largely statutory basis for these provisions, there is also not a great deal of influence that the autonomy principle can have. But neither does it appear necessary to demand a change for the purposes of our present analysis. Although some of AI's autonomy violations have been argued to be significant, it is not obvious that they, by themselves, warrant excessive amounts of compensation. Indeed, the relevant caps are more likely to be met by the eventuation of any accompanying physical injury. So long as the system of damages is prepared to recognise and compensate an individual for the autonomy-related implications of their injury, which is arguably the case regardless of MICRA's provisions, then the autonomy principle does not demand more of this element.

III. Conclusion

In conclusion, it has been claimed that California realistically provides only one common law mechanism, through which an argumentation based on the autonomy principle can effectuate a degree of protection against AI's novel challenges: negligence.

This is not to say that this principle has had no impact on the other relevant mechanism, the battery cause of action. It was found that several elements of this tort were in fact loosened out of considerations for patient autonomy. Yet this still had little practical effect, since the case law went

1354 California Civil Code section 3333.2, subdivision (a).

1355 *ibid* section 3333.2, subdivision (b).

1356 *ibid* section 3333.2, subdivision (c). Each of these capped sums is further broken down into maximum amounts of \$350,000 and \$500,000 respectively for healthcare providers, healthcare institutions and unaffiliated defendants – those who are not owners of another specified entity, in a joint venture with such an entity or have a contractual relationship with such an entity: *ibid* section 3333.2, subdivision (j)(3).

1357 *Bigler-Engler v. Breg, Inc.* (2017) 7 Cal.App.5th 276, 321-322.

on to impose narrowly and rigidly drawn conditions on the validity of the patient's consent. Their consent would only be invalidated on the basis of substantial changes in physical characteristics of the procedure, changes in professional identity or changes that clearly undermined the therapeutic motivation of the professional. Ultimately, these requirements for a valid consent were ill-suited to address any of AI's unique challenges and they were entrenched on the basis of rule-specific considerations. A requisite amendment or extension of these rules through an appeal to principle was not realistically possible.

The negligence action, by contrast, was found to cover some kinds of cases involving AI's challenges. It was significant that the rhetoric underlying the informed consent analysis bore strong similarities to our analysis of procedural autonomy and, more concretely, precedent had already utilised this analysis to identify violations of informed consent obligations that could be developed to address the ML characteristics evaluated in Chapter 3. A patient could argue that various aspects of AI's risk-related status, their alterations of professional experience and their independent pursuit of objectives should be disclosed to them.

Overall, a plaintiff's claim would nevertheless be severely impaired by a relatively stringent adherence to negligence's doctrinal structure. Several elements are likely to preclude an effective protection of a patient's procedural autonomy, including: a limitation of claims to those involving: the eventuation of physical injury; AI use by primary care professionals – excluding institutions and secondary carers, and; a type of undisclosed information that would have been so significant as to cause a reasonable patient to avoid the injury.

This would result in a state of affairs where any protection against AI's interferences in a patient's procedural autonomy is, while not impossible, incidental at best. If a patient does succeed, then the award of compensatory damages under this claim will, additionally, be capped. Yet this is in line with general rules on the recovery of non-economic losses and does not, by itself, appear inadequate for the kinds of violations under discussion.

In comparison, a prophylactic effect – in the sense that professionals will seek to inform patients of AI features to avoid potential future liability – appears doubtful in light of the many well-established hurdles placed in the way of a successful claim and the more tangential relationship of most of the autonomy challenges to physical injury. In spite of the problematic nature of AI's novel features, their hidden and ancillary position will mean

that many patients will never have the opportunity to confront them and direct their thoughts and actions accordingly.

Chapter 8: Assessment of the comparison and of its wider significance

In reviewing the line of argumentation pursued throughout this work one must begin with the deployment of artificial intelligence (AI) or machine learning (ML) in medicine. As has been elaborated, this is a sophisticated novel technology that is rapidly gaining acceptance in the healthcare sphere and holds yet more promise. This should not be forgotten as one focuses on the series of problems that have been identified in relation to patient autonomy.

Chief among these challenges is the inherent opacity of AI decision making and its fraught relationship with established scientific knowledge. This is marked by the difficulty of carrying out established methods of performance evaluation to gauge the technology's functioning. As a result, even established uses of AI/ML exhibit an uncertainty that justifies analogies to forms of innovative or unlicensed treatment: they possess risk-related characteristics.

Another challenge emerged from AI's ability to pursue goal-directed action relatively independently. Without knowledge of these goals and this functioning, the choice to use certain devices could pre-determine a patient's decision, failing to facilitate their decision-making with regard to fundamental aspects of their care. To a lesser degree, ML devices would also be in a position to surreptitiously insert their own objectives into shared processes of decision making. These influences could be conceptualised as biases, nudges or manipulations that hinder a patient's ability to make theoretically rational choices that reflect their own values.

Lastly, the *raison d'être* of many uses of ML technologies was to mimic the capabilities of human experts. It was outlined how this resulted in various forms of cooperation: sometimes substituting a human specialist, sometimes providing them with a second opinion. It was judged that such cooperation could constitute an epistemological challenge for patients, who are normally justified in relying on the pronouncements of human experts whom they trust in their reasoning. This trust was found not to be transposable to machine-rendered outputs.

These normative, bioethical challenges were then argued to be of significance to the two selected legal systems. Both the UK and Californian com-

mon law have operationalised the concept of autonomy as a legal principle and both have appealed to the dimensions that make up our understanding of procedural patient autonomy. It was already noted at this stage, that the Californian courts were more prepared to impose rule-specific, doctrinal limitations on their reasoning with the autonomy principle. One may say that these countervailing concerns were generally given more weight.

Under the influence of this principle, the mechanisms of battery and negligence have been interpreted to anticipate the common law response to some aspects of AI's distinct challenges. This chapter will begin by comparatively evaluating the precise nature of these responses, assessing their adequacy and considering alternative forms of legal protection that may be necessary to overcome inherent shortcomings.

In the course of this, one should be mindful of the multifaceted ways in which the law is interacting with a technology that is precipitating widespread societal changes. There is not one reactive relationship between the two. Legal reasoning plays a proactive role, anticipating and guiding the appropriate responses that can be made to this innovation. An appreciation of this dimension runs throughout this chapter, culminating in a critical revision of the prominent assumptions identified in the law and technology literature in the introduction of this work.

I. The limits of the common law

Having examined the approaches of English and Californian common law, one can begin to discern some of the striking similarities and differences. The state of affairs after the application of the normative principle has been summarised in Table 1 and Table 2 below. These provide a side-by-side comparison of the battery and negligence mechanisms in the UK and California.

Table 1: Application of the battery action to AI in the UK and California

Jurisdiction	Classification	Elements of the claim				AI-related information required for valid consent			
		Contact	Intention	Hostility or unlawfulness					
UK	Outcome of autonomy-based argumentation	Direct, immediate interference	Intentional touching with a view to use AI	No further requirement	That significant outputs may be generated by the AI without the patient's subsequent ability to intervene				
	Circumstances still excluded	Decisions not to treat and non-physical interactions	Unintentional touching	None	AI use <i>per se</i>	AI's risk-related status	Shifts in human-AI expertise	Lesser AI influences	
California, U.S.	Outcome of autonomy-based argumentation	Direct and indirect contact	Inferred intention to deviate from patient consent	Unconsented to touching	No specific AI-related information				
	Circumstances still excluded	Decisions not to treat	Patient-imposed conditions on care of which the professional is unaware	Cases also excluded by consent element	AI use <i>per se</i>	AI's risk-related status	Shifts in human-AI expertise	AI determinations and lesser AI influences	

Key of autonomy-based argumentation	
	Strong argument for the stated position
	Available argument of uncertain strength for the stated position
	No reasonably available argument for a change of the stated position
	Corollary to an argued position

Table 2: Application of the negligence action to AI in the UK and California

Jurisdiction	Classification	Elements of the claim					
		Damage	Duty	AI-related information required to avoid breach		Causation	
UK	Outcome of autonomy-based argumentation	Physical harm or significant violation of a patient's autonomy	AI use by a medical professional	AI's risk-related status	That the selection of an AI can pre-determine a choice of alternative	That AI has led to a substantial depreciation of expected expertise	Individual's decision seriously affected by disclosure, or they would have delayed or avoided harm
	Circumstances still excluded	Lesser autonomy interferences	AI use only by institution	Lesser forms of AI influence		Individual's decision not seriously affected, or they would not have delayed or avoided harm	
California, U.S.	Outcome of autonomy-based argumentation	Physical harm	AI use by a primary caregiving professional	AI's risk-related status	That the selection of an AI can pre-determine a choice of alternative	That AI has led to a significant, non-obvious changes to status or expertise	Reasonable patient would have altered their decision after disclosure
	Circumstances still excluded	Autonomy violations <i>per se</i>	AI use only by a secondary caregiver and/or institution	Lesser forms of AI influence		Reasonable patient would not have altered their decision after disclosure	

Key of autonomy-based argumentation		
	Strong argument for the stated position	No reasonably available argument for a change of the stated position
	Available argument of uncertain strength for the stated position	Corollary to an argued position

A. Negligence

The point of departure for our comparative assessment are the informational obligations that were seen to be the most extensive and nuanced – and therefore the most relevant to medical ML devices – in both jurisdictions. These are those imposed under the negligence mechanism. After comparing the nature of mandated disclosure in both jurisdictions (under the breach element) and its relation to AI's challenges, we move on to consider the impact of the wider limitations imposed by the nature of these torts. Throughout, the strength and effect of autonomy-based argumentation is assessed.

