BELIEF-AUGMENTED JOINT MODEL: TIME-VARYING CAUSAL INFERENCE UNDER LATENT TRAJECTORIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Identifying and estimating time-varying treatment effects from observational healthcare data is challenging when clinicians act on beliefs about patient trajectories that are only partially observed. Calibration for observed covariates alone may leave residual confounding induced by these beliefs. We propose the Belief Augmented Joint Model (BAM), which reconstructs natural recovery trajectories by incorporating a belief-dependent process of clinical decisions. BAM represents clinicians' beliefs through a closed-form posterior over latent trajectory coefficients, derived from real-world observations under a linear-Gaussian trajectory model. We prove that under belief-driven confounding—the latent trajectory influences the decision only through this belief—the time-varying treatment effect is identified without estimating or parameterizing the underlying decision rule. In synthetic benchmarks, BAM outperforms eight standard models; the recovered decision channels successfully transfer to a rehabilitation cohort and MIMIC-III. These results establish a mathematically grounded framework for time-varying causal inference, demonstrating that explicitly accounting for clinical beliefs is essential for reliable decision support in long-cycle healthcare.

1 INTRODUCTION

Identifying and estimating time-varying treatment effects is crucial for personalized treatment decisions in inpatient rehabilitation. Such effects are increasingly inferred from electronic medical records (Feuerriegel et al., 2024). This requires methods that can estimate treatment effects over time from sparse, noisy observational data.

Existing approaches typically infer treatment effects from length of stay (LOS) (Ottenbacher et al., 2004). However, this induces two critical failures: first, it conflates treatment effects with natural recovery; second, it ignores individual recovery trajectories, failing to explain why some patients recover well after a short stay while others recover poorly after a long stay (Figure 1) (Hodel et al., 2021). The root cause is that LOS is driven by clinicians' decisions, which in turn depend on their subjective beliefs regarding each patient's latent trajectory (Kimmel et al., 2017). Two decades ago, Powell et al. (2002) recognized that unmeasured aspects of discharge decisions confound outcome evaluation; this fundamental barrier remains entirely outside the reach of standard covariate adjustment today.

To decompose observed recovery into natural recovery and treatment contribution, methods must infer the clinician's belief from sparse assessments and, in doing so, account for the uncertainty of that inference. Existing approaches to confounding adjustment—classical and neural alike—rely on observed covariates, external proxies, or learned representations of patient history (Robins et al., 2000; Miao et al., 2018; Bica et al., 2020a). As such, they implicitly assume that the clinician's belief about the latent recovery trajectory is observed, exogenous, or adequately captured by recorded covariates. However, in inpatient rehabilitation, this belief is unrecorded and directly guides daily treatment decisions, while functional outcomes are measured only sparsely. Crucially, because the belief must itself be inferred from data, treating it as a known input omits uncertainty that propagates to treatment effect estimates, leading to overconfident or biased inference. Recent neural methods can handle irregular observations and time-varying treatment decisions by weighting or balancing recorded patient histories (Seedat et al., 2022; Melnychuk et al., 2022; Frauen et al., 2025). However, when treatment decisions also depend on a clinician's belief about an unobserved

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

recovery trajectory, adjustment for recorded history alone does not establish identification. Other models infer latent recovery trajectories and their uncertainty, but do not connect those trajectories to treatment decisions to estimate causal effects (Hodel et al., 2021). To our knowledge, no existing framework identifies time-varying treatment effects while controlling for belief-driven confounding and propagating belief-inference uncertainty. (Table 1)

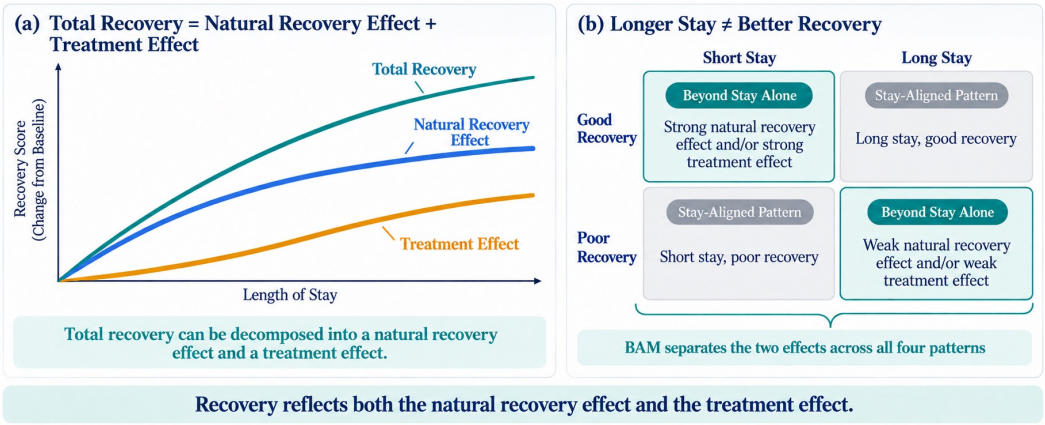


Figure 1: (a) Total recovery equals the natural recovery effect plus the treatment effect. (b) A longer stay does not necessarily imply better recovery.

