Accepted Version

Cognitive Bioprediction and Moral Responsibility: Foreseeability, Control and the Ethics of Anticipated Cognitive Impairment

Sinead Prince[1]*, Jamie Walvisch[2], Matthew Baum[3,4], Tasmia Hai[5,6], Daniele Pacaud[7], Julian Savulescu[1,8,9], Fergus Cameron[9]

[1] Centre for Biomedical Ethics, Yong Loo Lin School of Medicine, National University of Singapore, Singapore

[2] UWA Law School, University of Western Australia, Perth, Australia

[3] Faculty of Psychiatry, Brigham and Women's Hospital, Boston, Massachusetts

[4] Center for Bioethics, Harvard Medical School, Harvard University, Boston, Massachusetts

[5] Children's Hospital Research Institute of Manitoba, University of Manitoba, Manitoba, Canada

[6] Department of Educational and Counselling Psychology, McGill University, Montréal, Quebec, Canada

[7] Department of Pediatrics, Cumming School of Medicine, University of Calgary, Alberta, Canada

[8] Murdoch Children's Research Institute, Melbourne, Victoria, Australia

[9] Oxford Uehiro Centre for Practical Ethics, Faculty of Philosophy, University of Oxford, Oxford, UK

[10] The Royal Children's Hospital, Melbourne, Victoria, Australia

*Corresponding Author: Julian Savulescu: julian.savulescu@uehiro.ox.ac.uk

This is the accepted version of a paper that has since been published in AJOB Neuroscience and can be cited as:

Sinead Prince, Jamie Walvisch, Matthew Baum, Tasmia Hai, Daniele Pacaud, Julian Savulescu & Fergus Cameron (25 Jun 2026): Cognitive Bioprediction and Moral Responsibility: Foreseeability, Control and the Ethics of Anticipated Cognitive Impairment, A JOB Neuroscience, DOI: 10.1080/21507740.2026.2686620

Abstract

Bioprediction uses biomedical data, biomarkers and algorithmic tools to forecast future physiological states or impairments. As predictive technologies increasingly enable individuals to anticipate loss of cognitive capacity before it occurs, they raise a fundamental ethical question: how should moral responsibility be assessed when impairment becomes foreseeable in advance of its onset? We argue that biopredictive technologies can alter both the epistemic and control conditions that ground moral responsibility. In some cases, this generates obligations to respond to predictive warnings of cognitive decline; in more limited circumstances, it may also ground obligations to acquire and use biopredictive technologies where these represent a proportionate and least restrictive means of risk reduction. We further argue that emerging individual responsibilities generate corresponding obligations for states. Where access to biopredictive technologies is necessary to meet morally significant expectations of risk management, fairness and legitimacy require attention to affordability, accessibility and institutional support.

INTRODUCTION

Bioprediction uses biomedical data, biomarkers, and algorithms to forecast future physiological states or impairments (Baum 2016). Predictive genetic testing is one familiar example, but bioprediction now extends well beyond genetics to technologies that forecast seizure risk, hypoglycemia and other episodes of cognitive impairment. Biomarkers, machine learning (ML), and artificial intelligence (AI) are increasingly being used to generate personalized forecasts of events that may affect cognition, awareness, judgment, attention or executive functioning. This development is ethically significant because some forms of cognitive impairment have traditionally excused individuals from moral responsibility as awareness and agency are lost at the relevant moment. Bioprediction complicates that picture.

What distinguishes contemporary cognitive bioprediction from ordinary awareness of background medical risk is not simply improved accuracy, but the provision of temporally specific, individualized and actionable foresight. Some emerging technologies can identify periods of heightened risk while individuals still retain the cognitive capacity to understand the warning and respond to it.[1] This creates a meaningful separation between awareness of impending impairment and the later loss of control. That temporal shift may generate new opportunities for proportionate precaution, and so alter how moral responsibility is assessed.

The practical significance of this issue is already apparent. In epilepsy, seizure-forecasting systems use physiological and behavioral data to identify periods of elevated seizure risk, sometimes in advance of onset (Brinkmann et al. 2021; Aziz et al. 2025). Sub-scalp electroencephalography (EEG) devices (electrodes between the skull and scalp), when paired with smartphones and machine learning algorithms have been used to identify pre-ictal signals based on individual data and to predict seizures up to 10 and 20 minutes in advance of onset (Moghim and Corne 2014; Ra et al. 2023). Other studies propose seizure-forecasting algorithms that use biomarkers such as stress levels, heart rate and sleep quality to identify circadian cycles of seizure patterns (Stirling et al. 2021; Costa et al. 2024). These technologies allow people with epilepsy to *proactively* administer medication, increase stimulation from an implanted device, move to a safe location, alert care givers, and avoid hazardous activities. In diabetes, predictive continuous glucose monitoring (CGM) systems can warn users of impending hypoglycemia at a point when meaningful action remains possible.[2] These systems

[1] We use 'cognitive capacity' in a broader sense than the technical clinical concept of decision-making capacity. The term refers here to the cognitive and executive functions relevant to safe and competent agency, and is used because the paper concerns gradually emerging impairments that may affect practical agency before formal legal or clinical thresholds of incapacity are met.

[2] Hypoglycemia is generally defined as a blood glucose level below 4.0 mmol/L (Yale et al. 2018). It can impair cognition and motor control, including executive functioning, information processing, attention, memory, and visuospatial skills (Warren et al. 2007).

transmit glucose data to smartphone applications that use algorithmic modeling, often together with user-entered information such as food intake or insulin dosage, to forecast future glucose levels (Huang et al. 2025; Petrovski et al. 2023; Corbett et al. 2022; Boughton et al. 2026).[3] They may therefore enable preventive stepssuch as administering medication, adjusting insulin, moving to a safe location, alerting caregivers, or avoiding hazardous activities, including driving. Comparable predictive tools are also emerging in other domains, including psychiatric contexts, where they may forecast relapse or other forms of impairment with implications for responsibility and regulation (Cockerill 2020; Lee et al. 2025).

These developments raise a fundamental ethical question: how should moral responsibility be assessed when future cognitive impairment becomes foreseeable before it occurs? Cognitive biopredictive technologies (CogBTs) – algorithms that use biomarkers to forecast risks of cognitive impairment – matter ethically because they may alter both the epistemic and control conditions of moral responsibility. Where a person knows, or reasonably should know, that a predictive signal indicates a substantial risk of impairment that could harm others,[4] responsibility may attach not to conduct during incapacity itself, but to an earlier failure to take reasonable precautions.

This article makes three related contributions to the literature on responsibility, risk and bioprediction. First, it develops an account of moral responsibility for anticipated cognitive impairment that builds on, but goes beyond, existing responsibility-based approaches to risk management, including Matthew Baum’s (2016) risk-banding framework. Baum’s account helpfully explains why obligations may vary with the magnitude, likelihood, and foreseeability of the risks a person poses to others. Our claim is that CogBTs introduce a further and underexplored feature: they can change not only the degree of foreseeable risk, but the timing and structure of responsibility itself by creating temporally specific opportunities for precaution before impairment occurs. Second, we distinguish several different ways in which responsibility may be engaged in this setting, including failures to check available biomarkers, failures to respond to present warnings, failures to respond to predicted future risk, and failures to maintain or appropriately use predictive systems already adopted. Third, we argue that where access to CogBTs becomes necessary for individuals to meet morally significant expectations of risk management, corresponding questions of fairness, affordability and state support arise. Our claim is therefore not simply that CogBTs increase foreseeability, but that they reshape the timing, structure and distribution of responsibility in ways that have both interpersonal and regulatory significance.

[3] For example, US-based AI developer One Drop has developed a digital therapeutics platform that can predict hypoglycemia with accuracy of at least 94% at 4-hour intervals (Imrisek et al. 2022; Bagley 2020).
[4] Discussions of the impact of cognitive bioprediction and the risk of harm to solely oneself are outside the scope of this paper.

The discussion proceeds in three stages: it first examines how bioprediction affects the epistemic and control conditions of responsibility, then the individual obligations that may follow, and finally the implications for states, law and regulation. We develop this argument through a case study of diabetes management and the moral responsibility of people with diabetes when driving. Diabetes provides a particularly useful illustration because hypoglycemia can cause sudden and serious cognitive impairment, predictive technologies are already in widespread use, and driving creates obvious risks to others.

This case study is used not because diabetes is uniquely important, but because it offers a clear and already technologically mature instance of a broader ethical problem. As predictive tools increasingly make temporary impairments of cognition, awareness, judgment or self-control foreseeable in advance, questions about responsibility, precaution and fair access will arise across a widening range of contexts, including epilepsy, psychiatric relapse, and other conditions in which agency may be compromised episodically rather than permanently. The significance of cognitive bioprediction therefore extends beyond any single condition or regulatory setting. It concerns the ethical implications of a wider shift toward technologies that can render future impairments actionable while control is still retained, thereby reshaping how responsibility is allocated in medicine, law and public policy. The normative framework developed here is intended to be general, although the principal examples are drawn primarily from Australian, United Kingdom, and North American contexts, where the relevant regulatory and technological issues are especially well documented.[5]

Understanding how moral responsibility shifts in response to bioprediction is increasingly urgent. As predictive technologies become more accurate, more widely used, and more deeply embedded in legal, regulatory and clinical decision-making, ethical frameworks that treat medical incapacity as wholly exculpatory will become harder to sustain. The challenge is to recognize the moral significance of foresight without unfairly burdening those already disadvantaged by illness, cost, or unequal access to technology. This article develops one such framework.

