Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Patient responsibility and subsequent care in predictive medicine: a normative analysis

Xiongwen Yang

yangxiongwen@gz5055.com

Guizhou Provincial People's Hospital

Fang Wu

Guizhou Provincial People's Hospital

Lixia Liu

Guizhou Provincial People's Hospital

Chuan Xu

Guizhou Provincial People's Hospital

Research Article

Keywords: Predictive medicine, patient responsibility, medical ethics, causal inference, healthcare allocation, clinical decision-making, health data governance

Posted Date: October 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11163906/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Background

Patient-specific predictions increasingly inform prevention and treatment decisions. When retained, the same information may later be used to judge whether a patient should have acted differently and whether that judgment should affect subsequent care. We examine the additional justification required when predictive information moves from supporting a choice to attributing responsibility and imposing consequences.

Methods

We conducted a normative analysis informed by 108 purposively selected methodological, empirical and bioethical publications. We distinguished inferential targets and reconstructed the factual and normative premises required for responsibility attribution and institutional action. Two hypothetical clinical cases and a contrasting variant examined these distinctions. Published research informed the cases' access constraints and treatment burdens without validating the framework or establishing that predictive systems cause responsibility-based restrictions on care.

Results

The Gated Relevance Framework distinguishes clinical and deliberative uses from three responsibility-sensitive transitions: establishing a reasonably available opportunity to reduce risk, justifying a specified responsibility judgment, and independently justifying an institutional consequence. Estimated risk reduction does not establish that a particular adverse outcome would have been prevented. The transplantation case illustrates the absence of an available opportunity; the endocrine-therapy case shows why feasible treatment and informed refusal do not alone establish culpability. The contrasting variant demonstrates that even a defensible responsibility judgment does not automatically justify reduced priority for care. We define normative relevance migration as an unsupported change in the justificatory role assigned to information. Proposed safeguards preserve the purpose, uncertainty and context of predictive outputs and provide opportunities for patient challenge and accountable review.

Conclusions

Predictive information cannot by itself establish patient responsibility or justify adverse consequences. The framework identifies additional grounds required for these judgments while keeping clinical and deliberative uses distinct. It does not provide a sufficient justification for restricting care. Its feasibility, practical effects and potential unintended harms require empirical evaluation.

Clinical trial number: not applicable.

Background

Estimates of future health outcomes can help patients and clinicians assess disease risks and weigh preventive and therapeutic options. Predictive medicine brings such estimates into clinical decision-making; here, we focus on patient-specific estimates of disease risk, prognosis and outcomes under alternative treatment options. The tools producing these estimates have different inferential targets. Risk models estimate outcome probabilities, whereas treatment-effect methods seek to compare outcomes under alternative interventions. Patient-specific simulations, including emerging digital twins, may represent alternative health trajectories, although individualization alone does not establish causal validity[1–3]. Once a treatment decision has been made, however, predictive information may acquire a different role: an estimate originally used to support a choice may later be invoked to judge that choice. This raises the question of what such information can legitimately establish about patient responsibility and subsequent care.

Consider a hypothetical patient who declines a preventive treatment after discussing an estimate of its expected benefit. Years later, the patient develops the complication that the treatment was intended to prevent. The earlier estimate may help a clinician understand the treatment history and the patient's previous concerns. It might also be invoked to argue that the patient could have avoided the complication, bears responsibility for its occurrence, or should receive lower priority for subsequent treatment. Yet evidence that an intervention could have reduced risk does not by itself establish any of these conclusions. A treatment may have been inaccessible or excessively burdensome; a decision to decline it may have been reasonable; and even a justified attribution of responsibility may be insufficient to warrant restricting care.

Personal responsibility for health is an established concern in medical ethics. Existing accounts examine agency, opportunity and social circumstances[4, 5], and distinguish responsibility attribution from the justification for holding someone accountable[6]. Earlier work on predictive and personalised medicine has also considered whether greater knowledge of health risks may strengthen expectations of preventive action[7]. The present analysis focuses on a specific feature of this broader debate: the same patient-specific estimate can be used prospectively to support a choice and retrospectively to judge it. Its quantified, individualized form may appear to provide a concrete basis for claims about what this patient could or should have done. However, neither methodological validity nor specificity to the patient supplies the moral premises needed for those claims. The issue is therefore how the justificatory role of predictive information changes across decisions, including when the original estimate is sound.

Empirical research provides indirect grounds for examining this issue. Everett and colleagues found that responsibility considerations entered doctors' stated prioritization judgments in surveys and allocation vignettes, with variation across behaviors and proposed consequences[8]. Research on clinical documentation also shows why the interpretation of recorded information warrants attention. In a randomized vignette study, Goddu and colleagues found that stigmatizing descriptions of the same hypothetical patient elicited less favorable attitudes; resident physicians also selected less intensive pain treatment[9]. Sun and colleagues identified racial differences in the use of negative patient descriptors in clinical notes[10]. These studies concern stated allocation judgments, simulated

treatment decisions and observational documentation patterns, respectively. They do not establish that predictive systems increase blame or cause restrictions on care. They nevertheless motivate scrutiny of how recorded predictive information might interact with responsibility judgments when reused in later decisions.

Related evidence from AI-mediated communication highlights questions about how clinical information acquires authority. In an exploratory mixed-methods study involving 32 undergraduate nursing students, participants' responses to clinical scenarios identified tensions concerning interpretive authority, professional responsibility and the communication of uncertainty when patients used AI-generated health information[11]. These findings concern perceptions in an educational setting, rather than observed clinical allocation. They nevertheless motivate attention to how AI-mediated information is interpreted, while leaving open whether predictive estimates are used to attribute responsibility or restrict care.

The normative task is to distinguish the reasons that justify these different uses of information. A prediction may inform treatment or help explain a patient's values, concerns and support needs without supporting blame. Using it to attribute responsibility requires further consideration of the patient's opportunities, understanding, control and reasons for acting. Using that attribution to influence subsequent care requires an additional justification concerning the purpose and fairness of the institutional response. These questions matter because information can retain clinical or deliberative value while being insufficient to justify adverse judgments about a patient or restrictions on treatment.

This article examines two questions: under what conditions can patient-specific predictive information support an attribution of responsibility, and when, if ever, may that attribution legitimately affect subsequent care? We develop a three-gate normative analysis addressing evidential adequacy and reasonable opportunity, the conditions of responsibility attribution, and the justification of responsibility-sensitive institutional action. Published empirical literature and hypothetical clinical cases inform and illustrate the analysis. The aim is to make the additional premises required at each transition explicit and open to challenge, helping clinicians, clinical ethics services and institutional decision-makers assess the legitimate role of predictive information in subsequent care.

Methods

Study design and analytical scope

This study develops a normative analysis of the conditions under which patient-specific predictive information may justify responsibility attribution and influence subsequent care. The unit of analysis is a proposed inference or use of information: what an estimate supports, which additional factual and normative premises a further conclusion requires, and what remains justified when a necessary premise is absent or uncertain. We distinguish the clinical use of information to assess the benefits, burdens or

feasibility of care; its deliberative use to understand patients' values, reasons and support needs; its use in attributing responsibility; and its use in justifying responsibility-sensitive institutional action.

The analysis draws on methodological research, bioethical arguments and selected empirical studies. Its purpose is to examine justificatory relationships rather than estimate the frequency of responsibility-sensitive decisions or validate a clinical decision instrument. Technical validity and ethical justification are treated as related but distinct. A methodologically sound estimate may support one use without supporting another, while methodological uncertainty may itself constrain the conclusions that can reasonably be drawn.

Literature identification, selection and use

Literature searches were conducted on 1 August 2026 using PubMed, Web of Science and Google Scholar. These sources were used together to identify biomedical research and relevant scholarship in philosophy and bioethics, including work not indexed in PubMed. Searches addressed four questions: what patient-specific predictions and treatment-effect estimates can support; under what conditions health-related responsibility can reasonably be attributed; how responsibility judgments and clinical documentation may influence clinical reasoning; and how information reuse and responsibility-sensitive healthcare allocation can be justified. Searches also addressed treatment access and treatment burdens relevant to the clinical illustrations. The search strategies, selection approach and limitations of the available search documentation are described in **Additional file 1**.