1. Informational requirements under the breach element

Under the breach element of negligence quite some overlap has been identified between the two systems. This is not surprising given that the analysis was shaped by comparable autonomy principles. Nevertheless, distinctions already emerged regarding the specific norms or tests that had been elaborated by the leading cases of *Cobbs v. Grant* and *Montgomery v Lanarkshire Health Board* for California and the United Kingdom respectively.¹³⁵⁸

Beginning with the similarity of the two approaches, both tests for disclosure appealed to the figure of the patient as the primary decision-maker. Both courts were highly critical of medical paternalism, which would have been the consequence of a test catering exclusively to professional custom. Instead, their tests were directed at the needs of the patient and asserted the patient's right to assess relevant information by reference to their own values. This placed an onus on the professional to facilitate decision-making, while not overburdening the patient with information. In short, they combined dimensions of decisional and practical autonomy to arrive at an appropriate definition of reasonable disclosure.

But there were also revealing differences. The Californian rhetoric painted a picture of the patient as much more dependent on their professional, less able to find and evaluate evidence for themselves. This arguably led to a standard of care that remained more deferential to the experts and was developed in an incremental fashion, with a view to the certainty that

¹³⁵⁸ *Montgomery v Lanarkshire Health Board* [2015] UKSC 11, [2015] AC 1430; *Cobbs v. Grant* (1972) 8 Cal.3d 229.

the clinical profession required. In line with this, Californian common law relied more heavily on objective criteria such as the reasonable person and, to some extent, the reasonable professional.

The UK, by contrast, paired its recognition of the professional's facilitative role with a more independent view of the patient. The latter was able to find, evaluate and integrate evidence into their broader decision making and the advice of experts was subservient to their consumers' choice. This led to a more open elaboration of the test. Certain categories of information and restrictions upon these were highlighted, but disclosure obligations ultimately depended on the needs of the individual patient.

In this manner, both operationalisations of the informed consent standard reinforced the relevance of the procedural account of patient autonomy, albeit with different emphases. At the same time, and for their own distinct reasons, they indicated that autonomy-based arguments had to be combined with forms of analogical reasoning. In California this was necessitated by the express limitation of disclosure obligations by reference to established categories and the focus on incremental development. In the UK it emerged from the indeterminacy of the normative framework. Without more specific guidance in place, there was a wide discretion in the manner in which autonomy-based disclosure arguments could be framed, but their strength remained uncertain. To establish whether the common law would require a professional to advise a patient of AI-related characteristics the most forceful arguments appealed to both patient autonomy and concrete applications.¹³⁵⁹

i. Risk-relevant characteristics

A coalescence of principle and analogy was evident in relation to the risk-relevant status of medical AI. The disclosure of specific risks was seen to be the type of disclosure that enjoyed the most support by analogical reasoning in both jurisdictions – and to some extent it could be applied to

1359 As Green and Sales have pointed out the arguments from analogy and of principle are always intertwined: 'a judge may resort to reasoning by principle and/or analogy. That exercise is reciprocal. Reasoning by principle will allow the judge consistently to apply the law across analogous situations; reasoning by analogy will assist the judge in more accurately identifying and articulating the underlying principle': Green and Sales, 'Law, Technology and the Common Law Method in the United Kingdom' [2023](5) *Europäische Zeitschrift für Wirtschaftsrecht* p. 205, 211.

the novel technology. However, this was not extensive enough to adequately facilitate a patient's assessment of AI's relationship to risk. In other words, the best-supported argument by analogy did not offer adequate protection to the underlying principle of patient autonomy.

This is where the common law's adaptability and flexibility becomes relevant. Information concerning the innovative or unlicensed nature of procedures was not directly covered by the established disclosure norms of either jurisdiction. Nevertheless, with a view to the patient's need, there are lines of case law mandating such disclosure.

The strongest support for this argument was identified in California. Multiple cases on innovative or experimental procedures had affirmed the obligation of the professional to disclose these general characteristics. In the course of this, the court referenced a number of factors also found to be relevant under our conception of patient autonomy, such as: a lack of evidence regarding a device, its unapproved regulatory status, the availability of alternatives, and the plaintiff's preference for a conservative approach.¹³⁶⁰ In sum, by reference to existing case law and particularly through an implicit reliance on the autonomy principle, general risk-related characteristics have been held disclosable in California.

This provides a strong basis for the argument that AI's comparable status must likewise be disclosed. The argument by analogy may still be imperfect, in virtue of the different factual circumstances under consideration, but the argument from principle is made all the stronger by the existence of this line of cases and their reasonings' relation to the outlined principle.

In England the situation is somewhat different. There is no discernible practice of the courts extending risk disclosure obligations to generalised features, such as the innovative nature of the procedure. There is at least one case that has apparently been decided on this basis.¹³⁶¹ Yet the extent to which it employed reasoning related to the autonomy principle was limited, referring primarily to the uncertainty of a procedure and its lack of validation.

To supplement this shortcoming weaker forms of argumentations were introduced. Statements from other cases could be drawn upon that were

1360 *Daum v. SpineCare Medical Group, Inc.* (1997) 52 Cal.App.4th 1285.

1361 *Mills v Oxford University Hospitals NHS Trust*, [2019] EWHC 936 (QB), (2019) 170 BMLR 100.

either *obiter*,¹³⁶² or remained connected to a narrower form of risk analysis.¹³⁶³ Arguments advanced in the UK academic literature on this issue, which strongly supported the disclosure of risk-relevant characteristics, were also adduced. These factors provide indirect support for our line of argumentation and indicate that further extensions to AI characteristics will be arguable.

In summation, in both jurisdictions the general characteristics of AI, capable of impacting a patient's risk assessment, will have to be disclosed when assessed in light of the autonomy principle. In California, there is a forceful argument to this effect. In the UK there is an available argument of uncertain strength.

ii. Goal-directed action by AI

Another category of disclosure that was supported by a combined analysis of existing case law and principle related to the patient's control over choices in their care. In both jurisdictions the disclosure of alternative procedures has been firmly rooted in the need to connect a certain care pathway to the patient's unique values and preferences. This represented a clear and consistent appeal to the reflective dimension of autonomy, which was stated particularly powerfully in *Montgomery* in the UK and, to a lesser extent, in a series of Californian cases.¹³⁶⁴ The evident relevance of this principle provides a strong argument for the disclosure of the danger that an AI device pre-determines certain choices of the patient, by for example making a significant, serious diagnosis. This is reinforced by an analogy to well-established case law that holds it impermissible for a human expert to pre-determine a patient's choice by positively failing to provide them with information.

A possible limitation that must be conceded in both jurisdictions, is that a patient must have an available or reasonable alternative before a duty to advise arises. This was found to exclude at least some cases where the goal-directed action of AI had the potential to disconnect the patient from

1362 *Jones v Taunton and Somerset NHS Foundation Trust* [2019] EWHC 1408 (QB), [2019] 6 WLUK 193.

1363 *Webster v Burton Hospitals NHS Foundation Trust* [2017] EWCA Civ 62, (2017) 154 BMLR 129.

1364 The examples of *Mathis v. Morrissey* (1992) 11 Cal.App.4th 332 and *Wilson v. Merritt* (2006) 142 Cal.App.4th 1125 were given.

the pursuit of their individualised goals. Nonetheless, as it has been argued that a process of autonomous decision-making will only be disturbed if an identifiable decision is undermined, this requirement arguably represents an application of the autonomy principle, rather than a deviation from it. To the extent that the case-law tracks instances of ‘identifiable’ or ‘real’ decision-making accurately, and imposes a corresponding obligation to disclose the pre-determination of a relevant choice, there is no basis for censure.

iii. Human-AI expertise

Regarding shifts in human-AI expertise, the analogies that can be drawn in either jurisdiction remain relatively limited, although arguments from principle can be constructed on their basis. One possible interpretation of the English case of *Jones v Royal Devon and Exeter NHS Foundation Trust* was that: substituting the level expertise that the patient expected to be used in their procedure, without proper disclosure, violated their informed consent.¹³⁶⁵ This could be supported by limited *obiter dicta* in other cases and by arguments from principle, guiding the application of the higher-level *Montgomery* test. However, the scope of this argument is potentially quite limited. Disclosure may only be mandated where, as in *Jones*, the patient demonstrated a strong belief in an extraordinarily high level of expertise or, at least, where the shift in expertise was so substantial to be considered material information for the patient.

Similarly, in California the most relevant authority, *Davis v. Physician Assistant Bd*, stands relatively isolated in virtue of the fact that it assessed a disciplinary action and that it emphasised the difference in the professional’s status, rather than their expertise *per se*.¹³⁶⁶ As such it did not constitute a direct precedent for the claim that a breach of informed consent followed from a professional’s lack of expertise. Such relevance could only be inferred. At a broader level, there were conflicting indicators as to the relevance of a clinician’s expertise to a patient’s informed consent. It was possible to reference more targeted informed consent actions, concerned

1365 *Jones v Royal Devon and Exeter NHS Foundation Trust* [2015] Lexis Citation 3571.

1366 *Davis v. Physician Assistant Bd*. (2021) 66 Cal.App.5th 227.

with the risks of not seeking more expert advice for a specific procedure.¹³⁶⁷ A reconciliation of these separate strands with *Davis* reinforced the principled argument that such information must be disclosed to facilitate a patient's autonomous decision making.

In sum, case law in both jurisdictions touched on the problem that was highlighted by AI's ability to substitute human expertise. This can support the principled argument that a failure to disclose shifts in expertise, especially substantial shifts towards the application of technological expertise, can amount a breach of the patient's informed consent. However, the existent cases do not provide clear analogical support for this proposition – still less do they resolve the issue of whether a relative lack of expertise of a human specialist *vis-à-vis* an AI would be sufficient to trigger a disclosure obligation. Rather, it may be the case that an alteration in professional status (in California) or the expectation of a particularly high level of expertise (in the UK) is required. This is why there is an available autonomy-based argument of uncertain strength, encompassing at least a subset of cases where the use of AI devices alters human capabilities.

iv. Informational manipulation

The type of disclosure that was the most challenging to subsume under established legal categories, related to AI's ability to manipulate a patient's decision making. This was deemed to be an issue for procedural autonomy in so far as such manipulation did not engage the patient's own reflective reasoning about their therapeutic goals. In other words, where the nudging possessed a surreptitious non-therapeutic objective.