MORAL RESPONSIBILITY

Moral responsibility is commonly understood to depend on two core conditions: an *epistemic* condition, concerning what an agent knows or should know about the risks associated with their conduct, and a *control* condition, concerning whether the agent has the capability to act otherwise or to take reasonable precautions (Talbert 2024). Where either condition is absent, there is typically no responsibility. Harms that occur

[5] This matters because access to CogBTs, and the consequences of failing to use them, depend heavily on local arrangements for healthcare funding, insurance coverage, driver licensing, and fitness-to-drive regulation. An expectation that individuals acquire or use such technologies may be more realistic, and more normatively defensible, in systems where they are subsidized or widely available than in systems where access is costly or uneven.

during episodes of a lack of cognitive capacity are therefore often treated as excusable because the agent lacks control and epistemic awareness at the moment the harm occurs. A person who loses consciousness, experiences a seizure, or becomes cognitively impaired while engaged in an otherwise ordinary activity may be causally responsible for resulting harm, but not morally blameworthy if they did not foresee, could not have reasonably foreseen, or could not reasonably have prevented, the loss of control. In such cases, both the epistemic and control conditions fail at the time of the harmful act.

However, responsibility need not be assessed solely at the moment harm occurs. On "tracing" accounts of responsibility, an agent may be held responsible for harm caused during an involuntary episode where that harm can be traced back to an earlier negligent act or omission undertaken while the agent retained the relevant awareness and control (Zimmerman 1986; Fischer and Ravizza 1998). For instance, we generally hold people who cause harm whilst driving under the influence of alcohol responsible for their lack of cognitive capacity. This structure is often captured by the idea of a benighting act: an earlier voluntary act or omission that foreseeably results in a later state of ignorance or incapacity, and that grounds responsibility for subsequent harm (Smith 1983). This approach has a clear analogue in law, where responsibility may be grounded in the doctrine of prior fault, such that an individual may be held criminally or civilly liable for harm occurring during a state of incapacity if that incapacity was itself foreseeably brought about by earlier culpable conduct (Child 2014).

Tracing accounts can be applied to the *epistemic* condition through the idea of "culpable ignorance". Ignorance excuses only when it is non-culpable – that is, when an agent lacks relevant knowledge because they had no reasonable opportunity or reason to acquire it (FitzPatrick 2008). By contrast, ignorance may ground responsibility when an agent fails to inform themselves of relevant risks despite having reason to do so and a reasonable opportunity to acquire such information. For instance, a person with epilepsy may be morally responsible for harm when their ignorance of their increased risk of harm whilst driving can be traced to their refusal to attend a doctor after the first epileptic seizure. In such cases, responsibility is grounded not in what the agent knew at the moment of harm, but in their earlier failure to acquire or attend to information that would have made the risk foreseeable. Benighting acts therefore operate both epistemically and temporally: they explain how responsibility can attach despite later ignorance or loss of capacity.[6]

The control condition can be similarly cross-temporal. Control need not be exercised at the moment harm occurs, provided the agent had a meaningful opportunity to exercise control at an earlier point. As Glannon (1998) argues, advance physiological knowledge can enable agents to act anticipatorily, allowing responsibility to be grounded in earlier

[6] There have been several criticisms of tracing accounts, notably those by Levy (2009) and Vargas (2005).

We address these criticisms below in the section titled ‘Objections’.

decisions rather than later incapacity. What matters is whether the agent had a reasonable opportunity to respond to foreseeable risk while retaining cognitive capacity, not whether control was preserved throughout the entire causal sequence.

Importantly, neither the epistemic nor the control conditions of moral responsibility depend on perfect foresight or complete mastery over outcomes. Agents are routinely held responsible for managing probabilistic risks, provided those risks are sufficiently foreseeable and the burdens of precaution are reasonable (Baum 2016). As we will argue below, what is reasonable to expect is tied to what is proportional to achieve, taking into consideration the magnitude and likelihood of the risk, the severity of potential harm, the costs of precaution, the reliability of the prediction, and the agent's capability to act at the relevant time. Crucially, the standard for what an agent ought reasonably to know or do is objective: it turns on whether there were sufficient reasons for precaution given the risks involved rather than whether the person actually perceived such reasons (Sher 2009; Parfit 2011). When the potential harm to others is grave, the threshold for what an agent ought reasonably to know or do may be correspondingly higher.

These features of moral responsibility provide the framework within which CogBTs become ethically significant. By expanding what agents can reasonably be expected to know about their own future states, and by identifying points at which preventive action remains possible, CogBTs may affect how responsibility is assessed in cases of involuntary impairment. The sections that follow examine how this framework applies in practice, and what obligations it may generate for individuals and states in an increasingly biopredictive world.

BIOPREDICTION AND THE CONDITIONS OF RESPONSIBILITY

Biopredictive technologies can alter what agents can reasonably be expected to know about their future states and when they can meaningfully act on that knowledge: they fundamentally change how the epistemic and control conditions of moral responsibility are assessed.

Unlike retrospective medical assessment, bioprediction provides forward-looking information about future impairment, often while individuals still retain cognitive capacity. In this way, it can make risks that might otherwise remain diffuse or uncertain more salient and actionable. For instance, when CogBT alerts indicate a time-specific and substantial likelihood of future cognitive impairment that could foreseeably harm others, ignorance of that risk is less likely to be non-culpable. The control condition can also be affected: alerts can enable a person to take preventive action while their cognitive capacity remains intact. Even where impairment later undermines control entirely, predictive information may create a meaningful temporal separation between awareness of risk and loss of capacity.

Importantly, not all predictive technologies meaningfully enhance foreseeability or preserve control. Inaccurate, opaque, or poorly timed predictions may fail to generate genuine reasons for action, and in such cases responsibility remains limited. Predictive information grounds responsibility only where it is reasonably reliable, intelligible to the user, and provided at a point when proportionate preventive action remains possible. Some machine learning algorithms may increasingly satisfy these criteria in particular domains by generating more specific and temporally actionable predictions of elevated risk. The diabetes case study that follows illustrates how these assessment shifts of the epistemic and control conditions may generate anticipatory responsibilities in practice.

A CASE STUDY: THE DAYLESFORD CASE

In 2023, a 67-year-old man with long-standing type 1 diabetes drove into an outdoor dining area at the Royal Daylesford Hotel in Victoria, killing five people and injuring several others (Dalton 2024). At his committal hearing,[7] the prosecution contended that he remained responsible because his CGM had issued multiple low blood glucose level (BGL) alerts before the crash and he nevertheless continued driving. The prosecution argued that he knew of the risk because he tested his BGL at 4:00 pm, recording 7.2 mmol/L, and again at 5:19 pm, recording 2.9 mmol/L. After that second reading his phone received a “dismiss alarm check glucose” alert, followed by eight further alerts at five-minute intervals (Ore 2023). The man attempted to purchase food from a deli but was turned away because the venue was full. He resumed driving at 5:36 pm and crashed at approximately 6:07 pm.

His defense counsel argued that he was not criminally responsible because he was experiencing a severe hypoglycemic episode at the time of the collision – his BGL was less than 1.1 mmol/L – and therefore lacked the capacity for voluntary action. The defense contended that there was no evidence that he had perceived, comprehended, or interacted with the phone notifications in the 24 minutes before the final driving episode. This evidentiary uncertainty, together with medical evidence of the rapid decline in BGL and the lack of swerving or brake marks, led the magistrate to conclude that the prosecution could not prove a voluntary act beyond reasonable doubt, and all 14 charges were dismissed.

Importantly, the man’s CGM did not provide the kind of predictive warning on which this paper focuses. CGMs, and earlier forms of glucose monitoring such as finger-prick testing, increase epistemic access to an individual’s *current* physiological state but only provide limited insight into future impairment. Such technologies may alert individuals to *present* hypoglycemia, from which future risk may sometimes be inferred, but they do not reliably predict when cognitive capacity will deteriorate or create a genuine opportunity

[7] A committal hearing is a preliminary proceeding in which a magistrate determines whether the prosecution’s evidence is sufficient to justify sending an accused person to stand trial. If it is not, the

accused is discharged.

for precaution before control is lost, particularly given that impairment of driving capacity can begin before a person perceives it (Bell et al. 2010). As a result, in some cases a person may already lack cognitive capacity immediately prior to driving (when we reasonably expect them to test their BGL) or may lose that capacity while driving (when we reasonably expect them to cease driving). In such circumstances, the capacity required for responsibility may already be compromised by the time a warning appears. This aligns with legal doctrines concerning sudden medical incapacity: where impairment arises without warning before a meaningful opportunity for prevention, the conditions for responsibility are ordinarily not met.

