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.

When a better future becomes visible: digital twins, normative authority, and care after refusal

Xiongwen Yang

yangxiongwen@gz5055.com

Guizhou Provincial People's Hospital

Di Liu

Guizhou Provincial People's Hospital

Lin Yang

Guizhou Provincial People's Hospital

Bo Zhang

Guizhou Provincial People's Hospital

Chuan Xu

Guizhou Provincial People's Hospital

Yi Xiao

Third Affiliated Hospital of Sun Yat-sen University

Article

Keywords: Digital twins, Normative authority, Counterfactual reasoning, Informed refusal, Moral responsibility, Continuing care

Posted Date: November 3rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11060289/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

Simulation-oriented digital twins may enable comparisons of patient-specific health trajectories under alternative treatments or courses of action. Where these representations remain available after a decision, an alternative originally generated to inform choice may later become a reference point for evaluating refusal. This article examines what normative authority an unrealised but medically preferable simulated future should retain when an adverse outcome occurs and the patient requires further care. Through conceptual and normative analysis, it considers cases in which the simulation is sufficiently credible for deliberation and the patient has adequate understanding, decision-making capacity, and a reasonable opportunity to pursue the recommended option. Three connected limits are defended. Epistemic credibility does not by itself determine what the patient ought, all things considered, to choose. Medical preferability does not by itself establish a personal obligation to accept the associated option. Neither the occurrence of an adverse outcome nor the continued visibility of the alternative supplies the normative justification needed to convert refusal into moral failure. Independently justified reassessment remains possible, but retrospective salience cannot substitute for that justification. Continuing care also rests on the patient's standing as a person entitled to respectful clinical consideration, which is not earned through compliance. Past choices may inform current treatment without alone justifying blame or diminished moral standing. Care should remain responsive to present need, support renewed deliberation, and preserve relevant context without requiring patients to justify themselves as a condition of respect. Digital twins thus illuminate a broader problem concerning how technologically represented futures can inform human agency without becoming authoritative verdicts on past choices or present claims to care.

1. Introduction

Medicine has long sought to anticipate what may happen to a patient. Simulation-oriented digital twins may extend this endeavour by representing how an individual patient's health could develop under alternative treatments or courses of action. By combining patient-specific data with computational models, these systems may support comparisons of disease trajectories and treatment responses before their consequences unfold (Björnsson et al., 2019; Katsoulakis et al., 2024; Sun et al., 2023; Tudor et al., 2025). Their ethical significance also concerns what happens to these representations after a decision. Where a simulated alternative remains available, a future generated to inform choice may later become a reference point for judging the choice already made.

Consider a hypothetical patient offered a treatment that a sufficiently credible simulation indicates would improve expected survival but impose substantial burdens on everyday life. The patient understands the evidence, has a reasonable opportunity to receive treatment, and declines it because of those burdens. If the patient later deteriorates, the earlier simulation may remain available as a representation of the healthier future that treatment might have offered. It may then appear to support a judgment that the patient should have chosen differently. Yet the occurrence of the adverse outcome does not itself establish either that treatment would have prevented it or that accepting treatment was

obligatory. The ethical question concerns what the unrealised alternative can legitimately justify once the patient requires further care.

Existing scholarship on digital twins in medicine addresses privacy, consent, fairness, representation, autonomy, responsibility, and pressures towards health optimization (Bruynseels et al., 2018; Burr et al., 2026; Iqbal et al., 2022; Katsoulakis et al., 2024; Popa et al., 2021). These discussions provide important resources for assessing how simulations inform clinical decisions. Building on them, this article examines how the normative status of a simulated future may change as it moves from evidence within deliberation to a retrospective comparator after refusal. The contribution is to connect questions of evidence, obligation, and continuing care by following an alternative future across the decision it was generated to inform.

Two distinctions are necessary to examine this transition. An epistemically credible simulation may provide genuine reasons for action and strengthen a clinical recommendation without settling what the patient ought, all things considered, to choose. Likewise, a future that is preferable according to specified clinical outcomes need not be preferable in every respect that matters to the patient. Health is an important good, but its pursuit may compete with relationships, commitments, independence, and the burdens a person is willing to bear (Buyx, 2008; Entwistle et al., 2010; Resnik, 2007). Normative authority here concerns the stronger justificatory standing attributed to a simulated future when it is treated as determining what the patient ought to choose, rather than as evidence informing that judgment.

These distinctions remain relevant after refusal. An unrealised healthier future may help clinicians understand the patient's condition or reconsider available options. It may also invite a movement from the claim that an outcome might have been avoided to the claim that the patient ought to have avoided it. Debates about responsibility for health show why causal avoidability and moral responsibility require separate consideration (Entwistle et al., 2010; Sherwin and Winsby, 2011; Yang et al., 2026a). Patient-specific simulation gives this distinction a temporal setting because an alternative presented before choice may remain visible when the patient's actual outcome is known. Its greater salience at that point does not, by itself, supply the normative premise needed to justify blame.

This article therefore asks what normative authority, if any, an unrealised but medically preferable simulated future should retain after informed refusal. It develops a conceptual and normative argument, considering cases in which the simulation is sufficiently credible for deliberation and the patient has adequate understanding, decision-making capacity, and reasonable opportunity to pursue the recommended option. These assumptions allow the argument to address refusal without explaining it through model failure or compromised agency. The analysis is structured around three claims. Epistemic credibility does not by itself confer normative authority; medical preferability does not by itself establish personal obligation; and retrospective salience does not by itself confer retrospective normative authority. The argument examines why a later adverse outcome cannot, by itself, resolve the justification of an earlier choice.

The argument then examines what these limits mean for continuing care, where earlier choices may remain clinically relevant without independently warranting blame or reduced moral standing. Its broader significance concerns the relationship between technologically represented futures and human agency. Improved knowledge of possible consequences can inform a choice without determining how a person ought to live. Digital twins provide a focused case in which to examine how that distinction should be sustained after a decision, when an unrealised alternative remains visible alongside present need.