Following full-text reading, 108 publications were selected purposively for their relevance to these questions and their contribution to the analysis. Selection was directed toward clarifying concepts, identifying necessary premises, examining competing normative positions, and providing empirical context for the mechanisms and clinical circumstances under discussion. Inclusion did not imply agreement with a source's conclusions. Sources defending responsibility-sensitive policies and sources challenging their justification were both relevant to the argument.

Methodological sources informed distinctions between predicted risk, estimated intervention benefit and retrospective attribution of an outcome. Normative sources informed the examination of agency, reasonable opportunity, accountability, privacy and legitimate institutional action. Empirical studies provided evidence about expressed responsibility judgments, documentation practices, treatment access, treatment burdens, medication-taking behavior and supportive interventions. These sources served different argumentative purposes. Empirical findings could inform the assessment of circumstances and consequences, but could not independently establish whether a responsibility judgment or institutional response was justified.

Additional file 2 maps the included publications across seven analytical themes: predictive information and responsibility; responsibility, agency and allocation; treatment effects and causal inference; digital twins and model validity; stigma, documentation and governance; transplant access and adherence; and

endocrine therapy burden and adherence. These themes organize the literature by its role in the analysis rather than by retrieval source. Each entry reports the study design or source type, relevant findings or arguments, analytical contribution and limits of inference. Entries represent publications rather than necessarily independent studies or datasets. Reviews and related publications were therefore not treated as independent replications of the evidence they described.

Interpretation of empirical findings was bounded by the reported population, design and outcome. Survey responses were treated as evidence of expressed judgments, vignette experiments as evidence of responses under simulated conditions, and observational findings as associations unless the design supported a stronger inference. Qualitative studies informed understanding of experiences, reasons and practices without establishing their population prevalence. Trials of supportive interventions informed assessment of the interventions and outcomes studied, without establishing that support would be effective or accessible for every patient. Measures of medication possession, monitoring events and self-reported adherence were distinguished from direct evidence of medication ingestion, intentional refusal or culpability.

Published findings were distinguished from assumptions introduced in the hypothetical cases. This purposive literature selection supported a focused normative argument rather than an exhaustive evidence synthesis. No pooled effect estimate, formal risk-of-bias assessment or claim of comprehensive coverage is made.

Normative analysis and framework development

The analysis reconstructs the premises required to move between three distinct conclusions. The first concerns whether an estimated benefit corresponds to a reasonably available opportunity for the patient. The second concerns whether the patient's response to that opportunity supports a specified attribution of responsibility. The third concerns whether that attribution provides a legitimate reason for an institutional action affecting subsequent care. These transitions organize the three gates of the Gated Relevance Framework.

For each transition, factual premises are distinguished from normative premises. An estimate that treatment would reduce risk, for example, does not establish that treatment was accessible or feasible for the patient. Evidence that a feasible option was declined does not establish that accepting it was reasonably required. Likewise, attributing responsibility does not establish that a particular restriction on care would be justified. The analysis identifies the additional reasons required at each transition and examines how missing information, uncertainty or disagreement affects the conclusion.

The argument is guided by respect for patient agency, attention to unequal opportunities, and a requirement that institutions justify adverse consequences imposed on patients. These commitments direct attention to understanding, control, treatment burdens, competing obligations and institutional purposes. They do not presuppose that patients can never bear responsibility or that responsibility can

never be relevant to allocation. We distinguish claims that depend on contested accounts of patient obligations from the narrower claim that predictive information alone cannot supply the necessary justification.

Competing positions are used to examine both the requirements and limits of the proposed gates. Opportunity-based accounts inform the conditions under which responsibility might be attributed, while objections concerning structural disadvantage, privacy and professional standing challenge whether such attribution can be established or legitimately acted upon. The aim is to make these dependencies explicit rather than resolve every disagreement about responsibility in healthcare.

The resulting distinctions are examined for consistency across cases and for their ability to identify unsupported transitions. Particular attention is given to the differences between reducing risk and preventing an individual outcome, between responsibility for a choice and responsibility for a subsequent health condition, and between attributing responsibility and justifiably acting on that attribution. Treatment burdens are examined both for their effects on practical feasibility and for their relevance to whether accepting an otherwise feasible treatment can reasonably be required. The framework is presented as a structure for reasoned assessment, not as a validated scoring system.

Case construction and examination

Two hypothetical clinical cases and a contrasting variant are used to examine the implications and limits of the analysis. The kidney transplantation case addresses whether estimated treatment benefit represents a reasonably available opportunity when medication access and continuity of coverage are constrained. The adjuvant endocrine therapy case addresses whether declining a potentially beneficial treatment supports an adverse responsibility judgment when treatment burdens conflict with the patient's circumstances and priorities. The contrasting variant examines whether a stronger basis for responsibility attribution would be sufficient to justify an adverse decision about subsequent care.

Published studies inform the circumstances represented in the cases. For the transplantation case, relevant evidence concerns medication access, insurance coverage, administrative barriers and practical support. For the endocrine therapy case, relevant evidence concerns symptom burden, patients' treatment experiences, reasons for treatment interruption or discontinuation, and the possibilities and limits of supportive interventions. These findings inform the plausibility and interpretation of the scenarios without determining the normative conclusions.

Constructed patient narratives, numerical estimates and decision scenarios are distinguished from those studies' findings. They do not describe actual patients or deployed predictive models, and the numerical estimates are illustrative rather than clinically validated. Published studies are used as sources of evidence; their underlying participant-level datasets are neither obtained nor reanalyzed.

Each case is examined by identifying the proposed use of predictive information, specifying the factual and normative premises required for that use, and assessing which premises are supported, absent or

unresolved within the scenario. Comparisons across the cases and the contrasting variant examine how differences in opportunity, burden, agency and institutional purpose affect the conclusions that can be justified. These comparisons provide conceptual examination rather than empirical validation or evidence of clinical effectiveness.

The analytical output is a reasoned assessment of a particular use of information, including a conclusion that its grounds are insufficient or unresolved. Within the proposed framework, an unsupported or unresolved necessary premise does not warrant progression toward an adverse responsibility-sensitive decision. This asymmetric burden of justification reflects the normative commitments stated above. It applies to the proposed adverse inference, rather than to the provision of care: insufficient grounds for responsibility attribution do not make the information clinically irrelevant or preclude appropriate treatment, further investigation or supportive discussion.

Results

Distinct uses of predictive information require distinct reasons

The analysis distinguishes four uses of patient-specific predictive information: assessing the benefits, burdens or feasibility of care; understanding a patient's values, reasons or support needs; attributing responsibility; and justifying an institutional response on the basis of that responsibility. The first two uses are clinical and deliberative. Neither requires a finding that the patient is blameworthy.

An earlier interruption of medication may warrant investigation of a continuing access barrier. A previous refusal may help a clinician understand which treatment burdens the patient considers unacceptable. In both situations, the information can remain relevant to present care without establishing fault or reduced entitlement. Its relevance nevertheless requires reassessment: past circumstances and preferences need not describe the patient's current situation.

Responsibility-sensitive uses require additional premises. An estimate of treatment benefit does not itself establish that the alternative was reasonably available, that the patient ought to have taken it, or that foregoing it justifies a later disadvantage. The Gated Relevance Framework organizes these requirements into three transitions:

1. From estimated benefit to a reasonably available opportunity.
2. From that opportunity to a warranted responsibility judgment.
3. From that judgment to a legitimate institutional consequence.

These gates concern the justification of a proposed use, rather than stages that every patient must undergo. A prospective assessment of treatment feasibility does not become a responsibility judgment merely because it considers previous behavior. The distinction depends on the reason offered: whether

the information supports an assessment of present benefit and feasibility, or a claim that a past choice makes the patient less entitled to care.

Table 1 summarizes the transitions and their stopping conditions. Support at one gate does not establish the requirements of the next or independently authorize further investigation of the patient.

Table 1
Inferential requirements of the Gated Relevance Framework.