Seen in this light, the Daylesford case illustrates not judicial reluctance to assign responsibility in principle, but the evidentiary and temporal limits of non-predictive monitoring. Responsibility could not be traced to an earlier culpable omission because the technology did not provide sufficiently predictive and actionable notice while agency was still preserved. By contrast, had the driver received a reliable warning of impending hypoglycemia at a point when he still retained the capacity to understand the risk and take preventive action, the analysis may have been different. The relevant question would then not simply have been whether he acted involuntarily at the moment of collision, but whether he had earlier failed to take reasonable precautions in response to a foreseeable and temporally specific risk. As predictive CGM and related CogBTs become more reliable, the traditional approach to assessing responsibility for harms that arise without warning and beyond the agent's reasonable control will therefore come under increasing pressure. The following discussion examines how genuinely predictive systems may alter the epistemic and control conditions that underpin responsibility in cases of hypoglycemia-related driving harm.

CogBTs and anticipatory responsibility

Using ML algorithms, CogBT devices can analyze real-time BGL data, together with user input regarding food and insulin intake, to identify periods of elevated risk before impairment occurs – sometimes within 20 minutes, sometimes over the next two hours, and sometimes overnight – while the individual still retains cognitive capacity (Georga et al. 2013; Huang et al. 2025; Petrovski et al. 2023; Corbett et al. 2022; Boughton et al. 2026; Breton et al. 2012). If such devices are sufficiently reliable (such as trialled and affirmed through standard clinical or regulatory approval processes), their use will alter both what agents can reasonably be expected to know about their future states and when meaningful preventive action remains possible. One example is Accu-Chek SmartGuide CGM, developed by Roche, which sends BGL values every five minutes to the corresponding app. The app incorporates this real-time data, together with other physiological measures, to detect glucose patterns and predict future glucose levels. It can indicate hypoglycemia risk within the next 30 minutes, forecast how glucose levels will develop over the next two hours, and estimate the risk of nocturnal hypoglycemia.

Published evaluations report high levels of clinical accuracy (Mader et al. 2024; Parkes et al. 2000), though the example is used here to illustrate the temporal structure of contemporary predictive CGM systems rather than to endorse the performance of any particular device.

An individual who is alerted to a high likelihood of impending hypoglycemia – and therefore cognitive impairment – in the next two hours can take reasonable action to reduce the risk and to limit its potential impact on others, such as by delaying or canceling driving, seeking assistance, or adjusting food intake or insulin dosage. In such cases, responsibility attaches not to conduct during impairment itself, but to an earlier failure to take reasonable precautions in response to a foreseeable risk of future cognitive impairment.

Importantly, this does not entail a duty to eliminate all risk. Predictive systems operate probabilistically, and false positives and false negatives remain possible. Responsibility therefore depends on proportionality. The magnitude and likelihood of the predicted risk (e.g., 30% chance of severe hypoglycemia), the severity of potential harm to others (e.g., lack of consciousness whilst driving) the reliability and intelligibility of the warning (e.g., verified device by regulatory bodies as accurate 95% of the time), and the burdens imposed by precaution all remain relevant (e.g., cost of acquiring and maintaining the device). This anticipatory account of responsibility is not unique to diabetes. Comparable obligations might arise where seizure-prediction systems indicate elevated and temporally-specific seizure risk later in the day, or where relapse-prediction tools signal imminent loss of impulse control in psychiatric contexts.

Automated systems and redistributed control

Closed-loop or automated insulin delivery systems introduce further complexity. These systems integrate predictive algorithms with an “artificial pancreas” that automatically delivers insulin on the basis of real-time data and ML pattern recognition, through a subcutaneous cannula that may be tethered or wireless. These systems require user training and remain expensive, but they have shown the capacity to reduce glucose variability and lower the frequency of hypoglycemic episodes (Breton et al. 2012; Turksoy et al. 2013; Boughton et al. 2026; Grosman et al. 2021; Petrovski et al. 2023).

These hybrid systems redistribute, rather than eliminate, control. Monitoring of BGL and routine insulin adjustment are partially delegated to automated processes, while users retain responsibility for activities such as announcing meals, system calibration, and responding to alerts or system prompts (He et al. 2026; Boughton et al. 2026). For instance, Tandem T: Slim x2 Control IQ from Tandem Diabetes Care can predict glucose levels 30 minutes in advance, which has been found to significantly increase the amount of time participants spent within prespecified BGL ranges (Ekhlaspour et al. 2021). As a result, responsibility is not displaced wholesale but reallocated across human and technological components of the cognitive process.

Some recently developed systems are entirely closed-loop; they do not require manual input of meal or physical exercise data (Rashid et al. 2026; Templer 2022). For example, CamDiab has received FDA approval in January 2026 for the CamAPS FX fully closed loop insulin delivery system with research suggesting the algorithm, on average, improves time spent in predetermined blood glucose ranges and reduces time spent in hypoglycemia (Boughton et al. 2026).

In cases when an individual reasonably relies on a validated, properly functioning automated system, and harm arises from unforeseeable system failure, responsibility may be mitigated or partially displaced because the opportunity for reasonable intervention was reduced or removed. By contrast, responsibility may remain traceable when individuals disable alerts, neglect maintenance, ignore known system limitations, or continue to rely on automation despite clear warning signs. In such cases, responsibility again attaches to earlier capacity-retaining choices concerning design, maintenance, reliance, and response to prompts rather than to later impairment itself. Automation therefore complicates, but does not negate, anticipatory responsibility.

INDIVIDUAL OBLIGATIONS TO RESPOND TO BIOMARKERS

The case study shows that predictive technologies may generate several different kinds of individual obligation by making risk information available before impairment occurs. Similar questions are likely to arise wherever tools forecast epileptic seizures, cardiac events, psychiatric relapse or other forms of impairment that may place others at risk. More broadly, as CogBTs mature, responsibility for harm arising during involuntary impairment may increasingly turn on whether individuals responded reasonably to biomarker information while they still retained meaningful practical control. The discussion below distinguishes four related ways in which responsibility may arise in response to biomarker information.

First, an individual may *fail to check an available and relevant biomarker* where monitoring is reasonable in the circumstances. When checking is simple, low burden, and directly connected to preventing serious harm, failure to obtain readily accessible information can constitute a culpable omission. For instance, culpability may arise where a person with epilepsy fails to consult a readily available seizure-risk indicator before engaging in a hazardous activity, or where a cardiac patient ignores an accessible wearable alert before engaging in strenuous exertion. In such cases, responsibility for subsequent harm may be traced back to the earlier failure to check, provided the individual had the time, capability, and capacity to check their predictive device, not whether the individual perceived they had met such conditions. By contrast, if a biomarker comes to the agent's attention only once capacity is already compromised or if checking would not reasonably have prevented harm, responsibility does not arise.

Second, an individual may *fail to respond to a present biomarker that clearly indicates imminent risk*. Predictive information grounds responsibility only where it is reasonably reliable, actionable, intelligible to the user, and provided at a point when proportionate preventive action remains possible.[8] In these cases, the epistemic condition is satisfied because the risk is explicit, clear and salient. Where the individual also retains control, they may be expected to take immediate mitigating action. A driver who receives an alert indicating imminent hypoglycemia, for example, can self-treat and delay further travel. Similarly, a person who receives a real-time alert indicating extreme drowsiness or microsleep risk while operating a vehicle, or an oxygen-desaturation alarm in a high-risk work environment, may be expected to pause or withdraw from the activity.[9] Disregarding such information can ground moral responsibility and, in legal contexts, prior fault. By contrast, when awareness of the biomarker could only be identified concurrently with a significant loss of capacity, the control condition fails and responsibility will likely be absent.

Third, and most importantly for the purposes of CogBTs, an individual may *fail to respond to a predicted biomarker*. CogBTs alert users to a heightened likelihood of future impairment at a point when they still retain cognitive capacity and meaningful control over their actions. In these circumstances, both the epistemic and control conditions are satisfied in advance of the harmful event. In the diabetes context, a predictive CGM warning of impending hypoglycemia hours in advance may justify postponing or canceling a planned drive. Comparable obligations arise where a seizure-prediction system indicates elevated seizure risk later in the day. Ignoring a credible predictive warning in such circumstances may constitute a benighting act, grounding responsibility for subsequent harm in an earlier failure to take proportionate precautions.

Fourth, a distinct but related issue arises where an individual has already adopted a predictive system but *fails to maintain or use the system appropriately*. Where reliable technologies can autonomously manage risk, individuals may legitimately discharge some of their responsibilities through delegation. However, that delegation is conditional. Ignoring system alerts, neglecting calibration or update requirements, or

---

[8] Inaccurate, opaque, or poorly timed predictions may fail to enhance the user's epistemic position or practical capacity for action. Systems prone to excessive or low-specificity alerts may also generate "alarm fatigue" (Cvach 2012), reducing the salience of the warning and weakening the case for responsibility.