2. From Prediction to Simulated Patient Futures

The term digital twin covers heterogeneous technologies with different capacities for representation, monitoring, prediction, simulation, and decision support (Drummond and Gonsard, 2024; Iqbal et al., 2025; Katsoulakis et al., 2024; Tudor et al., 2025). This article focuses on applications that generate patient-specific alternative trajectories for clinical deliberation and allow those alternatives to remain available after a decision. The argument concerns these capabilities rather than the digital-twin label alone.

The relevant distinction is between estimating an outcome and comparing how outcomes might differ under alternative interventions. A model may estimate a patient's risk of deterioration without establishing how that risk would change if the patient received a particular treatment. Simulation-oriented digital twins seek to support comparisons of possible trajectories under specified treatments, preventive strategies, or other conditions (Björnsson et al., 2019; Katsoulakis et al., 2024). However, changing an input in a predictive model does not necessarily identify a causal intervention effect (Feuerriegel et al., 2024). Although emerging causal digital twins seek to support intervention-sensitive comparisons, valid individualised counterfactual inference should not be treated as an established property of digital twins as a class (Vallée, 2026).

A simulated patient future therefore refers here to a model-generated representation of how a particular patient's health may develop under specified conditions, actions, or interventions. Its evidential status may range from an exploratory scenario to a well-validated estimate. The analysis does not assume that current systems can reveal what would have happened to an individual under an unchosen option. Instead, the subsequent normative argument considers simulations sufficiently credible to inform deliberation, so that its conclusions do not depend on dismissing their predictions as unreliable. Even under this assumption, a simulated trajectory remains an estimate rather than a guarantee of an individual outcome.

Within this scope, four features define the applications at issue. First, the representation concerns an individual patient. Second, it permits comparison between alternative trajectories. Third, those alternatives bear on a clinical choice. Fourth, an alternative presented before the decision remains accessible afterwards. The first three features make simulation relevant to deliberation; the fourth enables an unchosen alternative to enter later assessments of the patient's actual course. Persistence is thus a condition of the problem examined here, not a capability attributed to every digital twin.

How an alternative persists also matters. A retained simulation records what the model represented at the time of choice. A later reconstruction may reproduce that comparison from the original information, whereas an updated simulation may incorporate evidence unavailable when the patient decided. These representations cannot be treated as interchangeable accounts of what the patient could have known. The central case considered here is an alternative presented during deliberation that remains available after refusal. If it is subsequently revised or reconstructed, information introduced later must be distinguished from the evidence available at the original decision. An updated comparison may inform present care, but cannot simply be treated as evidence that the patient previously knew what it now suggests.

The expression medically preferable future also has a restricted meaning. It denotes a simulated trajectory judged preferable according to specified biomedical or clinically recognised outcomes, such as longer expected survival, lower morbidity, or better disease control. Where outcomes conflict, the relevant criterion of preference must be identified rather than assuming a single medically best trajectory. Medical preferability does not mean that a future is better for the patient all things considered, nor does it imply an obligation to pursue it. The term identifies the comparison under examination while leaving its broader normative significance open.

Neither individualised prediction nor counterfactual reasoning is unique to digital twins. Their relevance here lies in the possibility of making an alternative explicit for a particular patient before choice and keeping it available afterwards. The problem may also arise in other technologies with these features, but digital twins remain the focal case. The question is how a representation generated to inform a decision may later be used to evaluate that decision, and what further justification is required when its clinical relevance is taken to establish what the patient ought to have chosen.

3. The Temporal Gap in Digital-Twin Ethics

Existing scholarship on digital twins in healthcare provides several starting points for examining the authority of simulated patient futures. Three strands are particularly relevant to this inquiry. They concern the limits of computational representation, the relationship between personalised recommendations and autonomy, and the distinction between causal contribution and moral responsibility. Read together, they help explain why a simulation may inform clinical judgment without settling what a patient ought to choose or how an earlier refusal should be evaluated.

Work on representation cautions against equating a detailed model of a patient with an adequate account of the person. Biological and clinical specificity may coexist with limited attention to relationships, lived experience, commitments, and the meanings attached to particular outcomes (Braun, 2021; Ferlito et al., 2024; Huang et al., 2022). This matters because the authority attributed to a simulation may extend beyond the questions it is equipped to answer. For the present argument, however, the issue is not exhausted by representational incompleteness. Even a model that incorporates

relevant patient preferences still requires a normative justification for treating its preferred trajectory as determinative of what the patient ought to choose.

A second strand examines how digital twins may support or constrain autonomy. Simulations can clarify consequences and make alternatives easier to compare, while personalised recommendations may also narrow the perceived space of legitimate disagreement (Bruynseels et al., 2018; Huang et al., 2022; Iqbal et al., 2022). These concerns motivate a distinction between improving the informational conditions of choice and establishing the authority to settle that choice. The question pursued here is what follows when a patient understands and accepts the evidence but nevertheless declines the medically preferable option. Adequately informed deliberation may clarify the consequences of a choice without producing agreement about what should be chosen.

A third strand concerns responsibility for health. This literature distinguishes causal contribution, avoidability, opportunity, and moral responsibility, with particular attention to social conditions that constrain individual action (Cappelen and Norheim, 2005; Schwan, 2021; Wikler, 2002; Yang et al., 2026b). These distinctions provide a basis for questioning the inference from a preventable outcome to personal blame. The present analysis examines whether that inference is justified even when the patient has adequate understanding and a reasonable opportunity to pursue the recommended option. These conditions do not, by themselves, establish an obligation to accept it. Without an additional normative premise, they therefore cannot establish that refusal constitutes a moral failure.