Gate	Question	Additional grounds required	Conclusion potentially supported	Stopping condition
1. Opportunity	Did the estimate identify an opportunity the patient could reasonably have taken?	An estimate appropriate to the claim, relevant uncertainty, and evidence about decision-time access and feasibility	A specified opportunity to reduce risk	Inadequate inferential warrant, an unavailable alternative, or unresolved material barriers
2. Responsibility	Could the patient reasonably have been expected to take that opportunity?	Understanding, relevant control, a defensible normative expectation, and consideration of burdens and competing reasons	A specified responsibility judgment concerning a choice, risk or outcome	A necessary factual premise or the normative basis of the expectation remains unsupported
3. Institutional consequence	Why should this responsibility judgment justify this consequence?	A defensible institutional purpose, substantive relevance of responsibility, proportionality, equal treatment and accountable review	A qualified justification for the specified institutional consequence	The consequence lacks substantive justification or adequate procedural safeguards

Note. The gates identify necessary justificatory requirements; satisfying one does not establish the next, and satisfying all three does not by itself provide a sufficient justification for restricting care. Responsibility for a choice or accepted risk does not establish responsibility for a particular outcome, which requires additional causal warrant. Stopping conditions apply to the proposed responsibility-sensitive inference or action, not to appropriate treatment, investigation or supportive discussion. The gates represent logical dependencies rather than a mandatory sequence of patient assessments.

Table 2
Conclusions supported by the cases.

Scenario	Gate 1: opportunity	Gate 2: responsibility	Gate 3: consequence	Continuing relevance to care
Medication-access barriers	The supported regimen was not reasonably available	The missing opportunity premise prevents the proposed attribution	A responsibility-based reduction in priority is unsupported	Reassess present access and support; examine institutional failures
Burdensome endocrine therapy	A risk-reduction opportunity is supported, conditional on model assumptions and stipulated feasibility	Access, understanding and support discussions do not establish a duty to accept residual burdens	Restriction based on the unestablished blame judgment is unsupported	Assess current treatment benefit and revisit preferences
Low-burden contrasting variant	The opportunity is supported by stipulation	Limited responsibility is supportable under the stipulated duty	A separate justification remains necessary	Assess present need, benefit and ordinary eligibility criteria

Note. The cases are hypothetical and informed by published literature; they do not report observed patient outcomes or validate the framework. Conclusions depend on the factual and normative assumptions specified in each scenario. The contrasting variant stipulates a limited duty to accept a low-burden opportunity to avert serious harm; it neither establishes a general duty to accept treatment nor attributes responsibility for the particular adverse outcome. Unsupported responsibility-sensitive inferences do not preclude appropriate clinical assessment or care.

Gate 1: From estimated benefit to a reasonably available opportunity

Gate 1 requires both inferential warrant and contextual availability. The estimate must support the comparison being proposed, and the alternative must have been realistically available to the patient when the decision was made.

A prediction of poor outcome does not by itself establish the benefit of an intervention[12]. Treatment-effect estimates require causal assumptions appropriate to the study design and intended use, including consistency for sufficiently well-defined interventions, exchangeability through randomization or adequate control of confounding, and positivity or sufficient treatment overlap[13]. Applying estimates or treatment rules to a target population beyond the study sample requires additional consideration of population differences and the assumptions supporting their transfer[14]. Personalization or a digital-twin label does not itself establish these conditions.

Even a credible estimate of treatment benefit generally does not reveal what the same individual would have experienced under both alternatives. A conditional average treatment effect is an average contrast

among individuals with specified characteristics, not an observed treatment effect for a particular patient[13, 15]. The supported conclusion may therefore concern an opportunity to improve the odds rather than proof that a particular adverse event would have been prevented.

The object of the claim must remain explicit. An avoidable exposure, a forgone risk-reduction opportunity and an avoidable individual outcome are different claims. Evidence sufficient to support one does not automatically support the others.

Contextual availability requires attention to conditions at the time of the choice. Studies of medication-taking after kidney transplantation and access to endocrine therapy describe practical barriers involving insurance, medication supply, refill arrangements, costs and support[16, 17]. These findings inform the circumstances requiring assessment; they do not establish which barriers affected a particular patient. Relevant considerations may also include communication needs, practical assistance, treatment burdens and competing obligations. These circumstances require contextual inquiry rather than inference from demographic or socioeconomic proxies. The patient's account is central, and absent documentation of a barrier should not be treated as evidence that no barrier existed.

This assessment draws on opportunity-based and socially situated accounts of responsibility[18, 19]. It also requires attention to institutional contributions. An alternative dependent on promised but undelivered support cannot be treated as though that support had been available. The literature on endocrine-therapy adverse effects identifies interventions for managing selected symptoms[20]. For the present analysis, however, the existence of such interventions does not establish that they were accessible, acceptable or sufficiently effective for the patient, or that the remaining treatment burden was negligible.

Gate 2: From opportunity to a warranted responsibility judgment

Gate 2 examines whether the patient could reasonably have been expected to take the opportunity. Establishing that an option was feasible and medically preferable does not establish a moral obligation to accept it.

Drawing on these accounts, we organize the assessment around four questions:

1. Did the patient have the decision-making capacity and understanding relevant to the choice?
2. Did the patient have sufficient control over the choice and its implementation?
3. What normative reason supports an expectation that the patient should have acted differently?
4. Do treatment burdens, the patient's reasons or competing obligations defeat that expectation?

These questions are proposed here as a deliberative structure, informed by existing accounts of knowledge, control, reasonable opportunity and justified accountability[4–6, 18]. They are not a validated instrument or a responsibility score.

The third question cannot be answered by improving predictive accuracy. A proposed expectation must identify its normative basis, such as a limited duty of reasonable self-care or an obligation to protect others, and explain why it applies in the circumstances. The framework does not assume an unrestricted duty to maximize health or minimize healthcare expenditure.

Treatment burden can be relevant at both Gates 1 and 2. At Gate 1, it may make an alternative practically unavailable. At Gate 2, an option may remain feasible while its burdens make accepting it unreasonable to demand. The same distinction applies to caregiving and other obligations: they may prevent implementation or provide a defensible reason to decline an implementable option.

Disagreement with a clinical recommendation does not by itself establish inadequate understanding. Decision-making capacity requires assessment of the abilities relevant to the particular decision[21]. Addiction, distress or disability likewise requires assessment of its actual implications for understanding and control, rather than an automatic conclusion of either full responsibility or incapacity[4].

Recorded medication-taking behavior also requires interpretation. Self-reports, electronic monitoring and dispensing records provide different information about medication use and have different sources of measurement error[22]. A recorded cessation therefore requires clarification of who initiated it and why. For example, Kuba and colleagues classified patient-requested stopping as discontinuation, while retaining some physician-directed cessation within their study-specific persistence category[23]. Such classifications must be understood before medication records are treated as evidence of an intentional choice, let alone a culpable one.

We propose that patients should have an opportunity to correct the factual account and explain their reasons. When a contested attribution is proposed for institutional use, review should involve a clinical ethics service or a comparably independent multidisciplinary process.

Gate 3: From responsibility to a legitimate institutional consequence

Gate 3 examines the proposed consequence separately from the responsibility judgment. This separation builds on accounts distinguishing being responsible from justifiably being held responsible through healthcare prioritization[5, 6].

The first question is whether responsibility is actually doing justificatory work. A current access problem may affect the feasibility of treatment and require a prospective clinical assessment. A proposal to reduce priority because the patient previously made an objectionable choice invokes a different principle. The institution must identify and defend that principle rather than allow a judgment of deservingness to reinforce an otherwise clinical assessment.

The proposed consequence must be specified, together with its purpose and the reason responsibility is relevant to that purpose. Further assessment concerns proportionality, consistent application, less

restrictive alternatives and effects on disadvantaged patients. Resource conservation alone does not explain why retrospective blame should determine access to care.