[9] Mitsubishi Electric announced a new AI-enhanced sensory technology that uses non-contact pulse, eye movement detected via camera, and vehicle control data (e.g., steering, acceleration) to identify alcohol intoxication and issue alerts to the driver or even take control of the vehicle by reducing speed. Those who ignore such alerts, or override car control, when intoxicated would be similarly morally responsible for any subsequent harm. However, the legal implications may be limited in jurisdictions where driving while intoxicated is already a strict liability offense regardless of warnings.

continuing to rely on a system despite known faults can constitute culpable omissions. In such cases, responsibility remains traceable to earlier choices concerning system use made with sufficient cognitive capacity, rather than to the later episode of impairment itself.

Across these cases, the unifying moral principle is proportionality. Obligations to check, respond to, or rely on biomarkers depend on the magnitude and likelihood of the risk, the severity of potential harm to others, the burden of precaution, and the agent's capacity at the relevant time. CogBTs do not create unlimited duties of vigilance. Rather, they help specify when ignorance becomes culpable and when loss of control ceases to be morally exculpatory. When predictive information is reliable, salient, and reasonably actionable, individuals may be expected to respond; when it is not, responsibility remains appropriately limited.

INDIVIDUAL OBLIGATIONS TO ACQUIRE BIOPREDICTIVE TECHNOLOGIES

Having established when individuals may be responsible for responding to predictive warnings, a further and more demanding question arises: are individuals ever morally obliged to *acquire* CogBTs? Unlike obligations to respond to existing information, obligations to acquire technology raise additional concerns about cost, access, and fairness. Risk levels vary widely across individuals and over time, and many people face significant financial, practical, or social barriers to acquiring CogBTs or restructuring their lives in ways that reduce risk, such as limiting driving or undergoing additional testing. The core normative issue is therefore not whether risk reduction is desirable, but under what conditions individuals may reasonably be expected to bear the burdens of achieving it.

To address this question, Baum (2016) proposes a graduated "risk-banding" approach, according to which the strength of an individual's obligations corresponds to the magnitude, likelihood, and foreseeability of the risk they pose to others. On Baum's view, individuals who know, or reasonably should know, that they face an elevated and recurrent risk of causing serious harm may bear stronger obligations to adopt available risk-mitigating measures than those whose risk is low, speculative, or transient. Crucially, risk is dynamic rather thanstatic. A person ordinarily regarded as low risk may temporarily enter a higher-risk category when predictive indicators signal imminent impairment, while a person classified as high risk may experience periods in which that risk is acceptably low. Obligations must therefore track both individual variation and temporal fluctuation in risk.

The diabetes context illustrates how Baum's framework might operate in practice. In some cases, individuals may have obligations to acquire CogBTs because additional biomarkers or clinical indicators identify them as posing an elevated and recurrent risk of future cognitive impairment. Research suggests that severe hypoglycemia-related driving incidents are highly concentrated within a relatively small subset of individuals with type 1 diabetes (Pedersen-Bjergaard et al. 2004; Bognetti et al. 1997; Rewers et al. 2002; Campbell et al. 2010). These higher-risk individuals can often be identified through clinical markers such as impaired awareness of hypoglycemia, prior severe episodes, or validated risk stratification tools (Cox et al. 2017; Karter et al. 2023; Matus et al. 2023).[10] For this group, the elevated risk of harm to others is not merely possible but reasonably foreseeable and recurrent. Where a person intends to engage in activities that pose significant risks of harm to others, and where effective CogBTs are available as the least restrictive means of substantially reducing that risk, a moral expectation to acquire and use them is more readily justified.[11]

By contrast, most people with type 2 diabetes, including many who use insulin, experience far fewer hypoglycemic episodes (Kovatchev et al. 2002) and do not appear to have elevated accident rates relative to the general driving population (Cox et al. 2003). For these individuals, the mere presence of a biomarker does not ground an obligation to acquire costly predictive technologies. This distinction illustrates a central feature of the argument: obligations to acquire CogBTs do not attach to diagnostic categories as such, but to demonstrable risk profiles that exceed socially tolerable thresholds.

The case is more complex for individuals whose risk is identifiable but comparatively low. New algorithmic tools can predict severe hypoglycemia even for people whose risk is not otherwise apparent (Tsichlaki et al. 2022; Giammarino et al. 2025; Cichosz et al. 2024). However, foreseeability alone cannot ground an obligation to acquire technology (Baum 2016). In a world of pervasive bioprediction, nearly everyone carries some non-zero risk of causing harm. Moral responsibility does not require agents to take on disproportionate burdens to mitigate highly unlikely risks. On Baum's risk-banding framework, obligations strengthen as both the probability and severity of harm increase, while remaining defeasible where risks are low and the costs of precaution are high.

Cost and burdens therefore matter. The financial burden of acquiring and maintaining CogBTs can be substantial, particularly in the absence of insurance coverage or public subsidy. Current unsubsidized predictive CGM costs may run into several thousand Australian or US dollars, depending on device type, jurisdiction, and insurance

[10] Cox et al (2017) identify a range of screening factors associated with increased driving risk among people with type 1 diabetes, including impaired awareness of hypoglycemia, a history of severe episodes or driving mishaps, extensive driving exposure, problematic management of hypoglycemia and related complications such as retinopathy or neuropathy.

[11] Baum (2016) notes that predictive systems can also serve as a "solution" by identifying low-risk windows for individuals who otherwise pose elevated risk. Predictive CGMs, for example, do not only warn users of impending danger but can also indicate periods in which driving is comparatively lower

risk, thereby enabling some high-risk individuals to meet their responsibilities without withdrawing from driving altogether.

arrangements.[12] These burdens may include not only upfront purchase costs, but also ongoing expenses associated with maintenance, calibration, software updates and replacement. Geographic and social context also matter. For individuals in rural or transport-poor regions, restricting or abandoning activities such as driving may impose severe social and economic costs not borne by those with access to alternative forms of transportation or support. Where the burden of acquiring and maintaining a technology is disproportionate to the level of risk posed, obligations to do so may be reasonably excused.

The principle of least restrictive means provides an important normative guide, defined here as the least harmful and least restrictive option that can reasonably reduce relevant risks (Baum and Savulescu 2013). Where an individual poses a sufficiently high and ongoing risk of serious harm to others, requiring the acquisition or use of a CogBT may be less restrictive, and more respectful of agency, than alternative measures that categorically prohibit participation in valued activities or roles. In many regulatory contexts, blanket exclusions are already imposed when health-related risks are deemed unmanageable, particularly in aviation, commercial vehicle operations, and military and law enforcement. If predictive technologies can reliably reduce risk while preserving conditional participation – such as continued engagement in work, mobility, or social life – requiring their use for higher-risk individuals may be morally preferable to more exclusionary alternatives. In such cases, an expectation that individuals bear significant, though not unlimited, costs (such as costs of predictive CGM devices) to mitigate foreseeable harm may be defensible.

A further concern is whether such obligations are *unfair* when compared with other tolerated sources of risk. Only a small proportion of serious harms in domains such as transport, employment, or public safety are attributable to identifiable medical conditions (Stork et al. 2006), and societies routinely tolerate many other risk-enhancing factors, such as youth, fatigue, or inexperience, without imposing comparable technological requirements. However, this does not undermine the legitimacy of differential obligations where individualized risk can be reliably identified and meaningfully mitigated. As Baum (2016) argues, probabilistic reasoning grounded in group-level and individual-level data can justify tailored obligations even where risks are not unique to a particular group. Biopredictive technologies refine, rather than transform, this familiar regulatory structure by enabling more targeted and proportionate risk management. Many social practices already operate on a conditional-access model, imposing additional requirements where risk exceeds acceptable thresholds ie., drivers

[12] Indicative consumer pricing suggests that unsubsidized predictive CGM costs may be between approximately AUD $3,000 and $5,000 per year in Australia and around USD $6,800 in the United States, although these figures vary substantially depending on device type, replacement frequency, insurance

coverage, and subsidy arrangements (Soronio 2025; Gardner 2025).

licensing regulations and medical assessments, firearm ownership and psychological assessments, and aged care or childcare criminal history checks.

In summary, individuals who pose a demonstrably elevated, foreseeable, and recurrent risk of serious harm to others bear moral obligations to acquire and use biopredictive technologies when doing so is a proportionate and effective means of risk reduction. These obligations are not absolute. They are constrained by cost, access, reliability and the availability of less burdensome alternatives.

Although the moral framework developed here is intended to apply across jurisdictions, its practical implications will vary depending on how healthcare access, insurance coverage, and regulatory design are organized in particular societies. The empirical illustrations used here are drawn primarily from Australian, UK, and North American contexts, where restrictions on driving with diabetes are well-documented and CogBTs are increasingly commercially available or subsidized for some groups. Affordability and accessibility differ significantly across health systems, and the corresponding state obligations identified below will therefore be more or less demanding depending on institutional setting. However, the underlying normative principles of proportionality, foreseeability, and the fairness of linking responsibility to access are not specific to any one system and are intended to inform policy deliberation across a range of institutional contexts.