The contribution developed here is to connect these questions across the course of a decision and its aftermath. A simulation retained from the original deliberation provides evidence about possible consequences before choice and remains a representation of an unchosen alternative after refusal. If an adverse outcome subsequently occurs, it may be invoked to evaluate the earlier decision. The relevant question then shifts from what the simulation gives the patient reason to consider to what it authorises others to conclude about the patient's refusal. This shift can occur without any improvement in the evidence originally available. It therefore requires an account of justification, not merely an account of predictive accuracy.

The argument proceeds by distinguishing epistemic credibility from authority to determine what should be chosen, before examining why medical preferability, even when accompanied by realistic opportunity, does not by itself establish personal obligation. These distinctions provide the basis for asking whether an adverse outcome and the continued visibility of the alternative can supply the justification that was absent prospectively. The analysis then turns to continuing care, distinguishing the clinical relevance of past choices from judgments about the patient's moral standing. Figure 1 brings these stages together. The aim is to explain how the justificatory limits of a simulated future should be maintained as its role changes across clinical time.

4. From Epistemic Credibility to Normative Authority

An immediate concern about digital-twin-generated futures is their evidential reliability, since a simulated trajectory may be inaccurate, poorly validated, causally misleading, or based on incomplete data (Drummond and Gonsard, 2024; Feuerriegel et al., 2024; Tudor et al., 2025). These concerns matter, but resolving them would not settle the normative question. Even if a simulation provides sufficiently credible estimates of the consequences of alternative actions, what does that credibility justify beyond confidence in those estimates?

Assume, for the purposes of the argument, that the model has been appropriately validated for the clinical use in question, uncertainty is adequately characterised, and its estimates of intervention effects are causally credible. Its predictions are sufficiently reliable to inform deliberation, without guaranteeing any individual outcome. Under these conditions, the simulation warrants substantial epistemic weight in assessing what is likely to happen under specified alternatives. Whether that weight also authorises a conclusion about what the patient ought to choose is a further question. The distinction between authority concerning clinical knowledge and authority over treatment decisions provides a useful starting point (Landmark et al., 2015).

Credible evidence has normative relevance. It can clarify consequences, reveal trade-offs, strengthen recommendations, and improve the informational conditions of autonomous choice (Entwistle et al., 2010; Montori et al., 2013; Parker and Chin, 2020). Evidence that a treatment is likely to prevent serious harm can strengthen a reason to accept it. The issue is whether the strength of the evidence also determines the weight of that reason relative to other considerations. A simulation may indicate that treatment A offers longer expected survival, whereas treatment B is more likely to preserve a function central to the patient's everyday life. Greater predictive accuracy can clarify this trade-off without determining how the competing outcomes should be valued.

This limit also applies when a model ranks the alternatives. A ranking depends on the outcomes selected, the weights assigned to them, and the burdens or constraints included in the comparison. Validation of the model's predictions does not, by itself, justify treating those evaluative choices as decisive for the patient. Incorporating the patient's expressed priorities may make the comparison more useful, but does not by itself justify treating the resulting ranking as decisive. The relevant question is whether the model adequately represents those priorities and whether its way of weighing them should govern the decision.

As uncertainty decreases, this boundary may become harder to recognise because disagreement with a recommendation may appear to require disagreement with the evidence supporting it. Yet a patient may accept the estimated consequences while assigning different importance to survival, function, treatment burden, or other goods. The role of patient preferences in evidence-based decision-making helps explain why agreement about likely outcomes need not produce agreement about treatment (Hoffmann et al., 2014; Montori et al., 2013; Yang and Xu, 2026). Where the patient understands and accepts the evidence, continued disagreement cannot simply be presumed to reflect an informational deficit.

Recognising this limit does not require clinical neutrality. Clinicians may recommend strongly on the basis of reliable evidence, explain the reasons for their recommendations, and help patients assess the consequences of available options (Epstein et al., 2004). Such recommendations combine evidence with evaluative judgments about benefits, burdens, and the patient's circumstances. Those judgments may be defensible, but their justification cannot be supplied by predictive accuracy alone. Normative authority, in the stronger sense used here, concerns the standing to settle what the patient ought, all things considered, to choose.

Proposition 1

Greater epistemic credibility or personalisation of a simulated health future does not by itself confer normative authority to determine what the patient ought, all things considered, to choose.

This proposition preserves the distinction between evidence that supports a recommendation and justification that makes it decisive. A simulated future may warrant serious consideration without settling the choice. A further question is whether its medical preferability, when the patient understands the evidence and has a realistic opportunity to pursue the associated option, is sufficient to establish a personal obligation to do so.

5. The Freedom to Decline a Medically Better Future

The question becomes more demanding when the simulation is credible, the recommendation medically well supported, and the patient understands the evidence but declines the recommended option. The issue is now whether the expected medical advantage, together with a realistic opportunity to pursue it, establishes an obligation to accept that option. A patient may have strong reasons to protect their health without those reasons necessarily being decisive.

The right of a decisionally capable patient to refuse recommended care provides an important starting point(Kleinman, 1991). Yet a right against unwanted intervention does not establish that every refusal is wise or morally justified. It can constrain professional interference while leaving open whether the patient ought to have chosen differently. The question examined here is whether medical preferability alone can establish an obligation, given the other goods and burdens at stake.

To isolate this question, the argument pairs the assumption of a credible model with an assumption of adequate agency. The patient has sufficient information, decision-making capacity, freedom from controlling influence, and a reasonable opportunity to pursue the recommended option. The patient also understands and accepts the simulation's estimates, including their uncertainty. Refusal in this case cannot be explained simply by misunderstanding, coercion, incapacity, or lack of access. These assumptions do not imply that such barriers are unimportant in practice. They allow the analysis to ask whether an obligation follows even when those barriers are absent. The opportunity concerns access to an intervention or course of action, not the ability to guarantee its predicted outcome.