We introduce the Belief-Augmented Joint Model (BAM) to separate natural recovery from the contribution of treatment over time. BAM represents each patient's untreated recovery by a latent trajectory and models the day-specific treatment effect with a flexible function $\beta(t)$. Under a linear-Gaussian trajectory model, sparse and noisy functional assessments yield a closed-form posterior over the latent trajectory coefficients. Its posterior mean serves as an explicit representation of the clinician's belief about the patient's recovery. BAM then links this belief to treatment decisions and estimates the trajectory, treatment effect, and decision components jointly. The outcome model retains the full daily treatment history, while the decision component uses treatment choices at observed assessment times.

Table 1: Capabilities for longitudinal treatment-effect estimation (✓: addressed; ✗: not addressed).

	Irregular observations	Time-varying confounding	Unobserved confounding	Confounder uncertainty in effect inference	Representative work
Discrete time (balancing)	✗	✓	✗	✗	CT (Melnychuk et al., 2022)
Discrete time (weighting)	✗	✓	✗	✗	Meta-learners (Frauen et al., 2025)
Continuous time (Bayesian CDE)	✓	✗	✗	✗	BNCDE (Hess et al., 2024)
Continuous time (weighting)	✓	✓	✗	✗	SCIP-Net (Hess & Feuerriegel, 2025)
Proxy-based	✗	✗	✓	✗	Wu et al. (2024)
Bayesian latent-class causal model	✗	✓	✓	✓	Liu et al. (2025)
BAM (ours)	✓	✓	✓	✓	—

Our main contributions are threefold:

1. **Conceptualizing Belief-Driven Confounding:** We formally introduce belief-driven confounding as a structural paradigm for longitudinal causal inference. By treating the clinician's belief as the mechanistic bridge between latent recovery and treatment choices, BAM successfully isolates natural recovery from time-varying treatment effects.
2. **Non-Parametric Identification:** We establish the non-parametric identification of time-varying treatment effects under belief sufficiency, requiring neither external proxies nor explicit parameterization of the clinician's decision rule. This transforms an unrecorded clinical driver into a precise, model-based identification condition.

3. **Operationalization and Uncertainty Propagation:** We operationalize this identification strategy for sparse rehabilitation cohorts. BAM computes clinician beliefs in closed form from noisy assessments and performs joint inference, ensuring that the uncertainty incurred during belief estimation is seamlessly propagated through to the final posterior uncertainty.

2 Related Work

Time-varying treatment effects. Marginal structural models adjust for time-varying confounding through weighting (Robins et al., 2000), while structural nested models express effects as history-specific blips (Vansteelandt & Joffe, 2014). Neural sequence models extend effect estimation to complex treatment and outcome histories (Bica et al., 2020a; Melnychuk et al., 2022); continuous-time models additionally accommodate irregular observations (Seedat et al., 2022; Hess & Feuerriegel, 2025). In contrast, BAM estimates day-specific treatment effects in sparse rehabilitation data—where binary treatment status is recorded daily but functional outcomes are assessed intermittently. It explicitly modeling the unobserved clinical belief that drives decisions, thereby isolating the true treatment effect from natural recovery.

Joint models of outcomes and treatment decisions. Shared-parameter models link longitudinal measurements to related processes through patient-specific latent trajectories (Wulfsohn & Tsiatis, 1997; Henderson et al., 2000). Closer to our setting, Sitlani et al. (2012) jointly analyze outcomes and treatment under noncompliance, Shardell & Ferrucci (2018) combine an outcome–exposure joint model with g-computation, and Liu et al. (2025) jointly model treatment and outcomes through visit-specific latent health classes. Unlike prior joint frameworks, BAM defines the decision-relevant shared state as an updated posterior belief regarding untreated recovery, computed from assessments available prior to each decision. Specifically, the latent trajectory governs natural recovery, the belief informs the treatment model, and the full treatment history enters the outcome model, with joint estimation unifying all three components.

Beliefs and latent confounding. Belief states represent decisions made under partial observation (Kaelbling et al., 1998). For causal effect estimation with hidden confounding, the time-series deconfounder learns latent substitutes from multiple treatments (Bica et al., 2020b), whereas proximal methods use measured proxies and bridge functions (Cui et al., 2024). In contrast, BAM derives the belief from sparse functional assessments without requiring separately measured proxies. Under belief sufficiency, it identifies the time-varying treatment effect without knowing the clinician's decision rule. Joint inference then propagates uncertainty in the inferred belief to the estimation of the effect curve.

3 Methods and Theory

3.1 Formalization and Mechanism

Patients $i = 1, \ldots, n$ are followed for T days ($T = 60$ in the rehabilitation cohort). Treatment $D_i(d) \in \{0, 1\}$ is recorded daily, with 0 indicating no treatment and 1 indicating treatment, whereas the standardized functional outcome Y_{ij} is observed only on irregular assessment days d_{ij}, with $t_{ij} = d_{ij}/T$. Let a_i be admission age and let $\bar{H}_i(t)$ contain outcomes and treatments recorded strictly before t. The latent vector $b_i = (b_{1i}, b_{2i})^\top \sim N(0, \Sigma)$ describes patient-specific deviations in untreated level and recovery rate. Within an assessment day we order $\bar{H}_i(t_{ij}) \to D_{ij} \to Y_{ij}$.