OBLIGATIONS ON STATES TO SUBSIDIZE BIOPREDICTIVE TECHNOLOGIES

A further question is whether states have corresponding obligations to subsidize access to CogBTs. State obligations may arise from several sources, two of which are especially relevant here: fair equality of opportunity and principles of fairness.

First, state obligations to publicly finance health-relevant technologies can arise from the public benefits they generate and from fair equality of opportunity — grounds that apply regardless of whether wealthier individuals could self-fund access. Daniels (1985), drawing on Rawlsian ideas of justice, argues that access to healthcare is foundational to fair equality of opportunity. On this view, health conditions that restrict an individual's ability to participate in social, political, or economic life are not merely matters of personal misfortune, but matters of justice.

This reasoning has particular force where access to CogBTs affects whether individuals can meet their moral obligations to others without sacrificing basic opportunities or social participation. If people are expected to manage foreseeable risks through technologies that are unavailable or unaffordable to them, holding them responsible for failing to do so becomes harder to justify. Baum (2016) similarly argues that access to predictive tools is especially important for those who are worst off, either because they face elevated physiological risk or because the consequences of harm and exclusion fall disproportionately on them.

Arguably, these justice-based considerations become especially important where CogBTs function as the least restrictive means of managing risk. Requiring the acquisition of a predictive technology may, in some cases, be more respectful of agency than categorical exclusion from valued activities or roles, such as access to entire categories of employment or to essential forms of mobility in regions without accessible or affordable public transport. In such cases, prioritarian principles of distributive justice may require state support for technologies that enable individuals to manage foreseeable risks without bearing disproportionate personal burdens. The concern is therefore not simply that some people cannot pay, but that unequal access to risk-management tools that condition social participation is itself unjust. On this view, state obligations to subsidize CogBTs need not be conditional on individual inability to afford them.

Second, state obligations may also arise from the risk of creating a two-tier system in which only the well-resourced can retain conditional participation while others are excluded. Although CogBTs may be cost-effective at a population level (Wan et al. 2018; Jiao et al. 2022), high upfront and ongoing costs continue to limit access in practice, particularly in the absence of insurance coverage or public subsidy. If access to CogBTs becomes a condition of meeting moral or legal expectations of risk management, inequalities in access may translate directly into inequalities in responsibility, opportunity and participation in socially important activities. In such circumstances, holding individuals responsible for failing to acquire technology risks compounding existing disadvantage rather than promoting fairness or safety. Where states require the use of CogBTs as a condition of participation in such activities, they incur corresponding obligations to subsidize access so as to avoid unjust exclusion from socially significant forms of participation.

OBJECTIONS

Reasonable Expectations

Much of our argument depends on the assumption that people can reasonably be expected to reduce foreseeable risks of future cognitive impairment and the harms that arise to others that may follow. However, several critics have highlighted the difficulty of distinguishing reasonable from unreasonable expectations about what agents can do to improve their epistemic position.

Rosen (2004) argues that we often lack justification for attributing ignorance to a clear-eyed case of akrasia – that is, knowingly failing to act on what one already recognizes as the right course. For example, when a person fails to undertake a relevant test of future risk, can we really be confident that they appreciated, at that moment, the risks they posed to others? Similarly, Vargas (2005) argues that closer inspection of the knowledge condition often reveals non-deliberative behaviors, such as dispositions, habits, or character traits, that contribute to harm but were themselves formed and maintained under conditions in which the person could not reasonably have foreseen their future consequences. The more cases we insist involve moral responsibility, Vargas suggests, the more we must relax the knowledge condition; otherwise, we must accept that responsibility is less common than we ordinarily assume.

Levy (2009) presses this concern further by arguing that in many cases we cannot reasonably expect people to remedy their normative ignorance. To hold them to such expectations, he argues, is to demand irrationality, and "it is not reasonable to blame people for failing to act irrationally" (p. 729).[13] A person may, for example, genuinely believe their BGL is fine and decide that a very small risk does not justify disruption to an ordinary activity. Since it is generally reasonable to tolerate some low-level risk in everyday life, we should be cautious about inferring culpability from every failure to test or reassess a risk (particularly daily risks).

One response is that individuals may nevertheless bear epistemic obligations to form *justified* beliefs. Robichaud (2014) argues that agents can be responsible for failures to take reasonable steps to improve their epistemic position. Strawson (1962) suggests that ignorance may be culpable where an agent's conduct reveals objectively inadequate regard for others.[14] FitzPatrick (2008) similarly argues that ignorance is culpable where a person deliberately keeps themselves in the dark, or fails to subject their beliefs to appropriate critical scrutiny despite having reason to do so. On this view, a failure to attend to morally salient risk may itself be culpable where the agent could reasonably have been expected to take corrective steps.

CogBTs strengthen this response in an important class of cases. Unlike the diffuse character-formation cases discussed by Vargas, or Rosen's concern about identifying clear-eyed akrasia, CogBTs can provide concrete, temporally proximate, and externally

---

[13] Levy argues that the sense of rationality is not whether the person has reason to avoid the act, but whether they have subjectively rational reasons to do so: "what it is rational for me to do—that is, what I can decide to do as the result of engaging in reasoning—is settled by my actual mental state" (p. 735). Levy defends this view of internal rationality as the threshold of culpability on the grounds that simply because we are challenged on our views, does not give us good reason to seriously reassess them – we rightly reject racist and Holocaust deniers without serious engagement with their views. Furthermore, even when we do seriously reassess our internal reasons, it is incredibly difficult to see them as mistaken. Given we can only blame actions for things we can reasonably expect them to have avoided, Levy argues that insofar as it is unreasonable to expect people to therefore change their moral views, they are not blameworthy: "Akrasia is, of course, a failure of practical rationality, and it is not reasonable to expect agents to do anything they can do only by way of a failure of practical rationality" (p. 739). Levy's objection is that culpability depends on what it is rational for the agent to do given their actual mental state, not simply on what there are objective reasons to do. On this view, if an agent could avoid the relevant ignorance only by acting against what is, for them, subjectively rational, then failure to correct that ignorance may not be blameworthy (Levy 2009, 735-739).

[14] Aquinas (*Summa Theologiae*, I–II, Q 6, art. 8) similarly suggests that ignorance can be culpable where a person could and should have considered the relevant moral norms but failed to do so through poor character or inadequate regard for others. Aquinas is also often associated with the thought underlying

the familiar maxim that ignorance of the law is no excuse.

generated alerts of future risk. In that respect, they may reduce some of the epistemic uncertainty that underlies skepticism about tracing responsibility to earlier omissions. Where a predictive system issues a salient warning – for example through visual, auditory, or haptic cues – at a time when the user still retains cognitive capacity and practical options for response, it becomes more plausible to say that the agent had a reasonable opportunity to recognize the risk and take precautions. This also narrows the force of Levy's concern that agents may act rationally on the basis of their actual mental states: where a clear alert indicates an elevated and imminent risk of harm, the claim that the agent reasonably believed no such risk existed is harder to sustain. None of this shows that every failure to respond to a CogBT alert is blameworthy. Much still depends on the reliability, intelligibility and timing of the alert, and on whether proportionate preventive action remained possible. But where those conditions are met, CogBTs may make it easier both to justify expectations of response and to identify the earlier point at which ignorance or inaction became culpable.

Intervening Factors

A common objection to expanding moral responsibility in response to biopredictive technologies is that such systems introduce intervening factors that may weaken the connection between an individual's conduct and subsequent harm. CogBTs often rely on algorithmic processes, ML models, or automated responses that operate with limited transparency and partial independence from their users. Where harm occurs following reliance on such systems, responsibility may appear diffused or obscured by technological mediation, rather than straightforwardly attributable to the user alone (De Villiers 2002; Martani et al. 2019; Meier et al. 2025).

However, technological mediation does not by itself displace individual responsibility. What matters instead is whether the intervening factors undermine the epistemic or control conditions that ground responsibility. Where a system failure is genuinely unforeseeable, undiscoverable, or occurs in a way that deprives the user of any meaningful opportunity to respond, responsibility may be attenuated or displaced. By contrast, when individuals silence alerts, ignore warnings, fail to maintain devices, neglect calibration or software updates, or continue to rely on systems despite known limitations, those earlier omissions may still ground responsibility. In both moral and legal terms, such conduct may amount to prior fault. These omissions or choices occur at a time when the individual retains awareness and control and may therefore ground responsibility for later harm, even if an automated system also plays a contributory role. In this respect, CogBTs do not displace responsibility; they reshape the temporal locus at which responsibility is assessed.

Mixed cases, where technological shortcomings and user omissions jointly contribute to harm, do not undermine this analysis. Moral responsibility is not all-or-nothing. It may be shared, distributed or mitigated depending on the relative contribution of different actors and systems. It is outside the scope of this paper to identify the multitude, and for legal purposes – mathematical – number of ways in which responsibility is divided in all contexts. Our argument is that where an individual unreasonably relies on a technology despite awareness of its limits, responsibility for harms may remain traceable to that earlier choice. Conversely, where a system fails in ways that exceed reasonable expectations of user oversight (e.g., faulty device), responsibility may appropriately shift away from the individual and toward manufacturers, clinicians or regulatory bodies.