Medicine has strong reasons to value survival, disease prevention, relief of suffering, and preservation of function. People also organise their lives around relationships, responsibilities, independence, identity, and meaningful activities (Entwistle et al., 2010; May et al., 2009; Nordenfelt, 2013). Pursuing better health may support these goods, but it can also require burdens or sacrifices that affect them (Björnsson et al., 2019; Katsoulakis et al., 2024; Mair and May, 2014; May et al., 2009). A clinically optimized trajectory and a life worth choosing are overlapping but non-identical objects of evaluation (Hausman, 2015; Nordenfelt, 2013).

Consider again a treatment expected to prolong survival while imposing substantial burdens on daily life. A patient may understand its benefits yet reasonably give greater weight to maintaining independence or fulfilling responsibilities that treatment would disrupt. Another may accept shorter expected survival to preserve a function central to a valued activity. Such choices can reflect an informed assessment of competing goods rather than ignorance of medical consequences. Accounts of person-centred care provide grounds for taking these priorities seriously within deliberation (Entwistle and Watt, 2013; Epstein and Street, 2011). Their relevance does not make every expressed preference decisive, but it prevents biomedical benefit from being treated as the only consideration requiring justification.

Incorporating these priorities into the model can improve the comparison without eliminating the evaluative question. Measures of quality of life, treatment burden, and expressed preferences may help make recommendations more responsive to the patient (May et al., 2009; Montori et al., 2013). Nevertheless, priorities may conflict, resist comparison on a single scale, or change as illness unfolds (Sprangers and Schwartz, 1999). A ranking that reflects previously expressed preferences remains open to reconsideration. The patient need not regard a model's representation of those preferences as a final determination of what matters now.

The claim that a patient can pursue an option should therefore be distinguished from the claim that they have strong reasons to pursue it and from the further claim that they are obligated to do so. Feasibility removes one possible obstacle to obligation; it does not establish obligation. Strong health-related reasons may support a forceful recommendation, but whether they establish a duty depends on their weight relative to competing considerations and on a further account of what the patient owes to themselves or others. Neither feasibility nor medical benefit, taken alone or together, supplies that account.

This distinction gives substantive importance to informed disagreement. Respect for autonomy requires room to consider that a refusal may express a defensible ordering of goods, rather than merely allowing the patient to decline a course already assumed to be obligatory (Entwistle et al., 2010; Olthuis et al., 2014). Clinicians may explain risks, challenge misunderstandings, explore the patient's reasons, and recommend strongly. The argument limits the inference from medical non-optimality to failure; it does not require agreement with every refusal. Risks to others, obligations arising from particular relationships, resource constraints, and professional responsibilities may introduce additional

considerations. Cases involving impaired capacity, coercion, or serious misunderstanding fall outside the assumptions adopted here.

Proposition 2

The medical preferability of a simulated future, even when the option associated with it is realistically available and supported by strong reasons, does not by itself establish that the patient is obligated to choose that option, because biomedical optimisation does not exhaust the goods relevant to a person's life.

Proposition 2

identifies a limit on what has been justified at the time of choice. It does not establish that every refusal is justified or that no obligation could be established on further grounds. Its temporal significance emerges when the recommended option is declined, the simulated alternative remains available, and an adverse outcome subsequently occurs. The alternative may then become more salient without that outcome alone resolving the earlier question of obligation. The question is whether this change in perspective can legitimately support a stronger judgment about what the patient should have chosen.

6. From Prospective Choice to Retrospective Judgment

The temporal problem arises when a patient declines the recommended option and the simulated alternative remains available as the actual course unfolds. Neither predictive credibility nor medical preferability alone determines what the patient is obligated to choose. If an adverse outcome subsequently occurs, however, the earlier simulation may acquire a different significance. A representation introduced to inform deliberation may now be used to evaluate the refusal.

A simulation presented before a decision may remain available after the patient has declined the recommended treatment. Prospectively, it indicates that accepting treatment may offer a better expected health outcome. Following refusal and deterioration, it may be read as a representation of the healthier future the patient might have had. The original estimate need not have changed for its use to change. The question shifts from what the patient has reason to consider to what the outcome is taken to reveal about the earlier choice.

Counterfactual thinking and its relationship to causal attribution and hindsight are well-established subjects of inquiry (Roese, 1997; Roese and Olson, 1996). Digital twins do not create the possibility of asking whether an outcome might have been avoided. Their significance in the case examined here is that the alternative was individualised, made explicit before the decision, and communicated to the patient. When accompanied by documentation of the discussion, the retained simulation may make both the recommended option and the patient's recorded understanding readily available for later review. This provides a specific setting in which evidence supporting informed choice can later be invoked in support of blame.

Two distinctions are necessary. First, an adverse outcome does not by itself verify the unchosen alternative. Deterioration following refusal may be consistent with the original risk estimate, but it does not establish that accepting treatment would have prevented this patient's deterioration. An updated simulation may also incorporate information unavailable at the time of choice. Such information may improve present clinical understanding, but it cannot simply be attributed to the patient's earlier knowledge. The persistence of a comparison and the credibility of its causal interpretation remain separate matters.

Second, even a well-supported causal comparison does not establish a corresponding obligation. Suppose there are sound reasons to conclude that the recommended intervention would probably have reduced the risk of the adverse outcome. That conclusion may help explain the patient's condition and inform subsequent treatment. It does not alone show that the patient was obligated to accept the intervention, including its burdens and competing consequences. Scholarship on responsibility for health provides grounds for distinguishing causal contribution from judgments of responsibility and blame (Friesen, 2018; Schwan, 2021; Wikler, 2002). To characterise refusal as a breach of duty requires a justified account of the duty allegedly breached; to attribute blame may require further judgments about the patient's knowledge, control, and circumstances.