For a treatment history $\bar{d}$, let $Y_i^{\bar{d}}(t)$ be the potential outcome. The primary estimand is the effect of adding treatment on day d,

$$\mathbb{E}\Big[Y_i^{\bar{d}+\delta_d}(t) - Y_i^{\bar{d}}(t)\Big] = \beta(d/T) = B(d/T)^\top \beta, \qquad t \geq d/T, \tag{1}$$

with zero effect before d. Here $B(\cdot) \in \mathbb{R}^K$ is a prespecified origin-rebased spline basis. The cumulative daily-treatment contrast is $\Gamma(t) = \sum_{d \leq Tt} \beta(d/T)$. In Section 4, we use the day-42 single-day treatment effect $\beta(42/T)$ as a reporting point for comparison. The complete notation used by BAM is provided in Appendix A.

Figure 2 provides an overview of BAM. The model represents a latent untreated recovery trajectory inferred from the observed outcome history, summarizes this information as a posterior belief

that informs treatment decisions, and estimates treatment effects through their influence on future outcomes. Across time, past treatments contribute to subsequent outcomes, while newly observed outcomes update the belief used for later decisions.

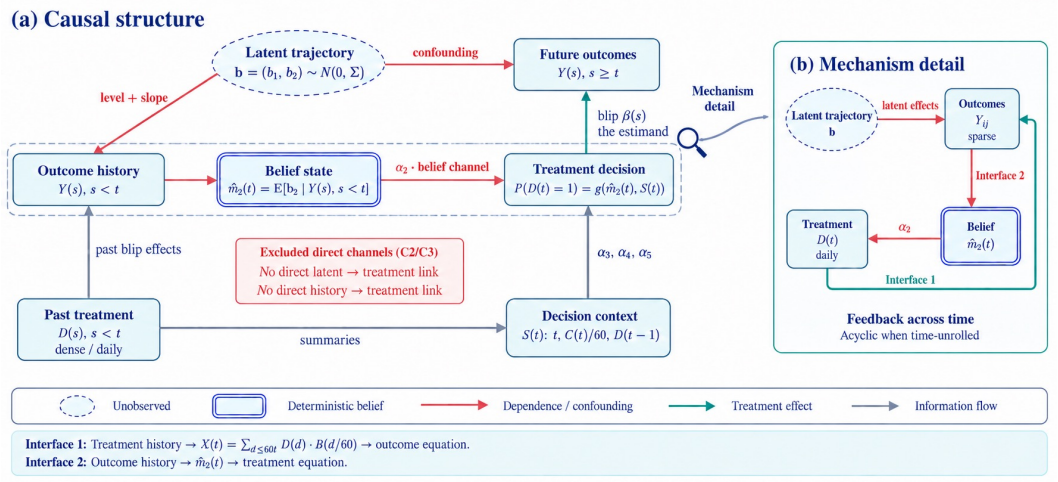


Figure 2: BAM causal structure and belief-mediated feedback.

3.2 BAM

The outcome model retains the complete daily treatment history:

$$Y_{ij} = \mu(t_{ij}) + \gamma a_i + X_i(t_{ij})^\top \beta + b_{1i} + b_{2i} t_{ij} + \varepsilon_{ij}, \qquad \varepsilon_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2). \tag{2}$$

Here, $\mu(t)$ is the population untreated recovery trajectory, γ is the admission-age coefficient, and ε_{ij} is the residual outcome error with variance σ^2. The prespecified spline basis $B(d/T)$ and coefficient vector β define the single-day treatment-effect curve. The accumulated treatment design $X_i(t) = \sum_{d \le Tt} D_i(d) B(d/T)$ summarizes the complete daily treatment history up to time t.

Before time t, we subtract the fixed and treatment components from previously observed outcomes to form the historical residual vector $r_i(t)$. Let $G_i(t)$ be the corresponding random-effects design matrix, with rows $(1, t_{ij})$ for assessments before t. Gaussian conditioning gives the exact history-filtered belief

$$\widehat{m}_i(t;\eta) = \mathbb{E}_\eta[b_i \mid \bar{\mathcal{H}}_i(t)] = A_i(t;\Sigma,\sigma^2) r_i(t), \qquad A_i = \left\{\Sigma^{-1} + \sigma^{-2} G_i^\top G_i\right\}^{-1} \sigma^{-2} G_i^\top, \tag{3}$$

where $A_i(t;\Sigma,\sigma^2)$ maps the residual outcome history to the posterior mean of b_i, and $\eta = (\mu, \gamma, \beta, \sigma^2, \Sigma)$ collects the outcome-side parameters. The vector $\widehat{m}_i(t;\eta)$ is the posterior mean of the patient-specific untreated trajectory, and $\widehat{m}_{2i}(t;\eta)$ is its slope component. It is a model-based representation of the decision maker's belief, not a direct measurement of an internal clinical state. For outcomes on which higher scores are better, a negative value of $\widehat{m}_{2i}(t)$ indicates slower untreated recovery than the population trajectory, whereas a positive value indicates faster recovery. Under the Gaussian model, this belief is available in closed form and is differentiable with respect to the outcome-side parameters.