Reinforcing stigma and stereotypes

A further objection is that predictive information may create the impression that individuals exercise meaningful agency over future impairment when, in reality, control remains partial, fragile or unevenly distributed. Assigning responsibility on this basis risks overstating individual agency and unfairly attributing blame for outcomes that remain substantially shaped by biological processes, chance, or technological limitations. There is also a broader concern that predictive responsibility frameworks may reinforce familiar narratives of personal failure or deficient self-control, particularly where the underlying condition or biomarker is already socially stigmatized (Broom and Whittaker 2004).

Several considerations limit the force of this objection. First, the framework developed here does not treat possession of a biomarker, or membership in a diagnostic category, as a basis for responsibility. Responsibility comes from failing to respond reasonably to a *specific and temporally situated risk* while retaining cognitive capacity. In that respect, the framework targets voluntary conduct rather than diagnostic status and avoids treating biological vulnerability as a moral failing (Broom and Whittaker 2004).

Second, CogBTs may in some contexts reduce rather than exacerbate stigma by enabling more precise and individualized assessments of risk. Where risk can be identified temporally and personally, regulatory and moral expectations need not be imposed broadly on all individuals associated with a particular biomarker. This stands in contrast to blunt group-based rules that implicitly treat all members of a category as equally dangerous. More granular prediction can therefore narrow the scope of intervention, limiting burdens to periods of demonstrably elevated risk and allowing individuals to participate more fully in social life when their risk is acceptably low.

Third, concerns about illusory control underscore the importance of institutional context. Where individuals lack access to predictive technologies, lack training to interpret them, or face structural barriers that undermine their ability to respond, expectations of control may be unjustified. As argued in the preceding section, corresponding obligations on states to ensure fair access and appropriate regulatory design are therefore essential. Without such support, appeals to individual responsibility risk becoming empty or punitive rather than expressions of fair moral expectation.

In sum, biopredictive technologies do not confer complete control over future impairment. They may, however, provide morally significant opportunities for anticipation and precaution. Where such opportunities exist, assigning responsibility rests on a defensible account of agency under conditions of risk, uncertainty, and partial governance. Where they do not, responsibility remains limited for the same reason it has traditionally remained limited: the relevant epistemic or control conditions are not satisfied.

IMPLICATIONS FOR FUTURE LAW AND POLICY

We have argued that individuals may bear heightened moral responsibilities in relation to the use, and in more limited circumstances the acquisition, of CogBTs where these technologies enable proportionate anticipation and mitigation of foreseeable harm arising from cognitive impairment. These claims have important implications for future law and policy. While moral responsibility and legal liability are not coextensive, the framework developed here helps illuminate how existing legal doctrines may evolve as CogBTs become more accurate, more widely used, and more deeply embedded in regulatory decision-making.

Judicial Impact

Many areas of criminal and civil law already rely on concepts closely aligned with the temporal structure of responsibility advanced in this paper. Doctrines of prior fault, negligence, recklessness, and foreseeability routinely ask whether an agent failed to take reasonable precautions at an earlier point in time, thereby contributing to later harm. Courts have long distinguished between conduct occurring during incapacity and antecedent failures to manage known risks. In this respect, CogBTs do not require a fundamental reorientation of legal responsibility. Rather, they supply new forms of evidence relevant to familiar inquiries about what an individual knew or should have known, and whether reasonable steps were available to prevent harm.

This continuity is evident in judicial approaches to medical incapacity. In the United Kingdom, courts have distinguished between loss of control arising without prior fault and impairment resulting from failures of risk management. In *R v Quick*[15], hypoglycemia induced by insulin was treated as an external factor capable of grounding automatism, whereas in *R v Hennessy*[16] hyperglycemia resulting from failure to take insulin was treated as an internal condition, leaving responsibility potentially traceable to the defendant's earlier conduct. Although the doctrinal classifications have been contested, the underlying structure reflects a concern with whether the defendant could reasonably have anticipated and managed the relevant risk.

[15] [1973] QB 910.
[16] [1989] 1 WLR 287.

Canadian courts have applied similar reasoning in driving cases involving medical impairment. In *R v Maki*[17] a driver with diabetes was convicted of dangerous driving where he knew his BGL was low upon testing but nevertheless chose to drive, resulting in harm. Responsibility attached not to his conduct during hypoglycemia, but to the earlier decision to drive in the face of a known and foreseeable risk. Similarly, in *R v Noureddin*,[18] a defendant with diagnosed dyscognitive epileptic seizures was convicted of dangerous driving causing death on the basis that he was aware of his condition, had received medical advice not to drive, and nevertheless assumed the risk.

By contrast, Australian courts have declined to impose criminal responsibility where loss of capacity was sudden and unforeseeable. In one widely reported case, a driver who fatally struck an eight-year-old child was acquitted on the basis that his hypoglycemia arose without warning, leaving no opportunity for anticipation or preventive action (McPhee 2022). Taken together, these cases illustrate that responsibility in this domain already turns on foresight, retained capacity, and the availability of proportionate precaution.

Biopredictive technologies may sharpen these assessments by providing more granular evidence about risk awareness and opportunity for prevention. Data generated by CogBTs can render aspects of risk anticipation and response more readily reconstructible, shifting what were previously highly subjective questions – what a person noticed, appreciated, or anticipated – into evidentiary records capable of more structured and potentially more objective assessment. CGM data, for example, may reveal not only that an alert occurred, but when it was issued, whether it was acknowledged, and whether any remedial action was taken. This does not eliminate interpretive judgment, nor does it guarantee reliable inference in every case. It may, however, reduce reliance on retrospective self-report and enable more precise evaluation of whether warnings were issued and ignored at a time when meaningful preventive action remained possible.

What remains less settled is how courts should assess responsibility in cases where harm could have been avoided only if an individual had *acquired* predictive technology in advance. Existing legal requirements in contexts such as diabetes already mandate contemporaneous monitoring – for example, checking BGL immediately prior to driving – but do not require individuals to anticipate future glucose levels through predictive tools. Moving from failures to respond to available information to failures to acquire predictive technology would therefore mark a more substantial legal step. It would require courts to rely less on established fitness-to-drive rules and more on broader doctrines such as negligence, proportionality, and reasonable precaution. This distinction underscores the importance of separating obligations to respond to predictive information from the more

[17] Superior Court of Justice, unreported reasons of judgment, 8 Dec 2011.
[18] 2015 ONCA 770, 128 O.R. (3d) 23, 332 C.C.C. (3d) 114.

demanding question of whether individuals can reasonably be required to acquire predictive tools in the first place.

Regulation and Policy

Beyond adjudication, CogBTs raise important questions for regulatory design. Regulatory regimes often rely on categorical exclusions where health-related risks are difficult to individualize. CogBTs may enable more differentiated approaches by supporting finer-grained risk banding within populations traditionally treated as homogeneous, allowing conditional participation where risk is acceptably low while maintaining restrictions where it is not. In driving contexts, for example, accurate seizure prediction could replace diagnosis-based exclusions with more individualized assessments of risk, thereby restoring agency and social participation to some individuals who would otherwise be broadly excluded.

At the same time, reliance on CogBTs in legal and regulatory contexts raises concerns about fairness, access, and proportionality. When participation in socially important activities becomes contingent on access to predictive technologies, unequal access risks translating directly into unequal responsibility and opportunity. As argued earlier, this places pressure on states to ensure that responsibility-enhancing technologies are not available only to those with financial or institutional advantage. Without such safeguards, CogBTs risk becoming instruments of exclusion rather than tools of fair risk governance.

Finally, while this paper has identified circumstances in which individuals may bear moral obligations to acquire CogBTs, whether such obligations should be legally enforced requires careful political and institutional judgment. Legal mandates should arise, if at all, only where acquisition represents a proportionate and least restrictive means of reducing serious harm, where alternative measures are inadequate, and where access is reasonably assured. In many cases, encouragement, subsidization or conditional incentives will be more appropriate than coercive requirements. The central contribution of CogBTs lies not in expanding the scope of legal blame, but in enabling more nuanced, fair, and temporally sensitive approaches to responsibility and risk governance.[19]

CONCLUSION

[19] Courts have long engaged with biological explanations of behavior without treating them as either fully exculpatory or determinative of responsibility. For example, evidence concerning the MAOA gene has been raised in mitigation at sentencing on the basis that genetic predisposition, interacting with adverse environmental factors such as childhood maltreatment, may be associated with diminished impulse control (Widom and Brzustowicz 2006; Baum 2013; Baum and Savulescu 2013). Such cases illustrate that biological contributions to behavior are routinely treated as mitigating rather than excusing, reinforcing the relevance of temporally structured responsibility under conditions of partial

control.

CogBTs are beginning to reshape how moral responsibility is understood in cases of foreseeable cognitive impairment. Their significance lies not simply in increasing general foreseeability, but in changing when responsibility becomes assessable, what may reasonably be demanded before impairment occurs, and how responsibility is distributed between individuals, institutions, and technological systems. Where predictive signals provide a sufficiently reliable and actionable basis for precaution, responsibility may depend on whether proportionate steps were taken while meaningful agency was still preserved.