Drawing on accountability for reasonableness, institutional decisions should provide publicly accessible reasons, rely on considerations that fair-minded participants can accept as relevant, allow challenge and revision, and include oversight to ensure these conditions are met[24]. The relationship between these conditions and substantively fair outcomes nevertheless requires further argument[25]. In the present framework, institutional agreement and procedural compliance do not independently establish that responsibility is a legitimate basis for the proposed consequence. That substantive justification must also be supplied.

The accountable body is the institution responsible for the relevant allocation or eligibility policy. An individual clinician's note, a software field or a vendor's design choice should not independently create a responsibility-based allocation rule.

Gate 3 sets necessary constraints without supplying a complete theory of just allocation. Even when a limited responsibility judgment is defensible, the proposed consequence may remain unjustified. In that situation, the responsibility judgment should not be used to justify the proposed change in entitlement or priority. The framework does not authorize delaying appropriate care to investigate blame or introducing responsibility as an additional emergency triage criterion.

Evidence-informed examination through hypothetical cases

The cases distinguish published evidence about relevant clinical circumstances from constructed patient narratives and model outputs. All numerical estimates below are illustrative. They are not findings from the cited studies, outputs of deployed models or evidence that the framework has been clinically validated. For each hypothetical comparison, its stated uncertainty and assumptions remain relevant to any later use.

Case 1: A transplant recipient facing medication-access barriers

Hart and colleagues examined a historical US cohort of 78,861 Medicare-covered kidney-alone transplant recipients aged under 65 years who underwent transplantation between 2008 and 2014. Medication possession was assessed in a linked pharmacy-record subset. Compared with continued coverage, early and late Medicare loss were associated with lower immunosuppressive medication possession ratios and higher hazards of graft loss. Coverage loss around three years after transplantation was not associated with these adverse patterns[26]. These observational findings motivate inquiry into medication-access circumstances, but medication possession does not establish actual medication-taking behavior. The study does not establish individual causation, patient fault or the effects of current coverage policy.

Consider a hypothetical model estimating a five-year graft-failure risk of 20% under a supported medication regimen and 35% under recurrent interruption. After graft failure, this comparison is retrieved during assessment for another transplant. Contextual inquiry establishes that loss of coverage and unreliable transport prevented sustained access, and that the recommended support was not delivered.

The responsibility inference stops at Gate 1. Even assuming adequate methodological warrant for the comparison, the supported regimen was not reasonably available under the circumstances. The estimate also does not establish that this particular graft failure would have been prevented.

The interruption therefore cannot, on these facts, be treated as a culpable refusal. Gates 2 and 3 cannot repair the missing opportunity premise or justify lower priority on that basis. The clinical route remains relevant: present medication access and support require reassessment because they may affect future treatment feasibility. A new support arrangement may change that assessment.

The case supports contextual clinical evaluation and examination of failed institutional support. It does not determine automatic eligibility or automatic exclusion.

Case 2: Informed refusal of burdensome endocrine therapy

Moon and colleagues' qualitative study of 32 breast cancer survivors prescribed tamoxifen describes how participants weighed treatment benefits against side effects and identified information and support needs[27]. In a qualitative study of 30 patients with early-stage breast cancer, Walsh and colleagues described pain-related disruption, unmet management needs and temporary breaks from adjuvant endocrine therapy. Participants identified pain as a reason for treatment breaks, including clinician-directed breaks, rather than complete discontinuation[28]. These findings support eliciting patients' reasons and examining treatment burden. They do not establish the population prevalence of particular reasons, demonstrate that every decision was informed, or determine whether a refusal was morally justified.

Consider a hypothetical treatment-effect estimate of a five-year recurrence risk of 25% without adjuvant endocrine treatment and 15% with treatment. Treatment is accessible and practically feasible. The patient understands the comparison, has discussed relevant support options, and declines because the anticipated residual burdens conflict with a sustained caregiving commitment and the patient's priorities. Recurrence later occurs, and the earlier estimate is invoked in considering access to an expensive therapy. These circumstances, including the caregiving commitment and later allocation proposal, are constructed features of the case.

Conditional on adequate causal warrant and the stipulated feasibility, Gate 1 supports an available opportunity to reduce risk. It does not establish that treatment would have prevented this recurrence. If further inquiry instead showed that treatment demands made implementation impracticable, the opportunity premise would require reconsideration at Gate 1.

Gate 2 asks a different question: whether accepting the feasible treatment could reasonably have been required. Access, understanding and discussion of support do not eliminate the significance of residual burdens and competing obligations. The existence of symptom-management options cannot substitute for assessment of their availability and likely adequacy for this patient.

On the facts given, a blame judgment has not been established. Restricting later care on that basis is consequently unsupported. Present benefit, toxicity and ordinary eligibility criteria remain relevant through clinical assessment. Deliberation should also revisit the patient's current preferences rather than assume that the earlier refusal continues to apply.

The conclusion is specific to the proposed inference. Feasibility and understanding are insufficient to establish culpability; the case does not establish that every informed refusal defeats every possible responsibility judgment.

Contrasting variant: Responsibility without an automatic penalty

Modify Case 2 by stipulating an intervention with negligible burden, a substantial and well-supported reduction in risk, sustained access, adequate understanding and control, and no weighty competing consideration after fair inquiry. For this conceptual comparison, also stipulate acceptance of a limited duty to take a low-burden opportunity to reduce the risk of serious harm to oneself. These opportunity conditions draw on the golden-opportunity account, while the limited duty invoked here remains an explicit normative assumption of the comparison[18].

Under those assumptions, Gate 1 supports the opportunity and Gate 2 can support responsibility for knowingly foregoing it. The conclusion depends on the stated normative premise. A reader who rejects that duty can identify where the proposed responsibility judgment fails.

The stipulations still do not establish responsibility for the particular recurrence. Suppose an institution nevertheless proposes lower priority because the opportunity was knowingly forgone. Gate 3 requires an explanation of why that fact should affect access, together with substantive and procedural justification for the consequence[6]. If those grounds are absent, the proposed action remains unsupported even though the limited Gate 2 judgment stands.

This variant does not claim that an available cancer treatment has the stipulated profile. It demonstrates that resisting an automatic penalty need not depend on denying every form of patient responsibility[5].

Normative relevance migration and safeguards for information reuse

The analysis identifies normative relevance migration as an unjustified change in the role assigned to information. An estimate initially used to discuss treatment benefit may later be treated as proof of

avoidability and then as a reason to restrict access. These transitions require different premises. Reusing the same information does not supply them.

The problem concerns the reasons assigned to information, not only its movement between recipients or systems. A change can occur within one consultation when a prior refusal shifts from helping explain a patient's preferences to supporting a judgment of reduced deservingness. This focus complements contextual integrity by examining the justification of the new use[29]. It does not imply that every change of purpose is impermissible.

Prediction of treatment behavior presents a related problem. Balazard and colleagues examined endocrine-therapy patterns, toxicities and quality of life in the French CANTO cohort. Their prediction analysis included 5,282 patients, divided into training and held-out validation sets; the full-variable models achieved a C-index of 0.62 (95% CI, 0.56–0.68) in the 1,057-patient validation set[30]. These findings link treatment experiences with discontinuation risk, but predictive discrimination does not establish why a particular patient stops treatment or whether that patient is responsible. Even a more accurate model would require additional grounds for those conclusions. A predicted risk of discontinuation may prompt supportive inquiry; it does not by itself justify a judgment of unwillingness or reduced deservingness.

Support also requires evaluation rather than assumption. In a randomized trial of 304 participants, Graetz and colleagues found no statistically significant improvement in the primary outcome of one-year electronic pillbox-captured adherence with the tested app-based interventions compared with enhanced usual care, although the app-plus-feedback group showed improvements in some secondary outcomes[31]. This finding does not establish that monitoring is generally ineffective or that no individual benefited. For the present framework, the relevant implication is that offering an intervention cannot substitute for assessing whether the particular patient's barriers have been adequately addressed.

These studies do not demonstrate that predictive systems cause blame or reduced access. Together with the documentation and vignette studies discussed in the manuscript, they identify distinctions that governance should preserve. The following safeguards are proposals derived from the normative analysis rather than interventions shown to prevent harm.