At assessment times, the treatment decision model is

$$\operatorname{logit} \Pr\big(D_{ij} = 1 \mid \bar{\mathcal{H}}_i(t_{ij})\big) = \alpha_0 + \alpha_2 \widehat{m}_{2i}(t_{ij}) + \alpha_3 t_{ij} + \alpha_4 C_i(t_{ij})/T + \alpha_5 D_i(d_{ij} - 1), \tag{4}$$

where $C_i(t) = \sum_{d<Tt} D_i(d)$ is the number of treatment days before time t. The coefficients α_0, α_2, α_3, α_4, and α_5 represent the intercept and the associations of treatment with the posterior slope belief, calendar time, cumulative prior treatment, and previous-day treatment $D_i(d_{ij} - 1)$, respectively. Thus, outcome-side parameters determine the posterior belief entering the treatment model, while the complete daily treatment history enters the outcome model through $X_i(t)$.

3.3 Assumptions and Identification

A1 (Causal Consistency). Consistency, no interference, and no anticipation hold. Thus, the observed outcome under the realized treatment history equals the corresponding potential outcome,

one patient's treatment does not affect another patient's outcome, and treatment does not affect outcomes before it is administered.

A2 (Daily History-Conditional Ignorability). For every calendar day d,

$$D_i(d) \perp\!\!\!\perp \left(b_i,\, Y_i^{\bar{d}}(s) : s \geq d/T\right) \mid \bar{\mathcal{H}}_i(d/T). \tag{5}$$

Conditional on the recorded history before day d, treatment is independent of the latent untreated trajectory and future potential outcomes. This assumption also applies on days without a directly observed functional-outcome score, although BAM does not parameterize the treatment probabilities on those days.

A3 (Belief-State Sufficiency for Treatment Decisions). At an assessment time t_{ij}, there exists a measurable function g such that

$$\Pr\big(D_{ij} = 1 \mid \bar{\mathcal{H}}_i(t_{ij})\big) = g\{\widehat{m}_{2i}(t_{ij}), t_{ij}, C_i(t_{ij})/T, D_i(d_{ij}-1)\}\,. \tag{6}$$

Thus, the recorded history affects the modeled treatment decision through the posterior slope belief and the stated decision summaries. For identification, g may be any measurable function of these quantities. The logistic-linear form in Equation 4 is imposed only to obtain a finite-dimensional treatment model for joint estimation.

A4 (Overlap and Rank). On the target support, treatment probabilities are bounded away from zero and one. Both treatment states occur with positive probability on supported assessment days, and the corresponding spline vectors $B(t)$ and fixed-effect design have full rank.

A5 (Regularity and Assessment Mechanism). The model parameters lie in the interior of the parameter space, σ^2 is bounded away from zero, and Σ is uniformly positive definite. Assessment times are conditionally non-informative for the main model given the recorded history. A clear interpretation of the assumptions is provided in Appendix B.

Theorem 1 (Identification under belief-mediated treatment decisions)**.** *Under Equation 2, A1–A4, and the assessment-time condition in A5, any two parameterizations that induce the same observed-data law have the same treatment-effect coefficient vector β. Consequently, the single-day treatment-effect curve $\beta(d/T) = B(d/T)^\top \beta$ and the cumulative treatment contrast $\Gamma(t)$ are identified without specifying a parametric form for the decision function g.*

To see this, fix the recorded history immediately before a supported assessment. Under A2, conditioning on this history removes residual treatment selection through b_i and future potential outcomes. The treated–untreated conditional mean contrast is therefore $B(t_{ij})^\top \beta$, with the same conditional distribution of the latent untreated component in the two treatment groups. A3 specifies the decision-relevant state through which the recorded history enters the modeled treatment decision, but the identification argument does not require the functional form of g to be known. Under A4, variation in treatment over supported assessment days and the full-rank spline design identify β. Treatment effects on days without a directly observed outcome score are determined by the prespecified spline curve rather than by separate day-specific contrasts. Equation 4 is subsequently used to parameterize g for joint estimation, not to establish identification.

3.4 Joint Estimation and Inference

Let $\mathcal{J}_i$ index the assessment days of patient i. Define the decision covariate vector $z_{ij}(\eta) = \{1, \widehat{m}_{2i}(t_{ij};\eta), t_{ij}, C_i(t_{ij})/T, D_i(d_{ij}-1)\}^\top$ and let $q_{ij}(\eta,\alpha) = \mathrm{expit}\{\alpha^\top z_{ij}(\eta)\}$ denote the treatment probability in Equation 4. The parameter-relevant sequential likelihood is $L_i(\eta,\alpha) = L_{Y,i}(\eta) L_{D,i}(\eta,\alpha)$, with the outcome and decision components

$$L_{Y,i}(\eta) = \int \phi_2(b;0,\Sigma) \prod_{j\in\mathcal{J}_i} \phi\big(Y_{ij};\nu_{ij}(b;\eta),\sigma^2\big)\, db, \quad L_{D,i}(\eta,\alpha) = \prod_{j\in\mathcal{J}_i} q_{ij}^{D_{ij}}(1-q_{ij})^{1-D_{ij}}. \tag{7}$$

Under A2 and A5, the full observed-sequence likelihood equals Equation 7 multiplied by the decision laws on days without directly observed scores. These laws are variation-independent nuisance factors: conditional on the recorded history, they do not provide additional information about b_i. Treatments on those days nevertheless remain in $X_i(t)$, $C_i(t)$, and the lagged-treatment term. We

estimate $\theta = (\eta^\top, \alpha^\top)^\top$ by maximizing $\sum_i \log L_i(\eta, \alpha)$. Because q_{ij} depends on the estimated belief $\widehat{m}_{2i}(t_{ij}; \eta)$, the outcome and treatment-decision components are coupled through η.