In this respect it is interesting that the Californian Supreme Court case of *Moore v. Regents of the University of California* constructed a disclosure category aimed precisely at situations where human physicians' decision-making had been potentially subject to these same kinds of influences.¹³⁶⁸ An argument by analogy would theoretically support disclosure of AI biases that may affect professional and patient reasoning. However, this was not capable of providing practicable recourse to a patient who was unlikely to be able to prove that the AI had manipulated them in favour of certain non-therapeutic objectives. Nor would it align with negligence's focus on

1367 *Moore v. Preventive Medicine Medical Group, Inc.* (1986) 178 Cal.App.3d 728; *Scalere v. Stenson* (1989) 211 Cal.App.3d 1446.

1368 *Moore v. Regents of University of California* (1990) 51 Cal.3d 120.

the reasonable defendant, whereby the advising professional would have known, or ought to have known, of such objectives. Here the autonomy principle was ultimately of little assistance.

In the UK these practical hurdles would no doubt be present as well. On top of this, its negligence case law had not yet identified the non-disclosure of comparable influences to be breaches of a professional's duty to obtain informed consent. It had only recognised, at a more abstract level, that external pressures of a sufficient severity could undermine the patient's ability to give their (informed) consent.¹³⁶⁹

In consequence, there is no purchase for principled, autonomy-based argumentation to establish an AI's potential nudges – and the induced biases – as relevant information that must be disclosed. This is a gap that is left even by the present, forward-looking interpretation of the most autonomy-receptive element of the negligence action.

2. Non-informational requirements

The analysis of the breach element has revealed several differences between the two jurisdictions, not least in the strength of available forms of argumentation for the disclosure of different categories of AI-related characteristics. At the same time, a considerable overlap could be identified between them. This can be explained by the close association between rule-based and principled argumentation – whereby the test for the breach of disclosure obligations is receptive to considerations of autonomy – and by the specification of procedural autonomy principles in both jurisdictions.

Things are very different regarding the non-information-related elements of the negligence action: damage, duty and causation. The relevance of the autonomy principle to these elements is less readily constructed, but they serve as important checks on the success of any claim and thus the protection afforded by obligations to disclose information.

The Californian courts have exhibited a particularly strict adherence to rule-specific limitations. The relevant formulation of the damage element continues to require the eventuation of physical harm to the patient, the group of individuals who may have an obligation to obtain a patient's informed consent remains restrictive – primary caregiving professionals – and the causation test focuses on what the reasonable patient, rather

1369 *Thefaut v Johnston* [2017] EWHC 497 (QB), [2017] 3 WLUK 328.

than the individual patient with their specific values, would have done. Stated succinctly, the common law has here been reticent to sacrifice coherence and doctrinal stability in favour of flexibility and adaptation to the demands of the autonomy principle. The scenarios involving AI that have, as a result, been excluded from the imposition of informed consent requirements can be seen in Table 2.

The UK courts have taken a very different approach. Decided cases and the autonomy principle indicate that disclosure obligations concerning AI extend to the patient's entire care team. Beyond this, the significance attached to the realisation of the autonomy principle has, at the very least, generated considerable uncertainty around the nature of the damage and causation elements of the negligence action.

Arguments can reasonably be constructed that the English courts: (1) have recognised a sufficiently serious autonomy violation *per se* as a form of damage and/or (2) have found that if there is a breach of obligation relating to the patient's autonomy, then the orthodox but-for analysis under the causation element (linking breach and damage) is relaxed. Arguably, the patient need now only prove that their autonomy was sufficiently impaired, or that breaches of informational obligations would have delayed or changed a decision that led to the eventuation of physical harm. Moreover, it is beyond doubt that *Montgomery* has led to a causation test that is protective of an individual's reflective autonomy: i.e. their long-held, defensible preferences.

In the UK therefore, the common law has shown itself to be adaptable to shifts in wider society, favouring the stronger protection of the autonomy principle. Potentially this means that a much smaller class of AI uses, which are likewise outlined in Table 2, fall outside the purview of the discussed disclosure obligations. Yet this adaptation has, thus far, come at the cost of considerable uncertainty and incoherence. Although arguments in favour of stronger, more comprehensive autonomy protections are available by reference to the highest legal authority, it appears that they do not carry much force with the lower courts, who are still concerned to maintain doctrinal stability.

Finally, it is intriguing that one remaining non-informational shortcoming in both common law systems, the lack of stringent informational obligations on institutional users of AI, has been partially addressed by a specific statutory intervention in the UK. Namely, through the *UK General Data Protection Regulation* (UK GDPR) and *Data Protection Act* (DPA) 2018. But it was noted that, beyond the disclosure of AI use *per se*, the demands of

this statutory regulation are currently still unclear and in need of principled interpretation.

B. Battery

Our analysis of the battery action presents us, somewhat surprisingly, with a rather different picture than the one outlined above. Both in terms of the informational requirements imposed on valid consent and in terms of the impact of the autonomy principle on non-consent requirements.

1. Informational requirements for valid consent

The informational disclosure obligations that were so receptive to the autonomy principle in negligence, were much less so in battery. Rule-specific factors – encompassing the need to do justice to medical professionals in both jurisdictions and the limitation on damages for negligence, but not battery, actions in California – mandated a more limited obligation to disclose information.

In the UK this more limited duty meant that shifts in human-AI expertise did not have to be disclosed, in spite of potential analogies that could be drawn to established case law. The only feasible argument that could be advanced in favour of disclosure related to certain basic categories of information concerning the nature and broad purpose of a procedure. This was argued to correspond to decisionally necessary true beliefs, which demanded strong protection under the procedural account of autonomy, and to include the AI's potential pre-determination of a significant choice. To this extent, some protection could be afforded to autonomy. Still, even this argument was of a highly uncertain strength given the lack of clarity in this area of the law.

In California by contrast, the narrow delineation of relevant categories precluded any forms of disclosure that could help address AI's more nuanced autonomy challenges. For example, disclosure regarding shifts of expertise. In addition, a narrow focus on the physical nature of a procedure precluded the requirement of valid consent from providing protection for more fundamental, decisionally necessary beliefs. This included beliefs concerning the goal-directed nature of AI functioning.

2. Non-informational requirements

As in negligence, battery's non-informational requirements have the potential to limit the scope of the action severely. In the UK this is exemplified by the condition that there must be a direct interference with the patient's body and, in addition, that there must have been some intention to use an AI with an undisclosed, significant objective before such a direct interference takes place. The courts have given no indication that they would be prepared to abrogate these elements in light of considerations of patient autonomy. Quite the opposite.

By contrast, the rule-specific non-informational considerations in California have been framed fundamentally by reference to the autonomy principle. Violation of patient consent has established the unlawful nature of a procedure, it has allowed the courts to infer the requisite intent on behalf of the professional and it has provided the basis for an argument that indirect contact, crucially including the prescription and administration of medication, falls under the scope of battery.

Given the contrast with the approach taken by the two jurisdictions under negligence, these findings appear incongruous. It is possible that they exhibit a certain trade-off within the battery norm itself. If the common law opts to frame the consent requirement less restrictively, as in the UK, then other limiting factors must correspondingly be emphasised to keep it within acceptable doctrinal limitations. If the consent requirement is framed highly restrictively, as in California, then one can afford to relax other requirements, which depend on the patient first passing the hurdle of invalid consent. In the final analysis therefore, the adaptability of this tort appears to be kept within relatively tight bounds by both common law systems.

C. Conclusion

The above arguments have highlighted some of the strengths, but also some of the limitations of common law tort systems' ability to respond to novel factual circumstances and technological innovations. Flexible, adaptable categories of disclosure have been framed under the negligence claim in both jurisdictions. Indeed, these could be argued to cover a meaningful proportion of AI's novel autonomy challenges – that is, interpreted appropriately in light of the underlying autonomy principle. Nonetheless, the dif-

ferences between the Californian and English approaches have illustrated the issues associated with such flexibility.

If appeals to principle become too overt and untethered from more granular norms and guidance, then there is a real danger that effective protection is jeopardised by the resulting uncertainty. After a change in the standard of care, English law is in a state of flux and development and it has thus far failed to provide such norms. California, by contrast, has been more careful to operate by reference to incremental evolution and objective criteria. Arguably, this strikes a more appropriate balance between coherence and legal certainty on the one side and the protection of the autonomy principle on the other.

A similar argument may be transposed to the way in which both jurisdictions have handled the non-informational requirements of the negligence action. In several regards the UK has been prepared to infuse many elements of the action, which are designed to generate stability, with unorthodox considerations of patient autonomy. Californian law has refused to take this step. But again, one must make a context-sensitive analysis. The ability to protect patient autonomy in light of AI's more nuanced challenges is fundamentally dependent upon a loosening of these broader doctrinal requirements and striking the right balance arguably requires the development of different kinds of doctrinal limitations, ones which are receptive to principle. For example, through defining actionable damage by reference to a sufficiently restrictive, operationalised conception of patient autonomy – as the procedural account has been said to provide. This would require a proactive court to recognise that fundamental shifts in the nature of the negligence action have already taken place.¹³⁷⁰

An alternative, as some authors have proposed, may be the creation of a *sui generis* tort with distinct elements.¹³⁷¹ This would have the advantage of maintaining the structure of negligence, while also offering the principled

1370 As Green and Sales remark, the ability to make these kinds of fundamental decisions may be the preserve of the Supreme Court: Green and Sales, 'Law, Technology and the Common Law Method in the United Kingdom' [2023](5) *Europäische Zeitschrift für Wirtschaftsrecht* p. 205, 209. Recall also that Leggatt LJ noted that the Supreme Court may need to reconsider the area of damage and its relation to patient autonomy: *Duce v Worcestershire Acute Hospitals NHS Trust* [2018] EWCA Civ 1307, (2018) 164 BMLR 1 [92].