First, predictive outputs should retain the information needed to interpret their scope: model identity and version, date, target outcome, time horizon, intervention comparison where applicable, relevant validation context and material uncertainty. Later interpretations should remain distinguishable from the original output. An adherence-risk prediction, an estimated treatment benefit and a retrospective responsibility judgment should not be represented as interchangeable conclusions.

Second, a proposed responsibility-sensitive use should identify the responsibility claim, the intended consequence, the additional premises being relied upon and the accountable decision-maker. Where

contextual information is necessary, the record should distinguish the patient's account, established facts and unresolved matters. An earlier label should not silently become an eligibility criterion.

Third, patients should receive an understandable account of the proposed consequential use and have a meaningful opportunity to correct facts, explain circumstances and seek review[24]. Agreement to predictive decision support does not itself authorize later blame-based consequences. Equally, consent does not make an otherwise unjust consequence legitimate.

Finally, these safeguards must not turn routine care into a standing investigation of patient responsibility. Additional inquiry should be justified by a specific proposed use and limited to information necessary to assess it. The framework supplies no general license to collect sensitive information for moral appraisal. A proposed use, or the inquiry needed to support it, may remain objectionable on privacy or other substantive grounds[32].

The resulting governance approach preserves clinical and deliberative uses while requiring additional justification when information is invoked to attribute responsibility or impose an adverse consequence. Its implementation and effects remain questions for subsequent evaluation.

Discussion

The contribution of separating justificatory transitions

The central finding of this analysis is that increasingly individualized information does not collapse the distinctions between an opportunity to reduce risk, responsibility for a choice and the justification of a subsequent institutional response. Each conclusion requires reasons that the preceding conclusion cannot supply. The Gated Relevance Framework makes these dependencies explicit when predictive information is retained and assigned a different role in later care.

The underlying distinctions have substantial antecedents. Schmidt's account of health responsibility separates responsibility from punitive blame and treatment entitlement[5]. Work on golden opportunities examines how supported opportunities to adopt healthier behavior at reasonable cost may address epistemic and control conditions for responsibility-sensitive policies[18]. Clavien and Hurst distinguish practical and moral responsibility, risk-taking choices and illness, and the responsibilities of different parties[33]. The framework does not originate these distinctions. Its contribution is to organize them around the evidentiary scope of predictive outputs and the justification required when their use changes. A privacy objection may rule out an inquiry before contextual evidence is gathered[32].

The three gates correspond to different possible failures. Gate 1 identifies a mismatch between an estimated benefit and the opportunity actually available. Gate 2 identifies an unsupported move from a feasible option to an obligation to accept it. Gate 3 identifies an unsupported move from responsibility to an adverse consequence. Combining these questions into a single assessment of "responsible behavior" would obscure where a proposed inference fails and what kind of reason would be needed to defend it.

Their ordering is logical rather than a mandatory sequence of investigations. A privacy objection may rule out an inquiry before contextual evidence is gathered. An institution may reject responsibility as an allocation criterion without determining whether a patient was responsible. Conversely, evidence considered at a later stage may reveal that an earlier premise needs revision. The framework therefore provides an account of justificatory dependencies, not a linear procedure that all patients should undergo.

Gate 1: An estimated benefit is not yet a patient's opportunity

Gate 1 connects two questions that can otherwise become separated: whether the estimate warrants the proposed inference and whether the alternative was available in the patient's circumstances. Technical improvement addresses the first question only to the extent that the model's inferential target and assumptions match the intended use. Even a well-supported estimate of conditional average treatment benefit does not ordinarily identify the treatment effect for a particular individual[15].

This limitation matters because retrospective interpretation can change the object of the claim. "Treatment would have reduced the estimated risk" may become "the patient could have prevented this outcome." The second statement requires more than the first. Greater personalization can make the alternative trajectory appear more concrete without supplying the missing individual causal warrant. This is a concern about interpretation, not evidence that individualized displays necessarily produce blame.

The opportunity requirement also prevents an ideal treatment pathway from being substituted for the pathway the patient could actually follow. A comparison assuming continuous medication supply, transport and practical support cannot establish that a patient had that opportunity when these conditions were absent. In the transplantation case, the relevant failure is not simply that the model omitted another predictor. It is that the favorable alternative depended on support that was not delivered.

This point complements accounts of scaffolded responsibility, which recognize that individual agency and social support need not be competing explanations[19]. Attending to institutional contributions does not require denying patient agency. It requires identifying the conditions under which agency could be exercised. Equally, identifying an institutional failure does not automatically establish every claim about its causal effects or the responsibility of a particular professional.

The boundary with Gate 2 is important. A treatment burden may prevent implementation, making it relevant to availability, or it may leave treatment feasible while providing a reason why acceptance should not be demanded. These are different questions about the same circumstance. Their separation makes the argument more precise without assuming that practical feasibility and moral reasonableness can always be assessed independently[18].

Gate 2: Medical benefit does not establish a duty to accept treatment

Gate 2 contains the framework's most explicit normative disagreement. Once an opportunity is established, a responsibility judgment still requires an account of why the patient should have taken it. A recommendation based on expected health benefit does not itself establish a moral duty. Neither improved accuracy nor more detailed evidence of understanding can supply that duty.

Opportunity-based accounts provide a serious argument for responsibility in circumstances involving substantial benefit, adequate understanding and limited burden[18]. The contrasting variant preserves that possibility. It shows that a limited responsibility judgment can be supported if a defensible obligation is accepted, rather than making structural constraint or uncertainty defeat every attribution by definition. Andersen and Nielsen raise a more fundamental challenge: accommodating widely accepted responsibility-diminishing effects of genetic and social influences may lead towards broader skepticism about personal responsibility[34]. This is a contested philosophical argument, rather than an established conclusion that patients can never be responsible. The framework makes the implications of this disagreement explicit; it does not resolve it by presenting the relevant questions as a checklist.

The endocrine-therapy case illustrates why the normative premise cannot be omitted. Patients' accounts of tamoxifen use identify trade-offs involving benefits, side effects and support needs[27]. These accounts do not establish that every refusal is justified. These accounts identify considerations beyond understanding a recurrence estimate that require normative assessment. A person may understand the expected benefit and still have a defensible objection to the burdens involved.

Supportive options require the same care in interpretation. Their existence may change what is feasible or reasonable to expect, but an offer of support does not demonstrate that it was accessible, acceptable or adequate. Institutions should therefore avoid treating the provision of information or a monitoring tool as evidence that the remaining difficulty must lie in the patient's motivation.

The scope of responsibility must also remain explicit. Responsibility for accepting a risk and responsibility for a realized outcome are different judgments, not merely stronger and weaker versions of the same judgment[33]. A patient may knowingly forgo an opportunity without the available evidence establishing that this choice caused the later illness. Neither judgment warrants a global characterization of the person as irresponsible.

Finally, attention to patients' reasons must not become a requirement that patients disclose intimate circumstances to prove their deservingness. Silence, limited documentation or difficulty explaining a decision cannot simply be treated as evidence that no defensible reason existed. Because the institution proposes the adverse inference, it carries the burden of supporting it. This allocation of the burden reflects the framework's ethical commitments rather than a claim of neutrality between all accounts of health responsibility[32].

Gate 3: Responsibility does not determine what an institution may do

The independence of Gate 3 is necessary even if the disagreements at Gate 2 are resolved in favor of responsibility. A justified judgment about a past choice does not determine whether reduced priority, an additional condition of access or another consequence is permissible. The institution must explain why that particular consequence follows within a defensible account of healthcare allocation.

This distinction is especially important where clinical and responsibility-based reasons are mixed. Work on elective-surgery restrictions separates prognostic considerations from responsibility-based rationales[35], while analysis of Norwegian policy documents distinguishes prospective benefit arguments from retrospective fairness arguments[36]. A decision can invoke both. In such cases, an acceptable clinical consideration cannot compensate for an unsupported claim about deservingness.

Institutions should therefore identify the contribution made by each reason. If removing the responsibility judgment would change the proposed decision, responsibility is doing substantive work and requires its own defense. If the decision rests entirely on current benefit and feasibility, recording an adverse moral appraisal adds a potentially consequential claim without establishing its necessity. An unchanged decision does not, however, establish that responsibility played no role, because different reasons may independently support the same outcome.