The occurrence of the adverse outcome cannot, by itself, supply that normative account. Outcome bias research shows that knowledge of an outcome can affect evaluations of an earlier decision even when the information available to the decision-maker is held constant (Baron and Hershey, 1988). The normative issue here is what would justify such a change in evaluation. If the original case for treatment established a medical advantage but did not establish an obligation to accept it, the subsequent harm does not fill that justificatory gap. The realization of a medically worse future cannot, by itself, retrospectively supply an obligation that was not established prospectively.

This claim concerns the insufficiency of the outcome as a justification, not the impossibility of reassessing an earlier decision. Failure to establish an obligation during deliberation does not prove that no obligation existed. Later inquiry may reveal facts relevant to an independently justified obligation, or identify reasons that were previously overlooked. Any revised judgment must nevertheless explain the obligation and its application to the circumstances of choice. The occurrence of the predicted harm alone cannot substitute for that explanation. Likewise, new information may warrant a different recommendation now without showing that the earlier refusal was wrong.

The continued visibility of the simulation, together with documentation of the patient's understanding, can obscure these distinctions. Understanding the relevant risks supports the validity of an autonomous refusal (Entwistle et al., 2010; Kelley, 2005). After an adverse outcome, the same understanding may be invoked through the assertion that the patient "knew what would happen." This formulation collapses knowledge of a risk into knowledge of a certain outcome. It also treats awareness of a possible consequence as sufficient reason for blame. Knowledge can matter to responsibility where an obligation has been independently justified, but knowledge of a risk does not itself establish an obligation to avoid

it. Better-informed refusal therefore need not become more culpable refusal merely because the accepted risk materialises.

What changes after deterioration may be the emotional force, practical relevance, and prominence of the unrealised alternative. It may become harder for the patient or clinicians to set the comparison aside. The alternative can also remain clinically useful, provided its evidential limitations are respected. None of this automatically increases its authority to establish what the patient should have chosen. The normative vulnerability lies in allowing the transition from an anticipated risk to an experienced harm to make a previously unsupported judgment appear justified.

Proposition 3

The persistence and retrospective salience of a knowingly declined, medically preferable simulated future do not by themselves transform the earlier refusal into moral failure or diminish the moral standing of subsequent clinical need. Treating subsequent need as less worthy of a response requires a further argument about why past conduct should alter the patient's claims within care.

The two parts of this proposition concern distinct judgments. An allegation of wrongdoing requires justification of the obligation allegedly breached. A reduction in the moral standing of subsequent need requires a further argument about why past conduct should alter the patient's claims within care. A persistent simulation supplies neither argument by itself. This does not imply that previous choices are irrelevant to prognosis, treatment feasibility, or expected benefit. It means that their clinical relevance cannot simply be treated as a verdict on what care the patient deserves. The coexistence of present need and a visible representation of an unchosen alternative raises a further question about the positive responsibilities of continuing care.

7. Care after refusal: an ethics of unrealised futures

When a patient returns for care after declining a recommended option, an earlier simulation may remain available alongside the patient's present condition. The alternative may help clinicians understand the course of illness or assess what remains possible. Its persistence and retrospective salience do not, by themselves, justify blame or diminish the moral standing of subsequent need. A positive account of continuing care must also explain why clinicians should remain responsive to that need and how earlier decisions should inform their response.

The argument rests on the further ethical premise that a patient's standing as a person entitled to respectful clinical consideration is not earned by compliance with previous recommendations. Present suffering and vulnerability provide reasons for care, while the professional relationship gives clinicians responsibilities to assess and respond to them. Accounts of relational autonomy, person-centred care, and care ethics provide resources for this position (Entwistle et al., 2010; Entwistle and Watt, 2013; Morberg Jämterud, 2022; Vanlaere and Gastmans, 2011). These responsibilities do not depend solely on establishing that the patient's earlier decision was justified. Even where criticism of past conduct has

independent grounds, it does not automatically determine what care is appropriate now. Three requirements follow for situations in which an unrealised alternative remains visible.

The first is to preserve the moral significance of present need. Pain, fear, dependence, and functional limitation remain reasons for a clinical response even when an earlier choice contributed to their development. A representation of a potentially healthier course does not make that suffering less real or establish a reduced entitlement to respectful assessment. The distinction is between evaluating what care can now achieve and judging whether the patient merits concern. Questions about making patients responsible require more than evidence of a causal contribution to illness (Hansson, 2018).

This requirement does not imply identical treatment regardless of clinical circumstances. Earlier decisions may have changed prognosis, treatment feasibility, expected benefit, or the options still available. Those changes can justify a different care plan. The justification should identify their relevance to the current decision, rather than treating the earlier refusal as a reason to penalise the patient. Equal moral standing is compatible with different clinical responses to different needs and prospects. It does not guarantee access to every requested intervention, remove resource constraints, or require clinicians to provide treatment they cannot professionally justify. It does require that a visible history of refusal not substitute for an assessment of present need and appropriate care.

The second requirement is to reopen deliberation without making reconsideration depend on an admission of past failure. Respecting refusal does not permanently close a clinical question. Treatments, risks, circumstances, and patient priorities may change, and continuing care may properly include renewed recommendations (Dudzinski and Shannon, 2006; Entwistle et al., 2010). An earlier simulation may assist that discussion if its date, assumptions, and current applicability are understood. The relevant question is what the patient should consider now.

Patients may reconsider, continue to refuse, or express regret. None of these responses makes the adverse outcome sufficient proof that the earlier choice was irrational or morally wrong. Regret can reveal a change in priorities or an experience of consequences that the patient now evaluates differently. Clinicians may explore that change without requiring the patient to repudiate the earlier decision. Independently justified reassessment remains possible; the restriction concerns using the outcome alone as a verdict. Renewed deliberation should help establish a workable course of care, rather than function as a demand for acknowledgement that the original recommendation should have prevailed.