1371 Nolan, 'Damage in the English Law of Negligence' (2013) 4(3) *Journal of European Tort Law* p. 259, 376-382; Mulligan, 'A Vindictory Approach to Tortious Liability for Mistakes in Assisted Human Reproduction' (2020) 40(1) *Legal Studies* p. 55, 71-74.

protection that modern society appears to be demanding. However, currently the courts have given no indication that they would be prepared to take this step.

In the final analysis it must therefore be admitted that neither the Californian nor the English common law system currently strike the outlined balance. The former system emphasises rigidity and objectivity at the expense of the protection of principle. The latter seeks to protect the principle, but has thus far failed to convert this into sufficiently clear and specific norms.

As a result, the common law on informed consent is limited in its ability to respond to the novel, nuanced types of autonomy challenges posed by medical AI. At most the identified disclosure obligations under negligence's breach element could serve a signalling function, falling short of offering substantive protection. To be certain to avoid legal sanctions, medical professionals ought to be careful to address the following AI characteristics (where applicable to the specific device): their risk-relevant status, the substantial, non-obvious shifts they may cause in human-to-AI expertise and the ability of AI to pre-determine certain significant objectives in their relatively independent functioning. If a patient is not informed of these, however, they will probably lack recourse to the courts, unless they incidentally happen to fulfil the outlined additional requirements.

II. Guidance for a bespoke statutory scheme

In this section it is argued that lawmakers and regulators and, more generally actors who design modes of regulation complementing or substituting the common law, can derive useful guidance from the legal operation examined throughout this work. To be exact, it is considered how an approach building on this – a supplementary statutory scheme – could be designed to complement the identified weaknesses of the common law. By conceptualising a desirable legal solution to the autonomy challenges of clinical AI/ML, this sets the scene for the final, underlying analysis of how legal reasoning interacts with the phenomenon of technological innovation.

Of course, by recommending a legislative intervention for a subject that traditionally falls within the common law's purview, the present section taps into a much deeper discussion. Namely, the appropriate roles played by courts and legislatures in a modern society and – as addressed in Chapter

1 – the relative advantages and disadvantages of accommodating social and technological progress through either modality.¹³⁷²

While the preceding argumentation has served to explore primarily the adaptability of the common law, as well as the normative legitimacy of its structure, this section seeks to strike a middle path that capitalises on the strengths of both sources of law. By advocating for a legislative scheme that draws on the common law mode of reasoning, is supplementary to its informed consent framework and contains mechanisms for adaptation, it is hoped that the advantages of the developed approach can be broadly maintained. Simultaneously, the statutory form is not subject to the identified defects of the current system of judge-made law. It offers improvements in terms of specificity, clarity and in terms of monitoring compliance and enforcement. It is these advantages that this section additionally seeks to exploit, proposing a scheme that guarantees an effective protection of patient autonomy in relation to clinical ML devices.

Existing legislative structures will be considered in order to make this case and in order to determine the nature that such a scheme could take in the specific contexts of California and the UK. While legislators are much less constrained than judges by settled legal positions – and are technically at liberty to tailor solutions to the instrumental demands of innovation – it will be seen that they have nevertheless tended to institute incremental changes in this area. In effect, they are also liable to be directed by the existing legal material, both by the prevailing common law position and by the form of existent statutory and regulatory interventions.

In consequence, to inform the recommendations for a bespoke statutory scheme targeted at the resolution of AI's particular autonomy problems, we undertake two analyses in this section. First, we examine the nature of existing legislative informed consent schemes in the two jurisdictions. Second, we determine whether their underlying rationale can be transferred to the regulation of medical ML devices and, if so, what shape an adequate, bespoke solution would take. At this stage one can then draw on the results of the application of our legal mode of reasoning in the common law.

1372 For fundamentally different views on this matter, compare: Friedman, *Law and Society: An Introduction* (1977) 164-165; Calabresi, *A Common Law for the Age of Statutes* (1985) 2.

A. Existing consent statutes

As referred to in Chapter 1, both of the selected jurisdictions have imposed certain statutory regimes for the disclosure of information. For us, the most interesting precedent relates to mandated disclosure of certain pieces of information that may be additional to, or understood to reinforce, the kinds of disclosure mandated at common law.

1. United Kingdom

In the United Kingdom three instances can be cited in this respect. The statutory requirements for: clinical trials of investigational medicinal products; the storage and use of embryos and gametes, and; the removal, storage and use of human tissues.

The first example is provided by the consent requirements for clinical trials of investigational medicinal products. Under the *Medicines for Human Use (Clinical Trials) Regulations* 2004 an individual gives informed consent to participate in a trial only where they are ‘informed of the nature, significance, implications and risks of the trial’.¹³⁷³ This consent must be evidenced or recorded in writing.¹³⁷⁴ In addition, an investigator must arrange an interview with the subject where they have the ‘opportunity to understand the objectives, risks and inconveniences of the trial and the conditions under which it is to be conducted’,¹³⁷⁵ they must be provided with a point of contact for the purposes of obtaining further information,¹³⁷⁶ and they are to be informed of their right to withdraw from the trial at any time.¹³⁷⁷

Overall, in terms of the substantive categories of information that must be provided, the legislation largely replicates the common law. But in so doing it clearly demarcates and emphasises aspects that are more tangential to established forms of disclosure, such as: the objectives of the trial, its conditions and the right to withdraw. Furthermore, it establishes various formal or procedural mechanisms that frame the disclosure: it must be in writing, there must be an interview with the subject and there must be a point of contact for further discussion.

1373 *Medicines for Human Use (Clinical Trials) Regulations* 2004 schedule 1, part 1, paragraph 3(1)(a).

1374 *ibid* schedule 1, part 1, paragraph 3(1)(b).

1375 *ibid* schedule 1, part 3, paragraph 1.

1376 *ibid* schedule 1, part 3, paragraph 5.

1377 *ibid* schedule 1, part 3, paragraph 2.

The *Human Fertilisation and Embryology Act* 1990 can be considered apart from this. It requires that, for the purposes of the licence that is a necessary prerequisite for the ‘donation, storage and use of gametes and embryos’,¹³⁷⁸ certain consent conditions must be complied with.¹³⁷⁹ In particular, a consent to the use of an embryo must specify one of several possible purposes, in terms of treatment, training or research, and a consent to storage must specify the maximum period of storage and what is to be done with the gametes, embryo or admixed embryo if the person is unable to vary or withdraw consent.¹³⁸⁰ The licensing authority may specify further consent requirements in its directions.¹³⁸¹

More fundamentally, alongside such specific categories of information there are also general requirements regarding the fact that an individual must receive ‘such relevant information as is proper’ and that they receive counselling regarding the implications of their decision.¹³⁸² Consent must assume a certain form. Normally this means that it must be in writing and signed.¹³⁸³

In consequence, while there is some overlap with the common law requirements, indicated by references to ‘proper’ or ‘informed’ consent, one can see that there is considerable detailed, specific guidance that relates to the innovative nature of the subject matter. Moreover, the statute puts in place mechanisms regarding: the development of more elaborate guidance, the form of consent, and the need for counselling.

The last statute to be examined is the *Human Tissue Act* 2004. This is a wide-ranging statute that primarily governs the consent necessary from relatives for the removal, storage and use of human tissue from deceased persons.¹³⁸⁴ By contrast to the first two schemes, this largely leaves the consent requirements for the removal of material from living persons to

1378 *Jennings v Human Fertilisation and Embryology Authority* [2022] EWHC 1619 (Fam), (2023) 189 BMLR 17 [24].

1379 Human Fertilisation and Embryology Act 1990 (as amended) section 12(1)(c) and schedule 3.

1380 *ibid* schedule 3, paragraph 2(1)-(2). It also requires them to be informed about the conditions for varying or withdrawing consent: *ibid* schedule 3, paragraph 4-4A.

1381 *ibid* schedule 3, paragraph 2(3). Currently see: Human Fertilisation and Embryology Authority, ‘Code of Practice’ (2021 Ninth Edition) <<https://portal.hfea.gov.uk/knowledge-base/read-the-code-of-practice/>> accessed 26.3.2023.

1382 Human Fertilisation and Embryology Act 1990 (as amended) section 13(6)-(6A), and schedule 3, paragraph 3(1).

1383 *ibid* schedule 3, paragraph 1(1).

1384 Human Tissue Act 2004 schedule 1.

the regulation of the common law.¹³⁸⁵ Substantively, it does not elaborate upon vague notions of ‘appropriate consent’¹³⁸⁶ or ‘qualifying consent’.¹³⁸⁷ However, here too a procedure was created whereby an institution (the Human Tissue Authority) must provide more specific guidance.¹³⁸⁸

2. California

Beginning with the common ground that can be found in our U.S. case study, one can point to California’s statutory regime covering consent to medical experimentation.¹³⁸⁹ This posits stringent disclosure obligations, mandating a research subject to be provided with the experimental subject’s bill of rights and information on: ‘the nature and purpose of the experiment’; ‘an explanation of the procedures to be followed in the medical experiment, and any drug or device to be utilized’; discomforts and risks; benefits to be expected (if applicable); alternative drugs, devices or procedures, and; available forms of medical treatments in case of complications.¹³⁹⁰

In addition, a subject must also be given the opportunity to ask questions, they must be advised that they can withdraw from the treatment without prejudice, they must be handed a copy of the written consent form and they must make their decision free from ‘any element of force, fraud, deceit, duress, coercion, or undue influence’.¹³⁹¹ In this manner one can see that, alongside aspects of information disclosure, formal safeguards (the written nature of consent and specific rights to ask questions) have been put in place.