Let I denote the limiting per-patient Fisher information and partition it as $I = \left(\begin{smallmatrix} I_{\eta\eta} & I_{\eta\alpha} \\ I_{\alpha\eta} & I_{\alpha\alpha} \end{smallmatrix}\right)$. Define the Schur complements $S_\alpha = I_{\alpha\alpha} - I_{\alpha\eta} I_{\eta\eta}^{-1} I_{\eta\alpha}$ and $S_\eta = I_{\eta\eta} - I_{\eta\alpha} I_{\alpha\alpha}^{-1} I_{\alpha\eta}$.

Algorithm. BAM is fitted via joint maximum likelihood. At each candidate parameter value, Gaussian conditioning updates the history-specific belief $\widehat{m}_{2i}(t)$, after which the outcome and treatment-decision contributions are evaluated joinly. We initialize the two components using an outcome-only mixed model (Laird & Ware, 1982) and a logistic treatment model, and then optimize the joint objective from multiple starting values using a BFGS–Nelder–Mead–BFGS multi-stage sequence (Nelder & Mead, 1965; Nocedal & Wright, 2006). The solution yielding the largest unpenalized joint likelihood is retained, and uncertainty is computed from the full observed-information matrix. Further implementation and convergence details are provided in Appendix D.

Theorem 2 (Joint inference with an estimated belief state)**.** *Assume independent patient clusters, A2 and A5, correct specification of the factors in Equation 7, variation-independent nuisance decision laws on days without directly observed scores, identifiability and interiority of θ_0, positive-definite limiting information, and standard differentiability and moment conditions. Then the joint maximum-likelihood estimator satisfies*

$$\sqrt{n}(\widehat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, I^{-1}), \qquad \operatorname{avar}\left\{\sqrt{n}(\widehat{\alpha} - \alpha_0)\right\} = S_\alpha^{-1}, \qquad \operatorname{avar}\left\{\sqrt{n}(\widehat{\eta} - \eta_0)\right\} = S_\eta^{-1}. \tag{8}$$

Writing $p_{ij} = q_{ij}(\eta_0, \alpha_0)$ and $z_{ij} = z_{ij}(\eta_0)$, the cross-information induced by the belief channel is

$$I_{\alpha\eta} = \alpha_2 \, \mathbb{E}_i \left[\sum_{j \in \mathcal{J}_i} p_{ij}(1 - p_{ij}) z_{ij} \frac{\partial \widehat{m}_{2i}(t_{ij}; \eta_0)}{\partial \eta^\top} \right], \tag{9}$$

which vanishes when the belief channel is inactive, that is, when $\alpha_2 = 0$.

Theorem 2 establishes asymptotically valid joint inference for BAM when the belief state is estimated from the observed outcome history. The Schur complements characterize how uncertainty is coupled across the outcome and treatment-decision components, while Equation 9 isolates the coupling induced by the belief channel. This result provides a principled basis for constructing standard errors and confidence intervals for the treatment-effect curve from the full joint covariance matrix. In the experiments, we use this covariance matrix for uncertainty quantification and examine the practical contribution of joint uncertainty propagation.

3.5 A SHARP BOUNDARY

Proposition 1 (Direct latent selection)**.** *If Assumption A2 is relaxed so that treatment may depend directly on the unobserved slope b_{2i} after conditioning on the recorded history, then β is not identified in general. In a Gaussian threshold-selection construction, the observed treated–untreated contrast equals the causal blip plus a Mills-ratio selection term, and the two components cannot be separated from the observed law.*

Proposition 1 gives the identification result a precise scope and motivates the direct-latent-channel stress test in the experimental section. The resulting theory has one continuous line: BAM defines an exact belief state linking outcome and decision components; Theorem 1 establishes causal validity under recorded-history adjustment; Theorem 2 gives joint inference for the generated state; and Proposition 1 characterizes the failure boundary. Complete proofs of Theorems 1 and 2 and Proposition 1 are given in Appendix C.

4 EXPERIMENTS AND RESULTS

4.1 EXPERIMENTAL DESIGN

Data sources and preprocessing. We evaluate BAM in three complementary settings. The primary dataset is a multicenter institutional rehabilitation cohort derived from routinely collected clinical records at several Grade A tertiary hospitals in China. It includes 419 patients followed for up

to 60 days, with 3,674 irregular assessments and daily therapy records. The standardized neurological outcome combines domains from the ISNCSCI examination and established functional-independence measures (Kirshblum et al., 2011; Keith et al., 1987; Catz et al., 1997). Treatment is recorded daily, and each assessment-day decision is modeled using the history available before that day's treatment. We also construct cohort-calibrated simulations using the empirical assessment schedules and fitted characteristics of the rehabilitation data. Because the treatment-effect curve is known in simulation, effect recovery can be evaluated directly. BAM estimates the single-day effect $\beta(d/60)$ over days $d = 0, \ldots, 60$; we use $\beta(42/60)$ as a common reporting point for comparison. Finally, we use the openly distributed 100-patient MIMIC-III Clinical Database Demo (Johnson et al., 2016; Pollard et al., 2019) for four medication–outcome analyses: two furosemide–creatinine extractions, vasopressin–creatinine, and insulin–glucose. These analyses assess whether BAM's treatment-decision structure transfers to intensive-care data. We summarize simulation performance by bias and root mean squared error (RMSE) and evaluate decision channels in the observed cohorts using likelihood-ratio tests. The full data construction and processing pipeline is described in Appendix E.