The third requirement is to preserve relevant deliberative context across time. Work on patient-centred electronic records and the fidelity of goals-of-care documentation raises questions about how well clinical records retain the substance of conversations (Butler et al., 2020; Ma et al., 2024). Applied to the present case, this suggests that a simulated trajectory may remain readily accessible while the patient's reasons for declining the associated option become difficult to recover. This possible asymmetry is a conditional concern, not a claim that digital-twin systems routinely produce such an imbalance.

Where simulations are retained for later use, their interpretation should distinguish the evidence available at the original decision from information introduced subsequently. Relevant documentation may also include the burdens discussed, the patient's stated priorities, and unresolved uncertainty. Such context helps prevent a complex decision from becoming merely a recorded refusal beside a projected benefit. It also makes clear that an estimate represented what might occur under specified conditions, rather than what the patient was guaranteed to obtain.

Preserving context must remain proportionate and attentive to privacy and the patient's wishes about recording personal reasons. Patients should not have to disclose intimate commitments or provide an acceptable narrative to secure respectful care. Nor should an earlier account of their priorities become a permanent label that prevents those priorities from changing. Where reasons were not recorded or are no longer available, that absence is not evidence that the refusal lacked reasons. The purpose of documentation is to support fair interpretation and continuing deliberation, not to create a retrospective test of deservingness.

These requirements are particularly visible in relationships that continue through refusal, deterioration, and reconsideration. Nursing offers an important example of such continuity (Dudzinski and Shannon, 2006; Vanlaere and Gastmans, 2011), but the argument applies across the professionals and encounters involved in continuing care. A persistent simulation may connect those encounters informationally; sustaining the relevant ethical distinctions requires attention to how it is interpreted and used.

The broader significance concerns the authority of technologically represented alternatives after a decision has passed. Greater durability can make a comparison easier to retrieve without making it a sufficient account of the choice or a justified standard of moral evaluation. Continuing care therefore requires both clinical attention to the past and responsiveness to the person's present condition and agency. An unrealised future may help explain what happened and inform what should happen next. Its visibility alone cannot determine what the patient deserves from care.

8. Conclusion

This article has examined the normative status of a simulated patient future after the decision it was generated to inform. In the applications considered here, an alternative presented during deliberation remains available after refusal and may later become a comparator when an adverse outcome occurs. Its evidential content need not improve for its role in judgment to change. The contribution is to trace this transition and identify the justification required when a representation of what might have happened is used to establish what the patient ought to have chosen.

Three connected limits govern that inference. Predictive credibility does not by itself determine what the patient ought, all things considered, to choose. Medical preferability, even where the associated option is realistically available, does not by itself establish personal obligation. The realization of a medically worse future cannot, by itself, retrospectively supply an obligation that was not established prospectively. This does not exclude reassessment on independent grounds. It means that the

occurrence of harm and the continued visibility of an alternative cannot substitute for the normative justification that reassessment requires.

Continuing care also rests on the ethical commitment to recognising the patient's standing as a person entitled to respectful clinical consideration, regardless of compliance with earlier recommendations. Past choices may affect prognosis, treatment feasibility, and expected benefit, and may therefore justify changes in care. Their clinical relevance does not alone justify punitive treatment or diminished moral standing. Care should remain responsive to present need, permit renewed deliberation without demanding an admission of past failure, and preserve relevant context without making disclosure of personal reasons a condition of respect.

The argument applies to systems that make patient-specific alternatives available across a decision and its aftermath; it neither attributes these capabilities to all digital twins nor assumes that they inevitably generate blame. Its broader implication concerns how technological representations inform human agency over time. More accessible and credible simulations can support decisions while leaving their evaluative premises open to scrutiny. An unrealised future may remain useful for understanding the present and considering what comes next. Its persistence alone does not make it an authoritative verdict on the choice already made or the care now deserved.

Declarations

Data Availability

Data sharing is not applicable to this research as no data were generated or analysed.

Ethical Approval

This article does not contain any studies with human participants performed by any of the authors.

Informed Consent

This article does not contain any studies with human participants performed by any of the authors.

References

1. Baron J, Hershey JC (1988) Outcome bias in decision evaluation. J Pers Soc Psychol 54:569–579
2. Björnsson B, Borrebaeck C, Elander N, Gasslander T, Gawel DR, Gustafsson M, Jörnsten R, Lee EJ, Li X, Lilja S, Martínez-Enguita D, Matussek A, Sandström P, Schäfer S, Stenmarker M, Sun XF, Sysoev O, Zhang H, Benson M (2019) Digital twins to personalize medicine. Genome Med 12:4
3. Braun M (2021) Represent me: please! Towards an ethics of digital twins in medicine. J Med Ethics
4. Bruynseels K, de Sio S, van den Hoven F, J (2018) Digital Twins in Health Care: Ethical Implications of an Emerging Engineering Paradigm. Front Genet 9:31