At the same time, responsibilities of this kind are fair only if individuals have a realistic opportunity to meet them. We have therefore argued that emerging individual responsibilities may generate corresponding obligations for states and regulatory institutions to ensure that morally relevant biopredictive technologies are affordable and accessible. Without such support, expectations of individual responsibility risk becoming unfair, exclusionary, or practically unattainable.

Looking ahead, further work is needed to examine how emerging biopredictive capacities interact with doctrines of foreseeability, negligence, prior fault, and criminal culpability across different legal domains, including in relation to technologies aimed at predicting more permanent cognitive impairments such as dementia, where the moral structure of responsibility may differ in important respects. As predictive technologies become more accurate, more pervasive, and more deeply embedded in legal and regulatory decision-making, the central challenge will be to articulate principles of responsibility that recognize the moral significance of foresight while preserving fairness in the conditions under which responsibility is assigned.

Reference List

Aziz, Sarah, Amal A. M. Ali, Hania Aslam, et al. 2025. “Wearable Artificial Intelligence for Epilepsy: Scoping Review.” *Journal of Medical Internet Research* 27 (1): e73593. https://doi.org/10.2196/73593.

Bagley, Derek. 2020. “New AI-Powered Technology Increases Accuracy of Hypoglycemia and Hyperglycemia Predictions and Introduces Alert Capability for People Using Continuous Glucose Monitors.” *Endocrine News*, February 25. https://endocrinenews.endocrine.org/new-ai-powered-technology-increases-accuracy-of-hypoglycemia-and-hyperglycemia-predictions-and-introduces-alert-capability-for-people-using-continuous-glucose-monitors/.

Baum, Matthew L. 2016. *The Neuroethics of Biomarkers: What the Development of Bioprediction Means for Moral Responsibility, Justice, and the Nature of Mental Disorder*. Oxford University Press.

Baum, Matthew L., and Julian Savulescu. 2013. “Behavioral Biomarkers: What Are They Good For? Toward the Ethical Use of Biomarkers.” In *Bioprediction, Biomarkers, and Bad Behavior: Scientific, Legal, and Ethical Challenges*, edited by Ilina Singh, Walter P. Sinnott-Armstrong, and Julian Savulescu. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199844180.003.0002.

Bell, D., A. Huddart, and J. Krebs. 2010. “Driving and Insulin-Treated Diabetes: Comparing Practices in Scotland and New Zealand.” *Diabetic Medicine* 27 (9): 1093–95. https://doi.org/10.1111/j.1464-5491.2010.03084.x.

Bognetti, F., A. Brunelli, F. Meschi, M. Viscardi, R. Bonfanti, and G. Chiumello. 1997. “Frequency and Correlates of Severe Hypoglycaemia in Children and Adolescents with Diabetes Mellitus.” *European Journal of Pediatrics* 156 (8): 589–91. https://doi.org/10.1007/s004310050670.

Boughton, Charlotte K., Sara Hartnell, Hood Thabit, et al. 2026. “Hybrid Closed-Loop Glucose Control with Faster Insulin Aspart Compared with Standard Insulin Aspart in Adults with Type 1 Diabetes: A Double-Blind, Multicentre, Multinational, Randomized, Crossover Study.” *Diabetes, Obesity and Metabolism* 23 (6): 1389–96. https://doi.org/10.1111/dom.14355.

Breton, Marc, Anne Farret, Daniela Bruttomesso, et al. 2012. “Fully Integrated Artificial Pancreas in Type 1 Diabetes: Modular Closed-Loop Glucose Control Maintains near Normoglycemia.” *Diabetes* 61 (9): 2230–37. https://doi.org/10.2337/db11-1445.

Brinkmann, Benjamin H., Philippa J. Karoly, Ewan S. Nurse, et al. 2021. “Seizure Diaries and Forecasting With Wearables: Epilepsy Monitoring Outside the Clinic.” *Frontiers in Neurology* 12 (July). https://doi.org/10.3389/fneur.2021.690404.

Campbell, Laura K., Linda A. Gonder-Frederick, Donna K. Broshek, et al. 2010. “Neurocognitive Differences Between Drivers with Type 1 Diabetes with and without a Recent History of Recurrent Driving Mishaps.” *International Journal of Diabetes Mellitus* 2 (2): 73–77. https://doi.org/10.1016/j.ijdm.2010.05.014.

Child, John. 2014. “Prior Fault: Blocking Defences or Constructing Crimes.” SSRN Scholarly Paper No. 2401029. Social Science Research Network. https://papers.ssrn.com/abstract=2401029.

Cichosz, Simon Lebech, Søren Schou Olesen, and Morten Hasselstrøm Jensen. 2024. “Explainable Machine-Learning Models to Predict Weekly Risk of Hyperglycemia, Hypoglycemia, and Glycemic Variability in Patients With Type 1 Diabetes Based on Continuous Glucose Monitoring.” *Journal of Diabetes Science and Technology*, October 8, 19322968241286907. https://doi.org/10.1177/19322968241286907.

Cockerill, Richard G. 2020. “Ethics Implications of the Use of Artificial Intelligence in Violence Risk Assessment.” *The Journal of the American Academy of Psychiatry and the Law* 48 (3): 345–49. https://doi.org/10.29158/JAAPL.003940-20.

Corbett, John P., Liana Hsu, Sue A. Brown, et al. 2022. “Smartwatch Gesture-Based Meal Reminders Improve Glycaemic Control.” *Diabetes, Obesity and Metabolism* 24 (8): 1667–70. https://doi.org/10.1111/dom.14737.

Costa, Gonçalo, César Teixeira, and Mauro F. Pinto. 2024. “Comparison between Epileptic Seizure Prediction and Forecasting Based on Machine Learning.” *Scientific Reports* 14 (1): 5653. https://doi.org/10.1038/s41598-024-56019-z.

Cox, Daniel J., Linda A. Gonder-Frederick, Harsimran Singh, et al. 2017. “Predicting and Reducing Driving Mishaps Among Drivers With Type 1 Diabetes.” *Diabetes Care* 40 (6): 742–50. https://doi.org/10.2337/dc16-0995.

Cox, Daniel J., Jennifer Kim Penberthy, John Zrebiec, et al. 2003. “Diabetes and Driving Mishaps: Frequency and Correlations from a Multinational Survey.” *Diabetes Care* 26 (8): 2329–34. https://doi.org/10.2337/diacare.26.8.2329.

Cvach, Maria. 2012. “Monitor Alarm Fatigue: An Integrative Review.” *Biomedical Instrumentation & Technology* 46 (4): 268–77. https://doi.org/10.2345/0899-8205-46.4.268.

Dalton, Tyrone. 2024. “Court Hears Diabetic Driver Appeared ‘vague’ before Crash That Killed Five.” *ABC News* (Central Victoria), September 16. https://www.abc.net.au/news/2024-09-16/daylesford-driver-william-swale-ballarat-court-fatal-diabetes/104355420.

Daniels, Norman. 1985. *Just Health Care*. Cambridge University Press.

De Villiers, Etienne. 2002. “Who Will Bear Moral Responsibility?” *Communicatio* 28 (1): 16–21. https://doi.org/10.1080/02500160208537953.

Ekhlaspour, Laya, Melissa J. Schoelwer, Gregory P. Forlenza, et al. 2021. “Safety and Performance of the Tandem t:Slim X2 with Control-IQ Automated Insulin Delivery System in Toddlers and Preschoolers.” *Diabetes Technology & Therapeutics* 23 (5): 384–91. https://doi.org/10.1089/dia.2020.0507.

Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge Studies in Philosophy and Law. Cambridge University Press. https://doi.org/10.1017/CBO9780511814594.

FitzPatrick, William J. 2008. “Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge.” *Ethics* 118 (4): 589–613. https://doi.org/10.1086/589532.

Georga, Eleni I., Vasilios C. Protopappas, Diego Ardigò, Demosthenes Polyzos, and Dimitrios I. Fotiadis. 2013. “A Glucose Model Based on Support Vector Regression for the Prediction of Hypoglycemic Events Under Free-Living Conditions.” *Diabetes Technology & Therapeutics* 15 (8): 634–43. https://doi.org/10.1089/dia.2012.0285.

Giammarino, Flavia, Ransalu Senanayake, Priya Prahalad, David M. Maahs, and David Scheinker. 2025. "A Machine Learning Model for Week-Ahead Hypoglycemia Prediction From Continuous Glucose Monitoring Data." *Journal of Diabetes Science and Technology* 19 (5): 1353–61. https://doi.org/10.1177/19322968241236208.

Glannon, Walter. 1998. "Moral Responsibility and Personal Identity." *American Philosophical Quarterly* 35 (3): 231–49.

Grosman, Benyamin, Di Wu, Neha Parikh, et al. 2021. "Fast-Acting Insulin Aspart (Fiasp®) Improves Glycemic Outcomes When Used with MiniMedTM 670G Hybrid Closed-Loop System in Simulated Trials Compared to NovoLog®." *Computer Methods and Programs in Biomedicine* 205 (June): 106087. https://doi.org/10.1016/j.cmpb.2021.106087.