More generally, the Californian legislature has been much more proactive than its transatlantic counterpart in its complementation and substitution of common law informed consent requirements. The state imposes

1385 *ibid* section 45 and schedule 4.

1386 *ibid* section 3. For an analysis see: McHale, ‘Appropriate Consent’ and the Use of Human Material for Research Purposes: The Competent Adult’ (2006) 1(4) *Clinical Ethics* p. 195.

1387 Human Tissue Act 2004, schedule 4, part 1.

1388 *ibid* section 26.

1389 California Health and Safety Code sections 24170ff (the ‘Protection of Human Subjects in Medical Experimentation Act’).

1390 *ibid* section 24173 and section 24172, subdivisions (a)-(f).

1391 *ibid* section 24172, subdivisions (g)-(j).

statutory requirements regarding: the treatment of breast cancer,¹³⁹² the performance of hysterectomies,¹³⁹³ sperm or ova removal,¹³⁹⁴ and cosmetic implants.¹³⁹⁵ It is beyond the scope of this work to conduct a detailed analysis of these numerous provisions, but certain shared characteristics can be identified.

In this respect, it is notable that several common law classes of disclosure are replicated in the outlined provisions in various combinations, such as the nature of the procedure, risks, advantages and alternative treatment methods.¹³⁹⁶ At the same time, more specific pieces of information are also required. For example, in the case of breast cancer treatment it includes: available treatment options that are at the clinical trials stage, available telephone numbers for organisations that can provide information to the patient, a discussion of breast reconstruction surgery and statistics on the incidence of breast cancer.¹³⁹⁷ At least one of these categories – the disclosure of statistical information – clearly overrides the default common law position which was laid down by *Arato v. Avedon*, as examined in Chapter 7.¹³⁹⁸

Statute has also made provision for various kinds of formalities and procedural mechanisms. This includes the development of written summaries by departments, as well as their provision by professionals,¹³⁹⁹ and the necessity of written consent.¹⁴⁰⁰ Regarding the procedural operationalisation of consent requirements, there are then also mechanisms for the recurrent revision of such information summaries,¹⁴⁰¹ and for the imposition of automatic civil penalties for repeated failures of providers to obtain informed consent.¹⁴⁰² This serves to illustrate the broader institutional instruments available to the legislature in implementing and enforcing its mandates.

1392 *ibid* section 109275.

1393 *ibid* sections 1690-1691.

1394 California Business and Professions Code section 2260.

1395 *ibid* section 2259.

1396 California Health and Safety Code section 109275, subdivision (c); *ibid* section 1690, subdivision (a); California Business and Professions Code section 2259, subdivision (e).

1397 California Health and Safety Code section 109275, subdivision (c)(3).

1398 *Arato v. Avedon* (1993) 5 Cal.4th 1172.

1399 California Health and Safety Code section 109275, subdivisions (a), (c).

1400 *ibid* section 1690, subdivision (b).

1401 *ibid* section 109275, subdivision (c)(2).

1402 California Business and Professions Code section 2260, subdivision (e).

3. Rationale

To examine the intent underlying these provisions, the rationale for taking certain issues of clinical consent outside of the common law, one may look towards: any discernible legislative intent, the context for the provisions' passing and the nature of the legislation itself.

California does not generally record the hearings or reports underlying its legislative process. However, in so far as motivations can be deduced from the context, there are indicators that the common law was deemed to have been operating unsatisfactorily. For example, referring *inter alia* to California and the outlined regulation of consent to breast cancer treatment, Pope has noted: 'in the late 1970s and early 1980s, physicians were not disclosing less invasive treatment options to their breast cancer patients. In response, 14 states enacted statutes that require physicians to present the advantages, disadvantages, and risks of all medically viable alternative therapies'.¹⁴⁰³ In other words, the common law was deemed insufficient to ensure that the medical profession disclosed specific kinds of alternatives.

In spite of the UK's generally greater transparency regarding legislative intent, the relationship to the existing common law was only meaningfully addressed in the passing of the *Human Tissue Act* 2004. A public scandal, regarding the non-consented-to removal of tissue – primarily from deceased children – had led to the realisation that there was a *lacuna* in the pre-existing regime, governed partly by common law and partly by statute.¹⁴⁰⁴ Referring to this scandal in its Explanatory Note, the Act explicitly stated that 'the current law in this area was not comprehensive, nor as clear and consistent as it might be for professionals or for the families involved'.¹⁴⁰⁵ In response, it sought to establish the fundamental significance of consent in this area.

The other examples of statutory regulation of informed consent in the UK were not clearly based on the unsatisfactoriness of the *status quo*. The *Medicines for Human Use (Clinical Trials) Regulations* 2004 were not enacted as primary legislation, nor were they enacted *under* a targeted,

1403 Pope, 'Certified Patient Decision Aids: Solving Persistent Problems with Informed Consent Law' (2017) 45(1) *Journal of Law, Medicine & Ethics* p. 12, 17-18.

1404 Jones, *Medical Negligence* (Sixth Edition 2021) para 7-174; McHale, 'Appropriate Consent' and the Use of Human Material for Research Purposes' (2006) 1(4) *Clinical Ethics* p. 195, 195.

1405 Explanatory Notes to the Human Tissue Act 2004 paragraph 5.

enabling Act of Parliament,¹⁴⁰⁶ Rather, they were passed as a regulation implementing an EU directive and detailed discussions on informed consent remained limited. The focus was on the desirability of such reforms in general. Likewise, while the discussions underlying the *Human Fertilisation and Embryology Act* 1990 were much more extensive, they focussed on the more controversial aspects of the Bill. In so far as consent was discussed in relation to these two instruments, it was primarily seen as an added, legitimacy-conferring factor for a novel – otherwise morally and politically suspect – scheme.¹⁴⁰⁷

On the basis of these examples, one is left with two motivating factors. First, confirming the general phenomenon stated at the outset of this section: identifying a deficiency in the operation and application of the common law regarding consent can be an important prerequisite for a statutory intervention. Second, where a novel kind of statutory regime is introduced to regulate a certain clinical circumstance (such as experimentation or new uses of embryos), a legislative framing of consent can serve to bolster the legitimacy of the resulting framework. It is hard to derive a judgment on the adequacy of the common law's operation *per se* in such situations.

However, one can then derive more specific insights from the nature of the statutory provisions themselves. Rather intuitively, these appear to respond primarily to subject matters that are perceived to pose problems of particular social significance. In the informed consent context this includes above all challenges framed in terms of the value of patient autonomy.

For example, it seems that such interventions are thought most appropriate where there is a heightened danger that patient decision making will be subverted by the external imposition of choices upon them. Similarly, it

1406 The enabling Act was the European Communities Act 1972 section 2(2).

1407 In relation to the Medicines for Human Use (Clinical Trials) Regulations 2004, the need for consent was juxtaposed with safeguards for those without capacity to consent: Joint Committee on the Draft Mental Incapacity Bill, 'Draft Mental Incapacity Bill: Session 2002–03 Volume I' (2003) paragraphs 285–289. In relation to the Human Fertilisation and Embryology Bill 1990, it was stated in the House of Commons for example: 'The Bill is also about freedom of conscience. It does not seek to impose sanctions on those who think that research is wrong. It goes to great lengths to define consent. As Lord Hailsham put it in another place, those who would impose an absolute prohibition should ask themselves, what kind of right, in a free and liberal democracy, do they think that they have to say no to a highly responsible group of people working for the benefit of humanity and subject to the authority of Parliament?': House of Commons Debate 2 April 1990, volume 170, columns 953–953.

appears that the regulated circumstances are frequently ones where important interests of the individual are at stake. These kinds of concerns were succinctly stated by the Californian legislation on experimental treatment:

The Legislature hereby finds and declares that medical experimentation on human is vital for the benefit of mankind, however such experimentation shall be undertaken with due respect to the preciousness of human life and the right of individuals to determine what is done to their own bodies (...) There is, and will continue to be, a growing need for protection for citizens of the state from unauthorized, needless, hazardous, or negligently performed medical experiments on human beings.¹⁴⁰⁸

Quite similar statements can be made in relation to the regulation of the use and storage of embryos, of breast cancer treatment, or the performance of hysterectomies.¹⁴⁰⁹

The types of specific categories of disclosure that the statutory mechanisms elaborate, and the manner in which these relate to the existing common law, are also telling. Where such categories are additional to established informed consent precedents, or they clearly override them, this indicates an identified under-inclusiveness of the judge-made law. However, where they merely offer a concretisation of the more general common law standards – which appeared to be in the great majority of cases – it suggests that the common law, while applicable, was deemed too uncertain or conflicted.

In this manner the contribution of the creative process of legal reasoning is recognised. The outputs of this process maintain their relevance even

1408 California Health and Safety Code section 24171.

1409 In relation to breast cancer, see: *ibid* section 109250: ‘Despite intensive campaigns of public education, there is a lack of adequate and accurate information among the public with respect to presently proven methods for the diagnosis, treatment, and cure of cancer. Various persons in this state have represented and continue to represent themselves as possessing medicines, methods, techniques, skills, or devices for the effective diagnosis, treatment, or cure of cancer, whose representations are misleading to the public, with the result that large numbers of the public, relying on the representations, needlessly die of cancer, and substantial amounts of the savings of individuals and families relying on the representations are needlessly wasted. It is, therefore, in the public interest that the public be afforded full and accurate knowledge as to the facilities and methods for the diagnosis, treatment, and cure of cancer available in this state and that to that end there be provided means for testing and investigating the value or lack thereof of alleged cancer remedies, devices, drugs, or compounds, and informing the public of the facts found, and protecting the public from misrepresentation in these matters’.

in relation to (one may presume) much more instrumentally minded legislators who are much less bound by doctrinal limitations. Primarily these outputs are concretised or supplemented so as to allow for a more targeted, effective realisation.