5. Burr CD, Qian S, Winter P, Chico T, Rangel Smith C, Wagg D, Niederer SA (2026) Realising the digital twin: a thematic review and analysis of the ethical, legal, and social issues for digital twins in healthcare. AI Soc 41:5243–5267
6. Butler JM, Gibson B, Lewis L, Reiber G, Kramer H, Rupper R, Herout J, Long B, Massaro D, Nebeker J (2020) Patient-centered care and the electronic health record: exploring functionality and gaps. JAMIA Open 3:360–368
7. Buyx AM (2008) Personal responsibility for health as a rationing criterion: why we don't like it and why maybe we should. J Med Ethics 34:871–874
8. Cappelen AW, Norheim OF (2005) Responsibility in health care: a liberal egalitarian approach. J Med Ethics 31:476–480
9. Drummond D, Gonsard A (2024) Definitions and Characteristics of Patient Digital Twins Being Developed for Clinical Use: Scoping Review. J Med Internet Res 26:e58504
10. Dudzinski DM, Shannon SE (2006) Competent patients' refusal of nursing care. Nurs Ethics 13:608–621
11. Entwistle VA, Carter SM, Cribb A, McCaffery K (2010) Supporting patient autonomy: the importance of clinician-patient relationships. J Gen Intern Med 25:741–745
12. Entwistle VA, Watt IS (2013) Treating patients as persons: a capabilities approach to support delivery of person-centered care. Am J Bioeth 13:29–39
13. Epstein RM, Alper BS, Quill TE (2004) Communicating evidence for participatory decision making. JAMA 291:2359–2366
14. Epstein RM, Street RL Jr. (2011) The values and value of patient-centered care. Ann Fam Med 9:100–103
15. Ferlito B, De Proost M, Segers S (2024) Navigating the Landscape of Digital Twins in Medicine: A Relational Bioethical Inquiry. Asian Bioeth Rev 16:471–481
16. Feuerriegel S, Frauen D, Melnychuk V, Schweisthal J, Hess K, Curth A, Bauer S, Kilbertus N, Kohane IS, van der Schaar M (2024) Causal machine learning for predicting treatment outcomes. Nat Med 30:958–968
17. Friesen P (2018) Personal responsibility within health policy: unethical and ineffective. J Med Ethics 44:53–58
18. Hansson SO (2018) The Ethics of Making Patients Responsible. Camb Q Healthc Ethics 27:87–92
19. Hausman DM (2015) Valuing Health: Well-Being, Freedom, and Suffering. Oxford University Press
20. Hoffmann TC, Montori VM, Del Mar C (2014) The connection between evidence-based medicine and shared decision making. JAMA 312:1295–1296
21. Huang P-h, Kim K-h, Schermer M (2022) Ethical Issues of Digital Twins for Personalized Health Care Service: Preliminary Mapping Study. J Med Internet Res 24:e33081
22. Iqbal JD, Krauthammer M, Biller-Andorno N (2022) The Use and Ethics of Digital Twins in Medicine. J Law Med Ethics 50:583–596

23. Iqbal JD, Krauthammer M, Witt CM, Biller-Andorno N, Christen M (2025) A consensus statement on the use of digital twins in medicine. npj Digit Med 8:484
24. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, Achenie L, Liu H, Jackson P, Xiao Y, Syeda-Mahmood T, Tuli R, Deng J (2024) Digital twins for health: a scoping review. NPJ Digit Med 7:77
25. Kelley M (2005) Limits on Patient Responsibility. J Med Philos 30:189–206
26. Kleinman I (1991) The right to refuse treatment: ethical considerations for the competent patient. CMAJ 144:1219–1222
27. Landmark AMD, Gulbrandsen P, Svennevig J (2015) Whose decision? Negotiating epistemic and deontic rights in medical treatment decisions. J Pragmat 78:54–69
28. Ma JE, Schlichte L, Haverfield M, Gambino J, Lange A, Blanchard K, Morgan B, Bekelman DB (2024) Do goals of care documentation reflect the conversation? Evaluating conversation-documentation accuracy. J Am Geriatr Soc 72:2500–2507
29. Mair FS, May CR (2014) Thinking about the burden of treatment. BMJ 349:g6680
30. May C, Montori VM, Mair FS (2009) We need minimally disruptive medicine. BMJ 339:b2803
31. Montori VM, Brito JP, Murad MH (2013) The optimal practice of evidence-based medicine: incorporating patient preferences in practice guidelines. JAMA 310:2503–2504
32. Morberg Jämterud S (2022) Acknowledging vulnerability in ethics of palliative care - A feminist ethics approach. Nurs Ethics 29:952–961
33. Nordenfelt LY (2013) On the nature of health: An action-theoretic approach. Springer Science & Business Media
34. Olthuis G, Leget C, Grypdonck M (2014) Why shared decision making is not good enough: lessons from patients. J Med Ethics 40:493–495
35. Parker WF, Chin MH (2020) Epistemic Authority and Trust in Shared Decision Making About Organ Transplantation. AMA J Ethics 22:E408–415
36. Popa EO, van Hilten M, Oosterkamp E, Bogaardt MJ (2021) The use of digital twins in healthcare: socio-ethical benefits and socio-ethical risks. Life Sci Soc Policy 17:6
37. Resnik DB (2007) Responsibility for health: personal, social, and environmental. J Med Ethics 33:444–445
38. Roese NJ (1997) Counterfactual thinking. Psychol Bull 121:133–148
39. Roese NJ, Olson JM (1996) Counterfactuals, Causal Attributions, and the Hindsight Bias: A Conceptual Integration. J Exp Soc Psychol 32:197–227
40. Schwan B (2021) Responsibility amid the social determinants of health. Bioethics 35:6–14
41. Sherwin S, Winsby M (2011) A relational perspective on autonomy for older adults residing in nursing homes. Health Expect 14:182–190
42. Sprangers MA, Schwartz CE (1999) Integrating response shift into health-related quality of life research: a theoretical model. Soc Sci Med 48:1507–1515