Experimental design. The experiments are organized around four complementary questions. First, the rehabilitation cohort evaluates whether BAM recovers clinically interpretable treatment-decision channels and a time-varying effect curve in the target setting. Second, cohort-calibrated simulations compare effect-estimation accuracy with standard causal and longitudinal baselines under a known data-generating process. Third, trajectory misspecification, direct latent selection, belief-feature ablation, and the outcome-only LMM comparison separately examine robustness, the identification boundary, and the contribution of joint belief–decision modeling. Finally, the neural variants and MIMIC-III analyses assess representational flexibility and the portability of the decision structure beyond the primary cohort.

4.2 Experimental Results

4.2.1 Clinical Application

Direct application to the rehabilitation cohort. BAM estimated $\widehat{\alpha}_0 = -2.252$ (SE 0.090), $\widehat{\alpha}_2 = -0.151$ (SE 0.075), $\widehat{\alpha}_3 = -1.092$ (SE 0.258), $\widehat{\alpha}_4 = 12.040$ (SE 1.098), and $\widehat{\alpha}_5 = 0.737$ (SE 0.168) in the rehabilitation cohort; the likelihood-ratio statistics for the belief, time, accumulated-exposure, and previous-day treatment channels were 4.17, 19.55, 113.94, and 20.93, respectively. These results indicate that patients with a slower inferred untreated recovery rate were more likely to receive treatment, treatment probability decreased over follow-up conditional on the other decision variables, and patients with greater accumulated exposure or treatment on the previous day were more likely to continue treatment. These patterns are consistent with clinical practice in which rehabilitation plans are adjusted according to patients' recovery progress and implemented through coordinated, repeated treatment sessions (Winstein et al., 2016; McGlinchey et al., 2019; Foley et al., 2012). The estimated day-42 single-day treatment effect was $\widehat{\beta}(42/60) = 0.088$ (SE 0.033), indicating that one treatment day at day 42 increased the standardized functional outcome by an average of 0.088 standard deviations. The daily treatment effect over the entire rehabilitation period is shown in Figure 3.

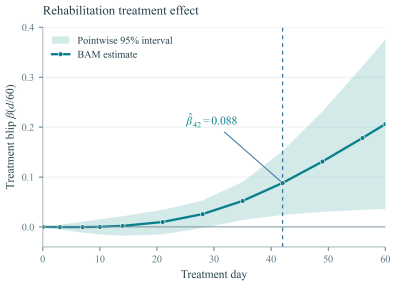


Figure 3: Trend of daily treatment effects over time.

4.2.2 Benchmark Comparison

Comparison with standard baselines. We compared BAM with inverse-probability weighting and its oracle-propensity version (IPTW and O-IPTW; Robins et al., 2000), random-forest and gradient-boosted g-computation (RF and GBT; Breiman, 2001; Friedman, 2001), a doubly robust learner (DR; Bang & Robins, 2005), causal forest (CF; Wager & Athey, 2018), and a gated recurrent unit sequence model (GRU; Cho et al., 2014). Across $B = 40$ cohort-calibrated replications, BAM achieved an RMSE of 0.036 at day 42, compared with 0.076 for the best-performing standard comparator, RF, and 0.079 for GRU. IPTW, O-IPTW, and DR produced RMSEs between 0.220

and 0.240, while GBT and CF achieved RMSEs of 0.154 and 0.203, respectively. BAM therefore reduced RMSE by more than one half relative to the strongest standard comparator. This result indicates that explicitly linking intermittent outcomes, the inferred untreated recovery trajectory, and daily treatment decisions improves treatment-effect estimation when outcomes are sparsely assessed but treatment changes over time. This finding is consistent with previous work showing that longitudinal causal estimation requires explicit adjustment for evolving treatment assignment and outcome history (Robins et al., 2000; Bica et al., 2020a; Melnychuk et al., 2022). The corresponding results are reported in the $B = 40$ columns of Table 2.

Table 2: Bias and RMSE for the single-day treatment effect at day 42, $\beta(42)$