Lastly, it is notable that the examined statutory schemes often outline formal and procedural elements that the common law generally does not, or cannot, lay down. The requirement of written information summaries and consent procedures would be one example. More substantial are the various implementation and enforcement mechanisms that have been discussed: the establishment of regulatory agencies, the issuance of guidance, the need for certain points of contact, automatic civil penalties, etc. The creation and operation of such measures is certainly beyond the capabilities of the common law judge and the adversarial justice system.

B. An informed consent statute for AI

Based on our analysis of AI's challenges, the common law's transformation of these challenges into legally cognisable categories of thought and the legislative solutions already in existence, it will be shown that a supplementary statutory scheme for AI offers one appropriate response to the emergence of this technology.

The starting point here must be the elaborated deficiencies in the common law's ability to meet AI's autonomy challenges. One principal difficulty related not to the fact that disclosure concerning ML devices fell directly outside of established categories, but rather that there was a relationship of uncertain strength between these categories and the specific informational needs generated by AI. Even where strong arguments were identified in principle and by analogy, these remained contestable and subject to general limitations that were, at best, incidental to patient autonomy. A statutory concretisation and 'streamlining' of such requirements would fit well with established precedent and it would draw on the insights of the common law analysis.

Mandating that a patient must be provided with categories of information *via* statute would be a readily available means for resolving the considerable legal uncertainty and the under-inclusiveness of the *status quo*. It should establish, specifically, that the patient be advised of the examined classes of information. This includes: the purpose of the AI use, its general functioning *vis-à-vis* the patient's goals and in relation to scientific know-

ledge, the relation between technology and expertise, its risk-relevant characteristics and, arguably, its potential for influencing decision-making. This is an important autonomy-maintaining baseline. Beyond this, the dynamic, rapidly evolving nature of the technology must be accounted for. More detailed disclosure obligations ought to be subject to mandatory review and revision by a designated authority at regular intervals.

Second, in the AI context the deficiencies of implementing informed consent through the operation of established common law mechanisms has proven a particular area of concern. It has been difficult to target its norms at the right actors, to fulfil the non-consent requirements of its actions, and to ascertain the appropriate amount of compensation that is to be offered in response to violations. At times, particularly regarding the subtle, manipulative influences of ML devices, the disclosure of information *per se* has also appeared to constitute a blunt instrument. The kinds of mechanisms that have been utilised under established legislative schemes would be suited to overcome these hurdles.

Comprehensive informational obligations could be imposed on institutions where they use the technology to engage with patients directly. Compliance with these duties could then be overseen, again, by designated authorities. As one must expect healthcare institutions to be closely involved in the development, selection and deployment of medical AI, this would provide a justifiable basis for the institution-wide regulation of consent requirements. Rather than insisting on the isolation of individual actors, there could be a broader obligation to ensure that healthcare teams using AI have proper consent-taking procedures in place.

Furthermore, one feature of AI is the pervasiveness of its use and its integration with other aspects of clinical care – an aspect that distinguishes it from the established legislated-for areas. It would therefore be of particular significance that related informational obligations are not too demanding. Requiring written consent for the use of ML devices, for example, could jeopardise the integration of the technology into healthcare workflows and the benefits that would accompany this. Practicability in the AI context should mean a sufficiently general summary in circumstances where substantial interferences with autonomy are possible.

As was seen in California and the UK regulation of clinical trials, a suitable means of accommodating these concerns would be mandated discursive engagements with the patient or the provision of wider decision-aids (written summaries). Points of contact where the individual can receive further information, if interested, can also be provided. In these cases,

the patient's decision-making is facilitated and they have the choice to proactively engage with the provided aids.

Conversely, when it comes to the issue of enforcement where information-facilitation procedures have failed, the two jurisdictions have further provided instructive precedents. Given the sometimes minor and potentially frequent infringements of consent requirements through AI deployment, it may be desirable to establish standardised awards, as exist under UK common law or in the statutory regulation of sperm and oval removal in California. As in the latter case, these could then also be targeted specifically at repeat offenders, without depriving patients of their existing civil remedies.

Indeed, as stated in the introduction of this section, it is regarded as desirable that, in so far as is possible, the statutory interventions remain supplementary to the common law, which will provide the broader informed consent framework. Pairing the autonomy principle with the precise norms of informed consent does not offer a perfect response to clinical ML devices in either jurisdiction. Nevertheless, it provides several, potentially adaptable, avenues for argumentation that offer promising starting points. In case of doubt as to the applicability of a specific legislative framework for AI devices, or as to the effect of a bespoke statutory scheme, it should be made clear that the common law remains an applicable, supporting instrument.¹⁴¹⁰

Likewise, the creation of tailored informed consent mechanisms would be aided by embedding them into a wider regulatory framework for clinical ML. This thesis has not purported to deal with many other problems associated with the technology that could inspire such a step, problems relating to medical device regulation, data protection, discrimination and more. However, in so far as AI, especially medical AI, is legislated for regarding one or more of these areas, it would create a real impetus, or even a need, for the legislature to confront the matter of patient consent. Political and legal actors –within, but also far beyond, the UK and the U.S. – have begun to engage with such reforms.¹⁴¹¹ The UK GDPR's explanation requirements

1410 Moses in Marchant, Allenby and Herkert, *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight: The Pacing Problem* (2011) 84-85.

1411 One broader example is provided by the recent statements of the German Ethics Council, *inter alia* on bias and informed consent: Deutscher Ethikrat, 'Mensch und Maschine – Herausforderungen durch Künstliche Intelligenz' (2023) 159-162. For the UK and the U.S. see Chapter 1, Section II.A.

for automated decision-making provide an early example establishing the desirability of this approach in one of our two jurisdictions.

If the legislature moves to provide a regime for the regulation of medical AI, then the arguments presented here consequently gain added force. In this process it should be ensured that patient's decision making, their fundamental interest in autonomy, is adequately facilitated. The creative contribution of common law reasoning offers valuable guidance for these purposes and it is envisioned that it will continue to perform a supporting role.

III. Legal reasoning and technological innovation

In the introduction of this work three prominent assumptions on the interrelationship between law and technological innovation were outlined and contested. These were: that legal reactions to technological change are to be judged by reference to extra-legal, policy-orientated standards; that technological progress is inevitably racing ahead of a largely inert legal system, and; that extra-legal regulatory modalities offer a better chance of realising pertinent instrumental objectives and keeping up with the regulatory needs of society. In light of the developed principle-based methodology and our assessment of its application to the autonomy problems presented by medical AI/ML devices, we are now at a stage to scrutinise these three assumptions on the basis of this concrete example.

A. Understated challenges of non-legal regulation

One argument that emerged from the literature was that direct regulatory interventions, such as altering the architecture of a program's code, provide an available and often preferable alternative to legal norms for the purposes of accomplishing various innovation-related policy goals. This aspect is evaluated first because it was an assumption that our work grappled with in Part I.

In the course of considering the nature of AI devices, and how they would realistically be integrated into their medical context, it became apparent that the functioning of the technology suffered from various shortcomings. These shortcomings included the black box nature of ML models – in particular a subclass of deep neural networks (DNNs) – and

the flawed nature of performance evaluations for the purposes of gauging the functioning of medical AI through standard clinical indicators.

Assessing whether and how these defects could be remedied through the design of the technology itself was an important prerequisite for understanding how a legal solution would function in relation to its factual subject matter. As such, the availability of non-legal alternatives assumed an orthodox position within the framework of the comparison. It was merely supplemented by a more dynamic, forward-looking perspective to anticipate whether technological solutions would be forthcoming in the near future. This matched the more general emphasis on anticipating the development of a technology that is still emerging and which has yet to be comprehensively implemented.

In taking this approach, it was examined whether it was realistic that the design or operation of medical AI/ML technology would be refashioned so as to avoid the various autonomy-related issues. This involved technological fixes that seemingly presented an immediate solution, as well as secondary legal interventions that would frame these fixes. For example, a regulation that would require developers to carry out, catalogue and respond to more suitable performance evaluations – i.e. ones which were tailored to ML – was considered as one possible solution. Similarly, the development of explainable or interpretable AI models was a direct way in which the technological architecture could be designed in order to accommodate the professionals' and, relatedly also, the patients' need for information.

A question that can now be asked is how such technological fixes are to be judged *vis-à-vis* the specific legal recommendations presented here: a combination of common law evolution and the prospect of a bespoke, supplemental statutory disclosure scheme. One can begin answering this question by engaging with the oft-cited benefits of technological solutions, above all the notion that technological fixes emerge with a rapidity that (better) tracks the speed with which corresponding technological challenges emerge.¹⁴¹² In the present scenario it was demonstrated that this was simply not the case.

As was noted in Chapters 2 and 3, considerable uncertainty persists concerning the best way(s) to evaluate ML performance, especially in the protean medical context. It is likely that an appropriate scheme has yet to be designed.

1412 See Chapter 1, Section II.C.3. in this regard.

With respect to explainable and interpretable ML methodologies, this flaw is even more strongly pronounced. These technologies are in a state of development, throwing up a whole host of complications that do not take one very much beyond the original problem. This is especially true for the, normatively more desirable, modality of interpretable AI. A rapid adoption and implementation of interpretable methods by designers of medical ML devices appears far-fetched at this moment.

Yet clinical AI's challenges to autonomy are already manifesting themselves. As a result, it is not inapposite to speak of a 'technological lag' or a 'phasing problem' for those seeking to regulate the challenges of innovation through extra-legal solutions. By contrast, the legal system offers an imperfect but normatively defensible approach that can be implemented now. Appropriately conceived, informed consent mechanisms are broadly applicable and will be relevant to many deployments of clinical ML devices.