Accountability for reasonableness informs requirements for publicly accessible and relevant reasons, opportunities for challenge and revision, and oversight[24]. These procedures cannot convert an indefensible allocation principle into a defensible one. Arguments based on reciprocity, incentives or stewardship would still need to establish their normative relevance and, where they rely on predicted policy effects, adequate empirical support. A claim that restrictions would encourage future adherence, for example, cannot be established by evidence that nonadherence predicts poorer outcomes.

Consequences also require attention to their effects on future opportunity. Davies's recursion objection identifies how responsibility-sensitive restrictions may contribute to further disadvantage and reduced capacity to protect health[37]. This is a substantive concern about the proposed policy, not something addressed merely by checking whether its rules were consistently applied.

Gate 3 therefore leaves open the possibility that a responsibility judgment is defensible while every proposed adverse response in the setting remains unjustified. The framework's value does not depend on identifying circumstances in which care should be withheld. It can clarify why responsibility fails to justify an action even when disagreement about the patient's earlier choice persists.

Uncertainty, stopping conditions and the legitimacy of inquiry

The stopping conditions could be read as making responsibility attribution impossible whenever uncertainty remains. That is not the intended requirement. Ethical and clinical judgments routinely rely on incomplete evidence. The relevant question is whether the premises are sufficiently supported for the specified conclusion and proposed use, not whether every uncertainty has been eliminated.

Different claims require different support. Evidence may justify discussing an increased risk without establishing individual causation. Contextual information may support a limited judgment about a choice while leaving its institutional relevance unresolved. The framework makes these differences visible rather than proposing a single confidence threshold for all three gates.

Nor does “unresolved” mean that any disagreement automatically vetoes a decision. It means that a necessary premise has not received an adequate defense for the use proposed. The framework does not specify numerical evidentiary thresholds, and the appropriate standards remain a substantive issue for further work. Any practical application would need to explain how sufficiency is judged and prevent unsupported assumptions from being treated as settled facts.

A separate objection concerns whether the inquiry should occur at all. Véliz argues that investigating personal responsibility for healthcare allocation is unjustifiably intrusive and exposes patients to additional privacy risks[32]. Her objection concerns the legitimacy of collecting and using information for this purpose, rather than merely the accuracy or security of the resulting record. Within the present framework, greater accuracy, consent and opportunities for appeal do not by themselves establish that legitimacy; the proposed inquiry still requires an independent justification.

The framework consequently should not be read as authorization to conduct routine responsibility assessments. Its requirements can expose the inadequacy of an existing attribution without licensing an institution to gather whatever information would be needed to strengthen it. If obtaining the necessary evidence would itself be unjustified, the adverse inference remains unavailable. The framework retains a critical role even in settings where responsibility-sensitive allocation is rejected altogether.

Implications for predictive systems and information governance

Normative relevance migration concerns a change in what information is taken to justify. It can occur without a new recipient or a formal transfer of data. A risk estimate or recorded refusal may acquire a moral significance during later interpretation. This focus complements contextual integrity by examining the reasons attached to information within as well as across contexts[29].

Related analyses of LLM-mediated communication examine how automated explanations redistribute interpretive authority and complicate professional accountability, while governance scholarship emphasizes patterns of reliance and institutional responsibility during routine use[38, 39]. The present analysis addresses a more specific question: what additional premises are required when patient-specific information is used to attribute responsibility to a patient or justify an adverse institutional consequence.

The empirical literature provides indirect support for attending to such interpretation. Goddu and colleagues show that stigmatizing notes can influence attitudes and hypothetical treatment choices in a vignette setting[9]. Sun and colleagues document unequal use of negative descriptors in clinical

records[10]. A multicenter study of GPT-4-assisted radiology-report simplification reported improved comprehension scores among volunteer readers and shorter simulated consultation times[40]. Such findings illustrate that the presentation of clinical information can affect understanding and communication, but do not establish effects on responsibility attribution or allocation decisions. Their relevance is to the plausibility of consequential framing, not proof of the complete sequence proposed here.

The governance response should preserve distinctions between model outputs, patient accounts and later judgments. It should also avoid creating permanent moral labels. A record that an attribution was considered during one decision may itself influence later clinicians, even if the proposed consequence was rejected. Collecting more contextual information and adding review fields could therefore introduce new risks alongside potential benefits[32].

Downstream safeguards also do not justify generating every possible counterfactual representation. Before producing or retaining an individualized alternative trajectory, its intended benefit, evidentiary basis and necessity require consideration. Retrospective estimates generated primarily to evaluate blame need justification independently of whether the technology can produce them. A less detailed representation may sometimes answer the clinical question without encouraging claims the evidence cannot support; whether particular designs achieve this remains an empirical question.

Finally, identifying a use as clinical does not establish its fairness. Prognostic criteria can disadvantage patients whose outcomes depend on unequal access to support. The direct clinical route avoids an unnecessary responsibility judgment, but it does not bypass obligations to assess model validity, remediable barriers and the justice of the allocation policy[35].

Limitations and priorities for evaluation

This analysis develops and examines a normative proposal. The 108 included publications were selected purposively and encompass methodological, normative and empirical work. Their number does not establish comprehensive coverage, independent replication or validation of the framework. The limits of the retained search documentation further constrain reconstruction of the retrieval process. Relevant ethical positions or empirical studies may have been missed.

The empirical evidence does not directly establish that patient-specific counterfactual predictions lead to adverse responsibility-sensitive allocation. Documentation studies, surveys, qualitative accounts and intervention studies support different parts of the contextual argument. Their findings cannot be combined into an untested causal chain. Likewise, the hypothetical cases examine inferential dependencies; they do not demonstrate feasibility, effectiveness or generalizability across services.

The framework leaves substantive questions open, including the content of reasonable self-care duties, standards for attributing individual outcomes, the legitimacy of responsibility-sensitive allocation and acceptable burdens of contextual inquiry. Its institutional recommendations also depend on access to review and patient support that may be unevenly available. Implementation could increase

documentation burden, encourage moral surveillance or give punitive decisions an appearance of legitimacy.

Evaluation should therefore examine both intended benefits and possible harms. Initial work could assess whether patients, clinicians and ethics services can distinguish the four uses of information and identify missing premises in standardized cases. Vignette studies could hold clinical information constant while varying access conditions, treatment burdens or the proposed use, and compare the framework with a simpler requirement to provide reasons. This would test whether the three-gate structure adds practical value.

Subsequent workflow evaluation could examine correction of factual errors, patient participation, review burden, delays and unequal effects across groups. A reduction in responsibility judgments alone would not establish success: the relevant outcomes include better-supported reasoning, appropriate access to care and avoidance of unnecessary investigation. Such evaluation would assess whether the proposed structure improves decisions without assuming that its practical benefits follow from the normative argument.

Conclusions

Patient-specific predictive information can inform care without establishing that a patient is responsible for a subsequent adverse outcome or that an adverse institutional response is justified. The Gated Relevance Framework distinguishes three questions: whether an estimated benefit represents a reasonably available opportunity, whether foregoing that opportunity supports a specified responsibility judgment, and whether that judgment justifies the proposed consequence. Each transition requires additional factual and normative grounds.

The framework preserves the distinction between clinical and deliberative uses and responsibility-sensitive uses, placing the burden of justification on those proposing adverse consequences. It provides a structure for identifying unsupported inferences and enabling accountable review, rather than a sufficient justification for restricting care. Its practical value, feasibility and potential unintended effects require empirical evaluation.

Abbreviations

AI
Artificial intelligence
EHR
Electronic health record

Declarations

Ethics approval and consent to participate

Not applicable. This article presents a normative analysis of published literature and hypothetical cases; it reports no research involving human participants, identifiable patient data or human tissue.