43. Sun T, He X, Li Z (2023) Digital twin in healthcare: Recent updates and challenges. Digit Health 9:20552076221149651
44. Tudor BH, Shargo R, Gray GM, Fierstein JL, Kuo FH, Burton R, Johnson JT, Scully BB, Asante-Korang A, Rehman MA, Ahumada LM (2025) A scoping review of human digital twins in healthcare applications and usage patterns. NPJ Digit Med 8:587
45. Vallée A (2026) From prediction to intervention: causal digital twins for personalized clinical decision support. J Translational Med 24:441
46. Vanlaere L, Gastmans C (2011) A personalist approach to care ethics. Nurs Ethics 18:161–173
47. Wikler D (2002) Personal and social responsibility for health. Ethics Int Aff 16:47–55
48. Yang X, Liang Y, Liu Y, Liu D, Yang L, Zhang B, Xu C, Wu Y, He J 2026a. When machines explain medicine: Nursing ethics and clinical communication. Nurs Ethics 0, 9697330261449460
49. Yang X, Xiao Y, Liu D, Sun Y, Zhou X, Yao J, Zhang B, Yang L, Guo W, Tang W, Tang X, Wu Y, Hu X, Wu F, Wang D, Xu C (2026b) Ethical tensions in AI-mediated clinical communication: a mixed-methods study of nursing students' perspectives. BMC Medical Ethics
50. Yang X, Xu C (2026) Ethical Governance of Large Language Models in Health Care: Trust, Responsibility, and Equity in Routine Use. J Med Internet Res 28:e93470

Figures

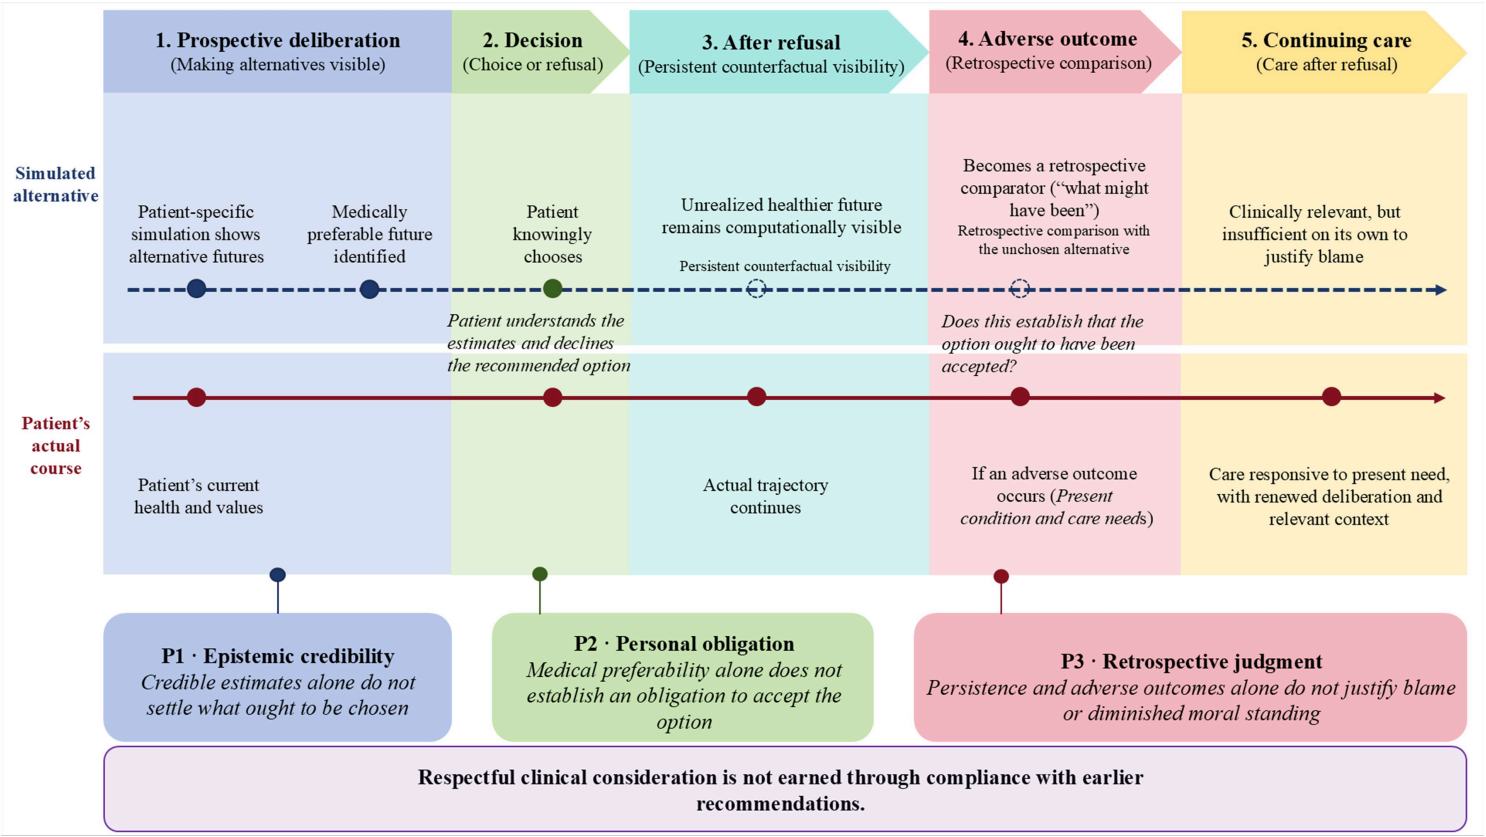


Figure 1

The changing role of a simulated patient future from deliberation to care after refusal.

The figure traces a patient-specific simulated alternative alongside the patient's actual course. Before a decision, the simulation provides evidence about possible outcomes. Epistemic credibility alone does not determine what the patient ought to choose (Proposition 1), and medical preferability alone does not establish an obligation to accept the associated option (Proposition 2). After informed refusal, the simulated alternative may remain available. If an adverse outcome occurs, it may become a retrospective comparator without thereby verifying what would have happened under the unchosen option. Its persistence and retrospective salience do not, by themselves, justify blame or diminish the moral standing of subsequent clinical need (Proposition 3). Continuing care may draw on clinically relevant information about earlier choices while remaining responsive to present need. The two tracks distinguish the simulated representation from the patient's lived course; they do not depict measured health trajectories or imply that refusal necessarily causes deterioration.