	Main benchmark		Trajectory misspecification										Identification-boundary stress test					
Method	$B = 40$		Control		t_3 errors		$\kappa = 0.3$		$\kappa = 0.6$		$\kappa = 1.2$		$\delta = 0.15$		$\delta = 0.30$		$\delta = 0.60$	
	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
BAM	−0.005	0.036	+0.009	0.031	−0.004	0.033	+0.005	0.033	−0.001	0.032	+0.000	0.048	+0.031	0.039	+0.044	0.049	+0.062	0.065
LMM	–	–	+0.009	0.031	–	–	–	–	–	–	–	–	+0.031	0.039	+0.042	0.047	+0.060	0.063
IPTW	+0.024	0.220	+0.033	0.267	+0.047	0.160	+0.018	0.267	+0.012	0.278	+0.010	0.338	+0.073	0.251	+0.096	0.232	+0.191	0.324
O-IPTW	+0.030	0.229	+0.054	0.271	+0.054	0.153	+0.042	0.274	+0.034	0.288	+0.024	0.342	+0.058	0.247	+0.062	0.228	+0.133	0.300
RF	−0.063	0.076	−0.059	0.078	−0.062	0.070	−0.064	0.074	−0.075	0.079	−0.076	0.080	−0.037	0.103	−0.035	0.099	−0.017	0.113
GBT	+0.014	0.154	+0.019	0.138	−0.002	0.139	+0.000	0.144	−0.024	0.120	−0.031	0.131	+0.032	0.171	+0.032	0.132	+0.087	0.206
DR	+0.036	0.240	+0.039	0.268	+0.042	0.171	+0.018	0.266	+0.012	0.280	+0.013	0.326	+0.084	0.229	+0.114	0.224	+0.163	0.273
CF	+0.082	0.203	+0.103	0.215	+0.091	0.173	+0.094	0.213	+0.068	0.202	+0.048	0.235	+0.193	0.272	+0.274	0.321	+0.450	0.494
GRU	−0.077	0.079	−0.077	0.078	−0.076	0.077	−0.077	0.078	−0.076	0.077	−0.071	0.074	−0.072	0.074	−0.070	0.072	−0.062	0.067

4.2.3 Robustness and Ablation

Robustness to trajectory misspecification. We replaced Gaussian outcome errors with t_3 errors and generated quadratic patient trajectories of increasing strength κ, while fitting BAM with the original linear-Gaussian belief model. Across the heavy-tailed and quadratic scenarios, BAM's bias remained within ± 0.005, and its RMSE ranged from 0.032 to 0.048, including an RMSE of 0.048 under the strongest trajectory departure at $\kappa = 1.2$. BAM also achieved a lower RMSE than every standard comparator at each value of κ. These results show that BAM can recover the day-42 treatment effect even when patient outcomes contain heavy-tailed variation or nonlinear recovery patterns that are not explicitly represented by the fitted belief model. This finding is consistent with previous evidence that fixed-effect estimates from mixed-effects models can remain approximately unbiased under moderate violations of their distributional assumptions, although estimation uncertainty may increase (Schielzeth et al., 2020). The corresponding results are reported in the t_3 errors, $\kappa = 0.3$, $\kappa = 0.6$, and $\kappa = 1.2$ columns of Table 2.

Stress test at the identification boundary. We added $\delta b_{2i}/\tau_2$ directly to the treatment-decision index, allowing treatment to depend on the unobserved recovery slope even after conditioning on the recorded history. This setting violates A2 rather than merely misspecifying the functional form of the decision model (Tchetgen Tchetgen et al., 2020). As δ increased from 0 to 0.15, 0.30, and 0.60, BAM's bias increased from 0.009 to 0.031, 0.044, and 0.062, while $\widehat{\alpha}_2$ shifted from -0.142 to -0.072, 0.020, and 0.151. The increasing bias and reversal of $\widehat{\alpha}_2$ show that unrecorded information directly affecting treatment can distort both effect estimation and interpretation of the belief channel. This stress test therefore characterizes the consequence of violating A2; it does not verify A2 in the observed cohort. The corresponding results are reported in the Control, $\delta = 0.15$, $\delta = 0.30$, and $\delta = 0.60$ columns of Table 2.

Belief-feature ablation. We next examined whether BAM's performance could be reproduced by treating the inferred belief as an ordinary additional covariate. The replication-specific estimate $\widehat{m}_2(42/60)$ was supplied to the propensity model and the doubly robust learner (Robins et al., 2000; Bang & Robins, 2005). As reported in Table 2, IPTW's RMSE changed from 0.220 to 0.222, while DR's RMSE decreased only from 0.240 to 0.231; both remained more than six times BAM's RMSE of 0.036. Thus, simply appending a belief estimate to an existing estimator did not reproduce BAM's improvement. The result supports using belief as a structured interface between the longitudinal

outcome and treatment-decision components rather than as an isolated predictive feature (Tsiatis & Davidian, 2004; Shardell & Ferrucci, 2018). Figure 4 shows the RMSE of BAM and standard baselines in the above experiments.

Comparison with an outcome-only LMM. We removed the treatment-decision likelihood to obtain an outcome-only LMM (Qian et al., 2020). On the $B = 20$ matched simulations, BAM and the LMM achieved similar RMSEs (0.0307 versus 0.0312), and both deteriorated under direct latent selection. The key difference is that the LMM models only longitudinal outcomes, whereas BAM additionally estimates how inferred recovery belief, follow-up time, accumulated exposure, and previous-day treatment shape clinical decisions. Thus, BAM retains high treatment-effect estimation accuracy while providing additional information about the treatment-decision process. The corresponding results are reported in the BAM and LMM rows of the Control, $\delta = 0.15$, $\delta = 0.30$, and $\delta = 0.60$ columns of Table 2.