Consent for publication

Not applicable. The clinical cases are hypothetical and do not describe identifiable individuals.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by the Talent Fund of Guizhou Provincial People's Hospital (Grant No. 2024-16), the Science and Technology Fund Project of the Guizhou Provincial Health Commission (Grant No. gzwkj2026-295) and the Basic Research Program of the Guizhou Provincial Department of Science and Technology (Qian Ke He Ji Chu) [Grant No. QKH-JC-MS (2025)492]. The funding sources had no role in the study design, analysis, interpretation, writing of the manuscript or decision to submit for publication.

Author Contribution

XY: conceptualization, methodology, formal analysis, investigation, funding acquisition and writing of the original draft. FW: writing of the original draft.LL: writing of the original draft.CX: conceptualization, methodology, review and editing, visualization, supervision, project administration and funding acquisition.

Acknowledgement

ChatGPT (OpenAI) was used solely for language editing and polishing during the revision of this manuscript. The authors retain full responsibility for the arguments, source accuracy, and final text. The tool is not an author and did not generate participant data or empirical results.

Availability of data and materials

No participant-level datasets were generated or reanalyzed. The analysis draws on the published literature cited in the article; the clinical scenarios and model outputs are explicitly hypothetical.

References

1. Desai RJ, Glynn RJ, Solomon SD, Claggett B, Wang SV, Vaduganathan M. Individualized Treatment Effect Prediction with Machine Learning - Salient Considerations. NEJM Evid. 2024;3(4):EVIDoa2300041. 10.1056/EVIDoa2300041.
2. Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, et al. Causal machine learning for predicting treatment outcomes. Nat Med. 2024;30(4):958–68. 10.1038/s41591-024-02902-1.
3. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, et al. Digital twins for health: a scoping review. NPJ Digit Med. 2024;7(1):77. 10.1038/s41746-024-01073-0.
4. Brown RCH, Savulescu J. Responsibility in healthcare across time and agents. J Med Ethics. 2019;45(10):636–44. 10.1136/medethics-2019-105382.
5. Schmidt H. Just health responsibility. J Med Ethics. 2009;35(1):21–6. 10.1136/jme.2008.024315.
6. Baeroe K, Albertsen A, Cappelen C. On the Anatomy of Health-related Actions for Which People Could Reasonably be Held Responsible: A Framework. J Med Philos. 2023;48(4):384–99. 10.1093/jmp/jhad025.
7. Resnik DB. Genetics and Personal Responsibility for Health. New Genet Soc. 2014;33(2):113–25. 10.1080/14636778.2014.905195.
8. Everett JAC, Maslen H, Nussberger AM, Bringedal B, Wilkinson D, Savulescu J. An empirical bioethical examination of Norwegian and British doctors' views of responsibility and (de)prioritization in healthcare. Bioethics. 2021;35(9):932–46. 10.1111/bioe.12925.
9. O'Conor APG, Lanzkron KJ, Saheed S, Saha MO, Peek S. Do Words Matter? Stigmatizing Language and the Transmission of Bias in the Medical Record. J Gen Intern Med. 2018;33(5):685–91. 10.1007/s11606-017-4289-2.
10. Sun M, Oliwa T, Peek ME, Tung EL. Negative Patient Descriptors: Documenting Racial Bias In The Electronic Health Record. Health Aff (Millwood). 2022;41(2):203–11. 10.1377/hlthaff.2021.01423.
11. Yang X, Xiao Y, Liu D, Sun Y, Zhou X, Yao J, et al. Ethical tensions in AI-mediated clinical communication: a mixed-methods study of nursing students’ perspectives. BMC Med Ethics. 2026. 10.1186/s12910-026-01602-4.
12. Kent DM, Steyerberg E, van Klaveren D. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. BMJ. 2018;363:k4245. 10.1136/bmj.k4245.
13. Hoogland J, IntHout J, Belias M, Rovers MM, Riley RD, Harrell E. A tutorial on individualized treatment effect prediction from randomized trials with a binary endpoint. Stat Med. 2021;40(26):5961–81. 10.1002/sim.9154.
14. Wu L, Yang S. Transfer learning of individualized treatment rules from experimental to real-world data. J Comput Graph Stat. 2023;32(3):1036–45. 10.1080/10618600.2022.2141752.
15. Post RAJ, Petkovic M, van den Heuvel IL, van den Heuvel ER. Flexible Machine Learning Estimation of Conditional Average Treatment Effects: A Blessing and a Curse. Epidemiology. 2024;35(1):32–40. 10.1097/EDE.0000000000001684.

16. Balakrishnan A, Weinmeyer R, Serper M, Bailey SC, Kaiser K, Wolf M. Navigating medication-taking after kidney transplant. Res Social Adm Pharm. 2022;18(10):3846–54. 10.1016/j.sapharm.2022.05.010.
17. Wester JR, Wagner RB, Motladiile B, Nkele I, Chinyepi N, Makhema MJ, et al. Access to and Utilization of Endocrine Therapy Among Breast Cancer Survivors in Botswana. JCO Glob Oncol. 2024;10:e2400180. 10.1200/GO.24.00180.
18. De Marco G, Douglas T, Savulescu J, Healthcare. Responsibility and Golden Opportunities. Ethical Theory Moral Pract. 2021;24(3):817–31. 10.1007/s10677-021-10208-1.
19. Levy N, Savulescu J. The Myth of Zero-Sum Responsibility: Towards Scaffolded Responsibility for Health. J Moral Philos. 2023;21(1–2):85–105. 10.1163/17455243-20233725.
20. Franzoi MA, Agostinetto E, Perachino M, Del Mastro L, de Azambuja E, Vaz-Luis I, et al. Evidence-based approaches for the management of side-effects of adjuvant endocrine therapy in patients with breast cancer. Lancet Oncol. 2021;22(7):e303–13. 10.1016/S1470-2045(20)30666-5.
21. Appelbaum PS. Clinical practice. Assessment of patients' competence to consent to treatment. N Engl J Med. 2007;357(18):1834–40. 10.1056/NEJMcp074045.
22. Dima AL, Allemann SS, Dunbar-Jacob J, Hughes DA, Vrijens B, Wilson IB. Methodological considerations on estimating medication adherence from self-report, electronic monitoring and electronic healthcare databases using the TEOS framework. Br J Clin Pharmacol. 2023;89(7):1918–27. 10.1111/bcp.15375.
23. Kuba S, Ishida M, Nakamura Y, Taguchi K, Ohno S. Persistence and discontinuation of adjuvant endocrine therapy in women with breast cancer. Breast Cancer. 2016;23(1):128–33. 10.1007/s12282-014-0540-4.
24. Daniels N, Sabin JE. Accountability for reasonableness: an update. BMJ. 2008;337:a1850. 10.1136/bmj.a1850.
25. Rid A. Justice and procedure: how does accountability for reasonableness result in fair limit-setting decisions? J Med Ethics. 2009;35(1):12–6. 10.1136/jme.2008.024430.
26. Hart A, Gustafson SK, Wey A, Salkowski N, Snyder JJ, Kasiske BL, et al. The association between loss of Medicare, immunosuppressive medication use, and kidney transplant outcomes. Am J Transpl. 2019;19(7):1964–71. 10.1111/ajt.15293.
27. Moon Z, Moss-Morris R, Hunter MS, Hughes LD. Understanding tamoxifen adherence in women with breast cancer: A qualitative study. Br J Health Psychol. 2017;22(4):978–97. 10.1111/bjhp.12266.
28. Walsh EA, Chabria R, Vranceanu AM, Park ER, Post K, Peppercorn J, et al. Understanding pain related to adjuvant endocrine therapy after breast cancer: A qualitative report. Eur J Cancer Care (Engl). 2022;31(6):e13723. 10.1111/ecc.13723.
29. Nissenbaum H. Privacy as contextual integrity. Wash L Rev. 2004;79:119.
30. Balazard F, Bertaut A, Bordet E, Mulard S, Blanc J, Briot N, et al. Adjuvant endocrine therapy uptake, toxicity, quality of life, and prediction of early discontinuation. J Natl Cancer Inst. 2023;115(9):1099–108. 10.1093/jnci/djad109.