A further factor advanced in favour of technological solutions was their direct, unmediated effect, which translated into the successful realisation of desired objectives. Within a certain subclass of cases, for instance where individuals are performing tasks exclusively through the operation of a software system, this aspect may hold true.¹⁴¹³ Yet, in the complex social relationships within which AI medical devices are embedded – cooperating with clinical professionals to facilitate patient decision-making – a design fix does not, and cannot hope to, have this effect.¹⁴¹⁴

Our normative analysis indicates that it remains an open question how improved performance evaluations and forms of interpretability would have to be utilised to further patient autonomy. Not all information is significant to an individual patient, too much will overwhelm them, and it depends on how the professional chooses to convey what they know.

Consequently, it is not enough for the law to frame the necessary code. To achieve desired, direct effects, behaviour guiding norms remain indispensable. In such situations the technological solution is, at most, the indirect facilitator of the more immediate, legal intervention which sets the appropriate bounds of disclosure. Once again one can identify a more complex interrelationship between law and technology, which departs from the

1413 Recall that a growing cyberspace was the primary context for such early claims: Reidenberg, 'Lex Informatica' (1997) 76(3) *Texas Law Review* p. 553; Lessig, *Code: Version 2.0* (2006); Berman, *Law and Society Approaches to Cyberspace* (2007).

1414 See: Chapter 2, Section IV. Consider especially the uncertain effects that different forms of architectural design ('nudges') had in this context.

narrative of law as a secondary facilitator of a better-suited technological fix.

This leads me to the final point, which relates to the subject matter of the next section. Even granting that technological solutions possess a rapidity and a direct, comprehensive effect that the law does not, it is questionable whether they are also targeted at the right ends and, crucially, how it is ensured that such a correct targeting is maintained. The immediacy of technological action is worth little, after all, if it does not actually engage with the relevant challenge.

The contrast between explainable and interpretable AI provides a good example of this fundamental point. Developers and commentators evidently understood and shared the concern for the preservation of user autonomy, but they perceived the underlying problem of opacity in several different ways. For many, especially those favouring explainability, it was evidently closely associated with the need to bolster user trust, since a lack of this hampered implementation. That the generated explanation was not truly revealing a model's functioning, or even exhibited manipulative tendencies, was not necessarily considered to be in conflict with the goal of autonomy enhancement. That goal could be stated at a sufficiently abstract level and it could be pursued in an uncontested fashion. This illustrates that those focussed on technology policy still do not provide a means for identifying the right ends in concrete situations. In the next section it is seen how the law does provide such guidance.

B. The significance of law's resistance to instrumentalisation

A further trend that was identified within the law and technology literature treated the legal system as an object that could be instrumentalised, in a discretionary manner and relatively comprehensively, in the pursuit of external ends. Most especially this was in the pursuit of policy objectives that were, rather broadly, associated with a desirable or effective regulation of innovation.¹⁴¹⁵ Our completed analysis relativises this assumption by appealing to the way in which the law has interacted with, and thereby transformed, a bioethical principle. Corresponding to our division between Part II and Part III of this work, this transformation had two noteworthy facets.

1415 See Chapter I, Section II.C.1.

Part II illustrated how the law is able to interact with a goal that is external to it and which may be normatively contestable. Yet, to operationalise this in legal reasoning, and to gauge how the legal system would react to its imperatives, one had to do much more than just take over, or put into effect, an external value. That value had to be aligned with the structure and content of the existing system and it had to be transformed by reference to this.

More precisely, our evaluation rejected – by focussing on the function of a legal concept – argumentation that appealed directly to external values. Such an appeal did not account for the structured manner in which legal norms interacted with each other and how they were subject to contestation and revision (the topic of the next section). Nor did it take seriously the contextual limitations that an appeal to any individual objective must be subjected to, whether these stemmed from other principles, rules or consequentialist modes of reasoning.

Only once one had transformed the extra-legal concept into a legal one, could these structural components be satisfactorily captured. In the present case, it was only by transforming the value of autonomy into a legal principle that its operation within the law could be properly accounted for. Any external policymaker seeking to utilise the legal system for their instrumental ends would do well to bear in mind this structural component and the broader, multi-faceted way in which the legal system reacts to the introduction of external norms. A favoured objective will have to be reconciled with the existing commitments that will continue to exert their force to achieve a certain balance of interests.

Admittedly, this aspect is most pertinent to an analysis focussing on the common law mode of reasoning. Nonetheless, it relates to a lesser extent also to other modalities, such as statutory interpretation. In our jurisdictions this application was explored in relation to UK GDPR. Hereby it was evident that the autonomy principle could be invoked to concretise the demands of several open-textured requirements. A similar insight also accompanied the examination of existent informed consent schemes in both California and the UK. One could deduce common autonomy-related rationales and a propensity for incremental development based partly on the existing common law.

Where the law is receptive to a given conception of a value that is to be advanced, such as procedural autonomy, it has been contended that the scope within which that value is utilised is a further important factor. It is a significant consideration that the legal system, while striving to achieve

coherence at a general level, additionally identifies areas or subject matters in which special concerns can be accommodated. This fosters coherence in a distinct field. In Part II it was apparent, especially in the UK, but to a lesser extent also in the American legal system, that patient autonomy was fashioned according to specific considerations that did not apply, for instance, to the autonomy concept at play in disputes of property or contract law. Medical law provided its own reference points, allowing for a more targeted, richer analysis.

On the instrumentalist view this compartmentalisation is an anachronism, hindering the development of a coherent approach to technology regulation.¹⁴¹⁶ But again, if these theorists and associated policymakers seek to insert their favoured interests into existing legal frameworks, it is surely indispensable that they account for these fields – whether or not they then seek to move beyond them.

In addition, it is maintained here that, without regarding these structures, one impoverishes the quality of the overall normative analysis. Arguably, one reason why bioethics and law have enjoyed a productive interrelationship, above all regarding the protection of patient autonomy, is that bioethicists have paid close attention to the specific legal mechanisms that have evolved to shape the realisation of this value.¹⁴¹⁷ Conversely, it has also been important that the health law field, in building a coherent set of norms, sought to adapt and transform bioethical concerns into legally cognisable forms.¹⁴¹⁸ In the course of this interaction a host of interests have been concretised that allow for, perhaps imperfect, but nevertheless significant forms of regulation.

This leads one to the final aspect that was exhibited, above all, in Part III. Considering legal reasoning to be transparently subject to external objectives would not only posit an overly malleable picture of the law – one which is likely to lead to misplaced, ineffective regulatory efforts. More than this, it would distort and jeopardise the substantial analytical resources provided by the outlined operation.

1416 Guihot, 'Coherence in Technology Law' (2019) 11(2) *Law, Innovation and Technology* p. 311, 319-321. See also the critical perspective of: Sommer, 'Against Cyberlaw' (2000) 15(3) *Berkeley Technology Law Journal* p. 1145, 1150-1151.

1417 Faden, King and Beauchamp, *A History and Theory of Informed Consent* (1986) 114-150.

1418 Wall in Phillips, Campos and Herring, *Philosophical Foundations of Medical Law* (2019).

Integrating a value such as patient autonomy into the law, examining its functioning and anticipating its impact offer a unique opportunity to develop a rich understanding of that value. It will require the identification of contested objectives, reflection upon preferable resolutions in concrete circumstances and the specification and revision of one's assumptions. Regulators who seek to invoke legal and extra-legal modalities to achieve their goals would lose these insights if the instrumental conception of the law were the only interpretation on offer.

One significant example that was cited in Section I. above and which underscores this point, relates to the differing concretisations of the procedural autonomy concept by UK and Californian courts – specifically regarding the standard of care in informed consent cases. Whereas both jurisdictions accounted for the necessary decisional and practical dimensions of our abstract autonomy concept, they then went on to realise them in diverging ways. The UK paired an independent, proactive perspective of the patient with a professional's supporting role. This generated a legal mechanism that was closely orientated towards the facilitation of their specific subjective needs. By contrast, Californian common law posited a relatively passive understanding of the patient, conceiving of them as broadly dependent on their physician's guidance. This led to an emphasis on objective criteria and on the need for legal certainty regarding the advice that must be provided by the professional. Not only does this generate insights into the possible specifications of procedural autonomy, it also exemplifies how different interests can be accommodated alongside relevant specifications and it further highlights the kinds of consequences that flow from the relevant choice.

Even in those instances where the operation of the common law appeared to entirely discount an otherwise promising normative approach, a productive interaction between abstract evaluation and specific legal argumentation could nevertheless be derived. For instance, the need to 'fit' AI's potential for informational manipulation under a Californian strand of case law that dealt with the non-therapeutic interests of human practitioners may have appeared to ignore the very novelty of the technology: that it was acting independently of human thought and mediation. However, this engagement directed the analysis towards concrete, pertinent considerations that maintained their instructive role for AI/ML devices. It was important to know that there was an existing, well-reasoned determination that an influence on patient decision-making did not have to be proven to have affected human judgment before it could become an appropriate subject

of censure. Further, it was significant that the presence of an overriding therapeutic objective could not justify the non-disclosure of entirely non-therapeutic influences.

Moreover, precisely the remaining limitations of the common law, in terms of monitoring and enforcement, pointed the way towards responses that appropriately targeted AI/ML functioning. The advantages of such a view were noted in this chapter, where it was considered how legislatures, although liberated from many specific constraints, do capitalise upon the guidance derived from the existent legal system's operation. It was noted how they in turn operate according to their own related rationales to offer complementary legal solutions to autonomy-related problems.

Perhaps these insights could be derived in an entirely abstract fashion, *via* a method that simply concentrates on understanding the nature of desirable ends and on anticipating their non-legal realisation. In this manner, as the technological instrumentalists appear to prefer, reliance on any particular system would be eschewed and their argument could remain aloof from doctrinal quirks. Yet, as the comparative lawyer knows, these quirks can constitute functional responses to complex problems. Learning how these things are in fact done, especially in highly dynamic and ever-changing environments, tends to yield unanticipated insights, rich accounts of the problem at hand and various possible methods for addressing them.