He, Xuecheng, Wei Huang, Wensheng Lin, et al. 2026. "A Wearable, Dual Closed-Loop Insulin Delivery System for Precision Diabetes Management." *Advanced Materials* 38 (8): e14945. https://doi.org/10.1002/adma.202514945.

Huang, Wei, Ivo Pang, Jing Bai, Binbin Cui, Xiaojuan Qi, and Shiming Zhang. 2025. "Artificial Intelligence-Enhanced, Closed-Loop Wearable Systems Toward Next-Generation Diabetes Management." *Advanced Intelligent Systems* Early View. https://doi.org/10.1002/aisy.202400822.

Imrisek, Steven D., Matthew Lee, Dan Goldner, et al. 2022. "Effects of a Novel Blood Glucose Forecasting Feature on Glycemic Management and Logging in Adults With Type 2 Diabetes Using One Drop: Retrospective Cohort Study." *JMIR Diabetes* 7 (2): e34624. https://doi.org/10.2196/34624.

Karter, Andrew J., Melissa M. Parker, Howard H. Moffet, et al. 2023. "Validation of a Hypoglycemia Risk Stratification Tool Using Data From Continuous Glucose Monitors." *JAMA Network Open* 6 (3): e236315. https://doi.org/10.1001/jamanetworkopen.2023.6315.

Kovatchev, Boris P., Daniel J. Cox, Linda Gonder-Frederick, and William L. Clarke. 2002. "Methods for Quantifying Self-Monitoring Blood Glucose Profiles Exemplified by an Examination of Blood Glucose Patterns in Patients with Type 1 and Type 2 Diabetes." *Diabetes Technology & Therapeutics* 4 (3): 295–303. https://doi.org/10.1089/152091502760098438.

Lee, Andrew, Daniel Buchman, Katrina Hui, et al. 2025. "Ethical and Legal Considerations of Artificial Intelligence Applications in Psychiatric Violence Risk Assessment: A Scoping Review Protocol." *PLOS ONE* 20 (10): e0334649. https://doi.org/10.1371/journal.pone.0334649.

Levy, Neil. 2009. "Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick." *Ethics* 119 (4): 729–41. https://doi.org/10.1086/605018.

Mader, Julia K., Guido Freckman, Thomas Forst, et al. 2024. “1924-LB: Clinical Performance of a Novel CGM System.” *Diabetes* 73 (Supplement_1): 1924-LB. https://doi.org/10.2337/db24-1924-LB.

Martani, Andrea, David Shaw, and Bernice Simone Elger. 2019. “Stay Fit or Get Bit – Ethical Issues in Sharing Health Data with Insurers’ Apps.” *Swiss Medical Weekly* 149 (2526): 2526. https://doi.org/10.4414/smw.2019.20089.

Matus, Austin, Anneliese J. Flatt, Amy J. Peleckis, Cornelia Dalton-Bakes, Barbara Riegel, and Michael R. Rickels. 2023. “Validating and Establishing a Diagnostic Threshold for the Hypoglycemia Awareness Questionnaire Impaired Awareness Subscale.” *Endocrine Practice* 29 (10): 762–69. https://doi.org/10.1016/j.eprac.2023.08.004.

Meier, E., T. Rigter, M. P. Schijven, M. van den Hoven, and M. A. R. Bak. 2025. “The Impact of Digital Health Technologies on Moral Responsibility: A Scoping Review.” *Medicine, Health Care and Philosophy* 28 (1): 17–31. https://doi.org/10.1007/s11019-024-10238-3.

Moghim, Negin, and David W. Corne. 2014. “Predicting Epileptic Seizures in Advance.” *PloS One* 9 (6): e99334. https://doi.org/10.1371/journal.pone.0099334.

Ore, Adeshola. 2023. “Driver Charged over Fatal Daylesford Pub Crash Ignored Blood Glucose Level Warnings, Court Hears.” Australia News. *The Guardian*, December 11. https://www.theguardian.com/australia-news/2023/dec/11/man-charged-over-daylesford-hotel-beer-garden-crash-that-left-five-people-dead.

Parfit, Derek. 2011. *On What Matters*. Vol. 1. Oxford University Press.

Parkes, J. L., S. L. Slatin, S. Pardo, and B. H. Ginsberg. 2000. “A New Consensus Error Grid to Evaluate the Clinical Significance of Inaccuracies in the Measurement of Blood Glucose.” *Diabetes Care* 23 (8): 1143–48. https://doi.org/10.2337/diacare.23.8.1143.

Pedersen-Bjergaard, Ulrik, Stig Pramming, Simon R. Heller, et al. 2004. “Severe Hypoglycaemia in 1076 Adult Patients with Type 1 Diabetes: Influence of Risk Markers and Selection.” *Diabetes/Metabolism Research and Reviews* 20 (6): 479–86. https://doi.org/10.1002/dmrr.482.

Petrovski, Goran, Judith Campbell, Maheen Pasha, et al. 2023. “Simplified Meal Announcement Versus Precise Carbohydrate Counting in Adolescents With Type 1 Diabetes Using the MiniMed 780G Advanced Hybrid Closed Loop System: A Randomized Controlled Trial Comparing Glucose Control.” *Diabetes Care* 46 (3): 544–50. https://doi.org/10.2337/dc22-1692.

Ra, Jee Sook, Tianning Li, and YanLi. 2023. “A Novel Epileptic Seizure Prediction Method Based on Synchroextracting Transform and 1-Dimensional Convolutional Neural Network.” *Computer Methods and Programs in Biomedicine* 240 (October): 107678. https://doi.org/10.1016/j.cmpb.2023.107678.

Rashid, Mudassir M., Laurie Quinn, and Ali Cinar. 2026. "Fully-Automated Insulin Delivery System." *Journal of Diabetes Science and Technology*, January 31, 19322968251409204. https://doi.org/10.1177/19322968251409204.

Rewers, Arleta, H. Peter Chase, Todd Mackenzie, et al. 2002. "Predictors of Acute Complications in Children with Type 1 Diabetes." *JAMA* 287 (19): 2511–18. https://doi.org/10.1001/jama.287.19.2511.

Sher, George. 2009. *Who Knew?: Responsibility Without Awareness*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195389197.001.0001.

Smith, Holly. 1983. "Culpable Ignorance." *The Philosophical Review* 92 (4): 543–71.

Stirling, Rachel E., Mark J. Cook, David B. Grayden, and Philippa J. Karoly. 2021. "Seizure Forecasting and Cyclic Control of Seizures." *Epilepsia* 62 (S1): S2–14. https://doi.org/10.1111/epi.16541.

Stork, Alexander D. M., Timon W. Van Haeften, and Thiemo F. Veneman. 2006. "Diabetes and Driving: Desired Data, Research Methods and Their Pitfalls, Current Knowledge, and Future Research." *Diabetes Care* 29 (8): 1942–49. https://doi.org/10.2337/dc05-2232.

Strawson, P. F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* (London), May 9, 187–211.

Talbert, Matthew. 2024. "Moral Responsibility." In *The Stanford Encyclopedia of Philosophy*, Fall 2024, edited by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2024/entries/moral-responsibility/.

Templer, Sophie. 2022. "Closed-Loop Insulin Delivery Systems: Past, Present, and Future Directions." *Frontiers in Endocrinology* 13 (June): 919942. https://doi.org/10.3389/fendo.2022.919942.

Tsichlaki, Stella, Lefteris Koumakis, and Manolis Tsiknakis. 2022. "Type 1 Diabetes Hypoglycemia Prediction Algorithms: Systematic Review." *JMIR Diabetes* 7 (3): e34699. https://doi.org/10.2196/34699.

Turksoy, Kamuran, Elif Seyma Bayrak, Lauretta Quinn, Elizabeth Littlejohn, and Ali Cinar. 2013. "Multivariable Adaptive Closed-Loop Control of an Artificial Pancreas without Meal and Activity Announcement." *Diabetes Technology & Therapeutics* 15 (5): 386–400. https://doi.org/10.1089/dia.2012.0283.

Vargas, Manuel. 2005. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29 (1): 269–91. https://doi.org/10.1111/j.1475-4975.2005.00117.x.

Warren, R. E., N. N. Zammitt, I. J. Deary, and B. M. Frier. 2007. "The Effects of Acute Hypoglycaemia on Memory Acquisition and Recall and Prospective Memory in

Type 1 Diabetes." *Diabetologia* 50 (1): 178–85. https://doi.org/10.1007/s00125-006-0535-6.

Yale, Jean-François, Breay Paty, and Peter A. Senior. 2018. "Hypoglycemia." *Canadian Journal of Diabetes*, Diabetes Canada 2018 Clinical Practice Guidelines for the Prevention and Management of Diabetes in Canada, vol. 42 (April): S104–8. https://doi.org/10.1016/j.jcjd.2017.10.010.

Zimmerman, Michael J. 1986. "Negligence and Moral Responsibility." *Noûs* 20 (2): 199–218. https://doi.org/10.2307/2215391.