31. Graetz I, Hu X, Kocak M, Krukowski RA, Anderson JN, Waters TM, et al. Remote Monitoring App for Endocrine Therapy Adherence Among Patients With Early-Stage Breast Cancer: A Randomized Clinical Trial. JAMA Netw Open. 2024;7(6):e2417873. 10.1001/jamanetworkopen.2024.17873.
32. Veliz C. Not the doctor's business: Privacy, personal responsibility and data rights in medical settings. Bioethics. 2020;34(7):712–8. 10.1111/bioe.12711.
33. Clavien C, Hurst S. The Undeserving Sick? An Evaluation of Patients' Responsibility for Their Health Condition. Camb Q Healthc Ethics. 2020;29(2):175–91. 10.1017/S0963180119000975.
34. Andersen MM, Nielsen ME. Personal Responsibility and Lifestyle Diseases. J Med Philos. 2016;41(5):480–99. 10.1093/jmp/jhw015.
35. Pillutla V, Maslen H, Savulescu J. Rationing elective surgery for smokers and obese patients: responsibility or prognosis? BMC Med Ethics. 2018;19(1):28. 10.1186/s12910-018-0272-7.
36. Traina G, Feiring E. Priority setting and personal health responsibility: an analysis of Norwegian key policy documents. J Med Ethics. 2022;48(1):39–45. 10.1136/medethics-2019-105612.
37. Davies B. Responsibility and the recursion problem. Ratio (Oxf). 2022;35(2):112–22. 10.1111/rati.12327.
38. Yang X, Liang Y, Liu Y, Liu D, Yang L, Zhang B, et al. When machines explain medicine: Nursing ethics and clinical communication. Nurs Ethics. 2026;0(0):9697330261449460. 10.1177/09697330261449460.
39. Yang X, Xu C. Ethical Governance of Large Language Models in Health Care: Trust, Responsibility, and Equity in Routine Use. J Med Internet Res. 2026;28:e93470. 10.2196/93470.
40. Yang X, Xiao Y, Liu D, Shi H, Deng H, Huang J, et al. Enhancing Physician-Patient Communication in Oncology Using GPT-4 Through Simplified Radiology Reports: Multicenter Quantitative Study. J Med Internet Res. 2025;27:e63786. 10.2196/63786.

Figures

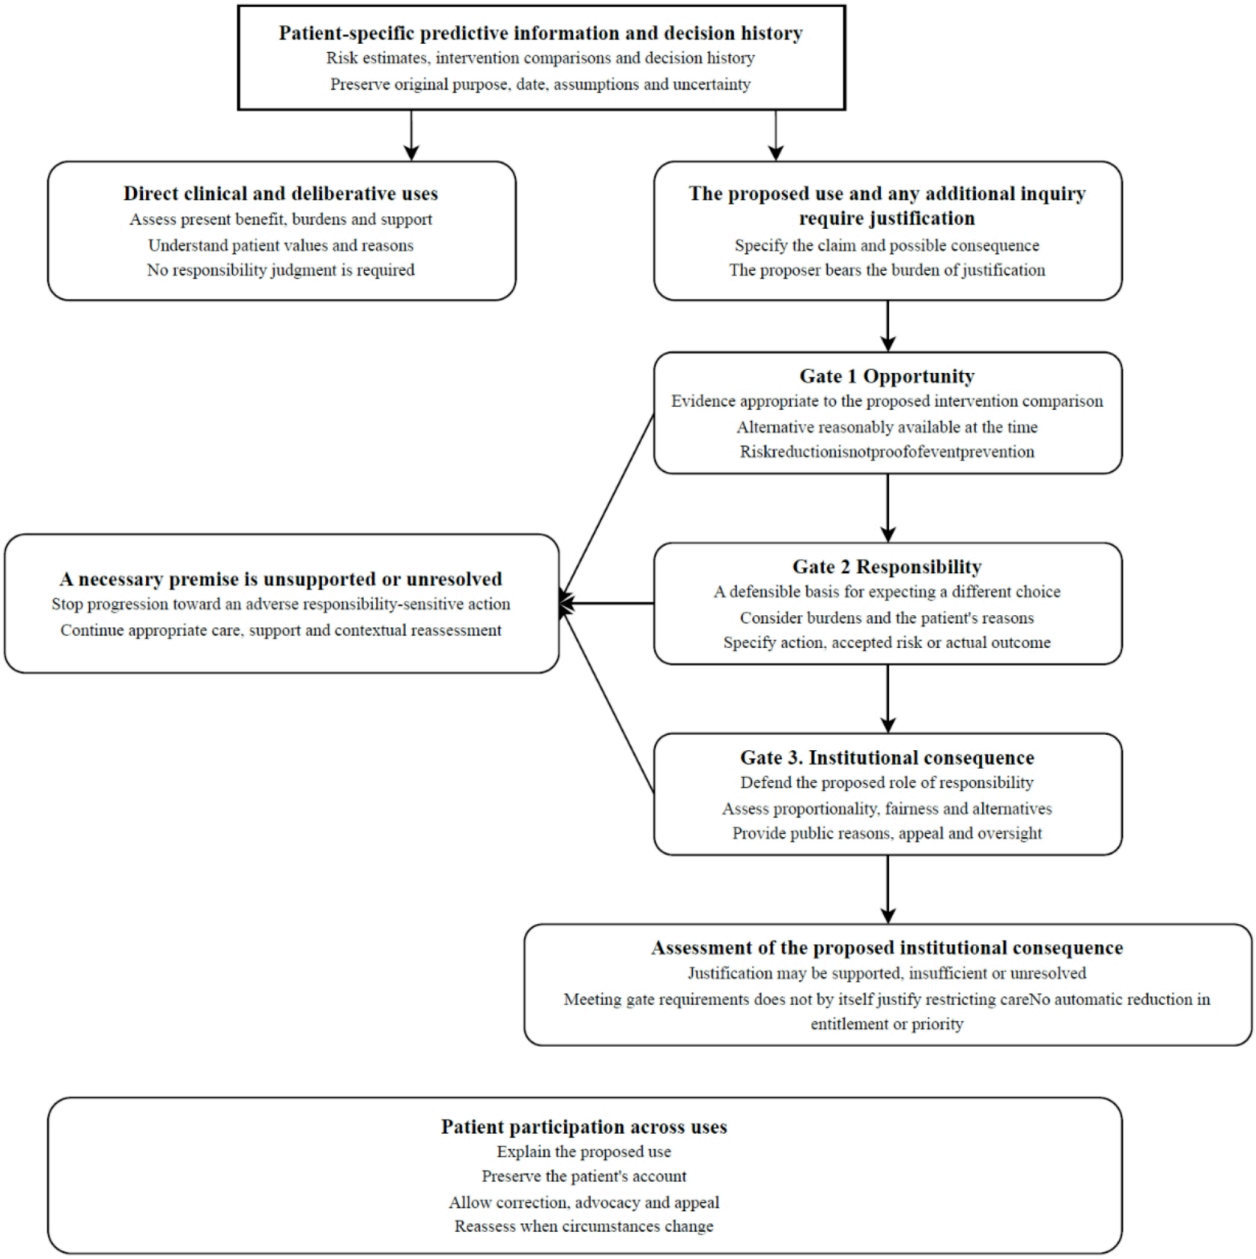


Figure 1

Distinct uses of predictive information and justificatory requirements of the Gated Relevance Framework.

Patient-specific predictive information may support clinical and deliberative uses without responsibility attribution. Responsibility-sensitive uses require additional grounds concerning a reasonably available opportunity, a specified responsibility judgment and the proposed institutional consequence. Arrows indicate justificatory dependencies rather than a mandatory sequence of patient assessments or

automatic authorization to proceed. An unsupported or unresolved necessary premise blocks the proposed adverse inference or action, not appropriate care or support. Responsibility for a choice or accepted risk does not establish responsibility for a particular outcome, which requires additional causal warrant. Meeting the gate requirements does not by itself provide sufficient justification for restricting care. Patient participation and opportunities for correction and review apply across uses; neither pathway exempts the proposed use from requirements of validity, fairness and privacy.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Additionalfile1.pdf
- Additionalfile2.pdf