C. Law's nuanced normative dynamism

The final argument that was identified in Chapter 1 holds that the law is purely reactive to, and must lag problematically behind, the challenges thrown up by technological progress.¹⁴¹⁹ To the extent that this constitutes a purely negative assessment of the law's functioning, this work has already departed from it. The previous section has sought to show one underappreciated positive facet of the law's limited, structured adaptation in response to external demands. Still, it was a further task of this thesis to examine how legal adaptation proceeds – and can proceed – from a methodological standpoint. On the basis of the specific assessment of AI/ML autonomy

1419 Chapter 1, Section II.C.2.

challenges one must ask how the law's responsiveness ought to be evaluated in this doctrinal respect.¹⁴²⁰

To make this assessment, this work did not purport to operate solely within the confines of existing rules, nor did it restrict itself to a narrow casuistic analysis. Rather, through its normative approach it departed from the state-of-the-art examinations of legal informed consent requirements and medical ML devices.¹⁴²¹ More broadly, it marked a departure from the law and technology literature that tended to argue either in favour of more liberal or more conservative approaches to rule interpretation – depending primarily on the instrumental demands of particular situations. By detailing the role that legal principles can play in the common law's adjustment to technological progress, a much more nuanced picture has emerged. This demonstrates that it would be misleading to adhere to a bifurcated approach, one that conceives of law's dynamism in the face of innovation as stemming from either a greater or lesser responsiveness to external, instrumentalist demands.

Under the principle-based approach it remains the case that certain avenues for change remain altogether foreclosed – several examples have been marked in Tables 1 and 2. For instance, where there is a well-established line of cases articulating the limited nature of a given class of rules, argumentation based upon the autonomy principle cannot be expected to overcome legal inertia. The narrow requirements pertaining to consent under the Californian battery doctrine provide an illustration of such a circumstance. It is an entirely orthodox proposition to hold that the weight of countervailing rigid and specific norms outweighs appeals to principle in such situations.

A similar conclusion is to be reached in cases where there is a conspicuous absence of authority. Namely, where a certain finding would constitute an altogether untested extension of existing doctrine, lacking even incidental support. Again, taking the generally more restrictive Californian system

1420 Practical difficulties of adaptation – such as: judge-made solutions being dependent on the right kinds of cases being brought forward, or the legislature functioning *via* relatively cumbersome information gathering and decision-making procedures – must be distinguished in this respect. The crucial question for present purposes is whether the law provides the conceptual tools to develop itself in the face of change. For another statement of this distinction, but addressing primarily the institutional component, see: Ard, 'Making Sense of Legal Disruption' (2022) *Forward Wisconsin Law Review* p. 42.

1421 Recall especially: Cohen, 'Informed Consent and Medical Artificial Intelligence' (2020) 108(6) *The Georgetown Law Journal* p. 1425.

as the example; arguing that actionable damage under negligence could be extended to autonomy harms *per se*, and that this would be an available adaptation to the pressures of AI's autonomy challenges, would be to misunderstand the limitations of principle-based reasoning. The autonomy principle itself is not capable of generating a cause of action and here there are significant countervailing concerns, albeit they are more abstract. They include above all: the principle of legal certainty and the consequentialist considerations that accompany it. In this situation the distinct considerations raised by medical AI are highly unlikely to call for a reassessment of these significant factors.

While this first class of cases may be described as pursuing a conservative approach, limiting the law's receptivity to desirable change, this would be to miss the point. It is a context-sensitive assessment that must be responsive, first and foremost, to the existent legal material. If this truly represents a strong constraint on reasoning, then arguing for a liberal or conservative mode of argumentation is of little avail.

Conversely, it was seen that there are instances where the emergence of medical ML technologies can be expected to exert strong pressures on the courts to review the state of existing doctrines in light of the demands of the autonomy norm. Under our principle-based analysis the relevant reconfiguration could take the form of novel interpretations of existing standards, of norm creation, of norm change or of exceptions that are made to otherwise unaffected rules. This too presents a more proactive and varied approach than is generally assumed in the wider literature.

Nevertheless, in the preceding examination, clear-cut illustrations of legal dynamism – demarcated in the above tables – also constituted a relatively small subset of the examined material. They emerged most prominently in relation to the interpretation of the breach element of negligence. For example, in both the UK and California a strong argument could be made that a requirement existed to disclose AI's ability to foreclose certain decision-making opportunities – even though this represented a novel phenomenon stemming from the technology's more independent functioning. Here, an extrapolation from the existing doctrinal framework was facilitated by a relatively open standard of care, which invoked autonomy-based reasoning, as well as by more specific rules that had been laid down in the case law. In particular, the courts had, by requiring the disclosure of available and reasonable alternatives, elaborated upon the significance of the reflective dimension of autonomy that also motivated the identification of the pertinent AI problem. It is these structural factors that pushed for a

liberal approach to emerge in this context, not a desire to realise an external objective.

In between these clear-cut opportunities for a dynamic adaptation of the law and those occasions where a rigid adherence to the existing rules was the most likely outcome, there was a third category of case. These were instances where it appeared likely that the introduction of AI/ML devices would require a reconsideration of existing norms, but the strength of available arguments was uncertain.

Under UK law, for example, it was important to consider whether interferences with individual autonomy constituted an actionable form of damage for the purposes of negligence, as this would allow the system to respond coherently to medical AI's challenges. This required an argument for norm generation, or at least adaptation, to be advanced. Although several significant judgments supported such a development, it could not be definitively asserted in light of prominent and persistent criticisms of courts and commentators. The latter were concerned, above all, with the doctrinal limitations of the tort of negligence and with the force of the principles of legal coherence and certainty.

But once more, rather than simply advocating for a liberal, instrumental approach that transcended these doctrinal concerns, a realistic argument in favour of legal dynamism required one to confront them. It was here that the autonomy principle, which had been developed to capture both AI's challenges and to fit within the legal system's specific requirements, became relevant. It was argued to be sufficiently concrete and objective to offer one acceptable conceptualisation of a cognisable form of damage. Therefore, by operating within the restrictions set by the pertinent norms, it was possible to make a structured argument that identified an opportunity for a proactive legal response to medical AI/ML.

This latter aspect – the way in which the dynamic or static nature of a system of norms is determined by the state of development of its more abstract principles – is the final manner in which this work has contributed to a more nuanced understanding of the law's adaptation to technological developments. Those who focus on specific rules appear to imply that the law is particularly prone to generate the 'pacing problem' because it is a cumbersome process to bring about the multifaceted, wide-ranging amendments that can be required by innovations that affect many distinct norms.

However, this does not account for the synergistic effects that a reconceptualisation of broader legal norms can similarly bring about. The poten-

tial for such an approach is evidenced by the impact that the autonomy principle has already had on the British negligence cause of action. Mutually reinforcing developments have spanned across specific norms, which encompass those pertaining to: actionable damage, the standard of care and causation. This demonstrates how the reconceptualisation of the legal system's higher normative commitments can have widespread knock-on effects on the shape of specific rules.

In relation to medical AI this means that, in a system such as the UK where principles pertinent to the regulation of the technology are experiencing a degree of flux, a greater adaptability to the demands of the innovation can be expected. The law is more open to mutually reinforcing adaptations, although knock-on effects on specific rules may still need to be worked out and the challenges posed by the technology may even help clarify the shape that these should take. Things are admittedly quite different in a more settled system, such as California, where a relatively stable balance between specific norms and abstract standards has been established.

Overall, it is to be noted that the law's propensity for development cannot be distilled into a simple formula aiming at the greater or lesser realisation of innovation-related objectives – it is intimately connected with the law's operation and the context-specific restrictions that this imposes. Accordingly, this work has subscribed to a nuanced differentiation that tracks the manner in which abstract and concrete norms interact to create space for creative legal activity. In spite of this context-specific focus, a more general insight can also be gleaned from the examination of medical AI under this lens: the potential for legal dynamism, at least at a doctrinal level, is much more extensive than is often recognised.

IV. Conclusion

Beyond evaluating the results of the common law comparison and outlining the lessons that can be learned from this for the broader regulation of medical ML devices, this final chapter has emphasised how the specific investigation relates to a much more fundamental discussion concerning the relationship between law and technological innovation. It has been demonstrated how the developed methodology provides a differentiated account that recognises the unique and persistent contribution of legal reasoning to this field.

To this end, it is necessary to question three prevalent assumptions concerning: the primacy of instrumental objectives, the relative inertia of the legal modality and an associated preference for extra-legal regulatory instruments. The preceding evaluation reveals that technological instruments may themselves lag behind, and be less effective than, available legal solutions. Moreover, the law continues to provide a multifaceted normative framework that assists in conceptualising the challenges presented by innovation, as well as indicating the nature of adequate solutions to these challenges. Lastly, in considering the law's ability to adapt dynamically one should take a nuanced approach that accounts for the pervasive role played by legal principles. It is the tensions caused by them that shape the creative, proactive role that legal reasoning assumes in accommodating the law to technological progress.

Returning to the specific research question posed in Chapter 1, the foregoing analysis demonstrates that, although the doctrines of informed consent may not themselves provide an adequate level of protection for patient autonomy as artificial intelligence is introduced into medicine, they nevertheless provide an available and inherently adaptable mechanism. The common law is in a position to respond, albeit imperfectly, to a rapidly emerging set of challenges. In addition, the mode of reasoning that characterises its operation points the way towards a coherent, more comprehensive solution. It has been suggested that this should take the form of a supplementary legislative scheme, functioning alongside the more general conditions of the common law and continuing to accommodate the evolving demands of the principle of patient autonomy.