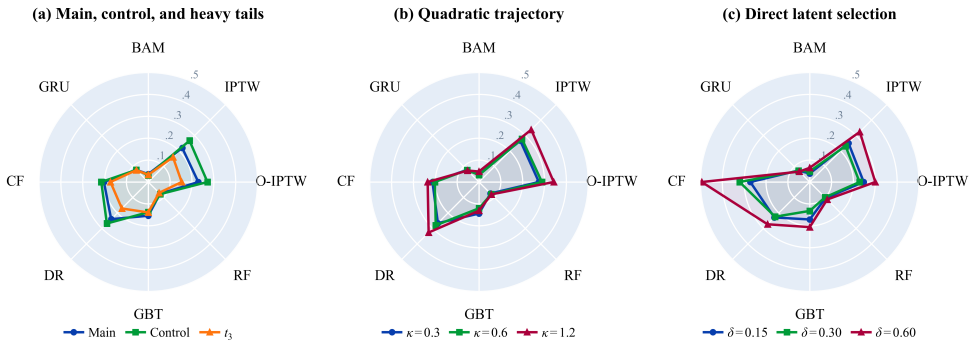


Figure 4: RMSE comparison of BAM with standard baselines across simulation scenarios.

4.2.4 Extensions and Portability

Neural variants. We evaluated two neural extensions without replacing the prespecified spline BAM used for the primary analysis. VI-BAM learned a belief representation without receiving the closed-form expression; its learned belief aligned closely with the closed-form belief (Pearson $r = 0.962$, Spearman $r = 0.953$) and gave $\hat{\beta}(42/60) = 0.075$. Replacing the spline outcome mean with a neural mean retained evidence for the belief channel ($\hat{\alpha}_2 = -0.163$, LR $p = 0.033$) and gave $\hat{\beta}(42/60) = 0.049$. These variants show that the belief-linked decision component can be combined with learned representations, while the spline BAM estimate remains the primary result.

Cross-cohort application to MIMIC-III. Finally, we applied the treatment-decision component to two furosemide–creatinine extractions, vasopressin–creatinine, and insulin–glucose derived from the MIMIC-III demo database (Johnson et al., 2016; 2019). The previous-day treatment coefficient was positive in all four analyses, with $\widehat{\alpha}_5 = 3.44$, 3.58, 4.18, and 4.82, and the corresponding likelihood-ratio statistics ranged from 54.0 to 156.4. Thus, previous-day medication consistently predicted current-day treatment across four medication–outcome settings, showing that the treatment-history channel captured by BAM is also present in intensive-care data. This suggests that the form of the decision model remains transferable to causal estimands in non-rehabilitation settings, while the magnitude of $\hat{\alpha}_5$ may partly reflect the construction of daily exposure from medication-order intervals. Supplementary results are provided in Appendix G.

5 Conclusion and Future Directions

We introduce BAM, a joint framework for longitudinal causal inference in settings where treatment is recorded daily but outcomes are assessed intermittently. BAM converts sparse outcome histories into posterior beliefs about untreated recovery, links these beliefs to treatment decisions, and retains the complete treatment history when estimating time-varying treatment effects. Across the rehabilitation cohort and cohort-calibrated simulations, BAM provided accurate effect estimates while revealing interpretable belief, time, accumulated-exposure, and treatment-continuity channels. We further discuss the limitations of the current study and directions for future work in Appendix L.

AI USE STATEMENT

In this work, we used generative AI tools, including OpenAI ChatGPT and Codex, to assist with checking and refining theorem proofs, modifying and debugging research code, and improving the clarity and presentation of the manuscript. All AI-assisted mathematical arguments were reviewed against the stated assumptions, and all modified code was executed and checked by the authors. The authors made all final methodological and interpretive decisions and take full responsibility for the content of this work.

Generative AI was not used to fabricate or alter clinical observations, autonomously adjudicate patient-level clinical information, or directly produce the numerical results reported in the paper. Cohort-calibrated synthetic data were generated by executable simulation programs implementing the data-generating mechanisms described in the paper, rather than being generated directly by an LLM. Tasks involving interviews, transcription, or qualitative and thematic analysis were not applicable to this work.

The authors reviewed all AI-assisted material. Mathematical statements and proofs were independently checked against the stated assumptions; generated or modified code was executed, tested, and examined using numerical diagnostics and reproducibility checks; reported values were traced to saved computational outputs; and references and literature-based claims were checked against their original sources. The authors made all final methodological, statistical, and interpretive decisions and take full responsibility for the content of this work, including all text, claims, code, and artifacts produced with the assistance of generative AI.

ETHICS STATEMENT

This study uses retrospective clinical records to investigate treatment-effect estimation and the structure of treatment decisions. The rehabilitation cohort data were de-identified, and the cohort does not include untreated patients; all patients received standard clinical care, and no treatment was withheld for research purposes. The rehabilitation cohort contains sensitive patient information and cannot be made publicly available. Only aggregate results are reported, and patient-level records will not be included in the released materials; precautions are taken to prevent re-identification. The MIMIC analysis uses the publicly available, de-identified MIMIC-III Clinical Database Demo v1.4, consistent with its original data license. The findings are intended for methodological research rather than direct treatment recommendations. Their interpretation depends on the stated identification and modeling assumptions, and differences in patient populations, documentation practices, and treatment policies may limit generalizability. The analyses do not establish clinical utility or equitable performance across patient subgroups, nor do they evaluate potential biases or disparate impacts. Further validation is required before clinical application. The authors take responsibility for the final manuscript and all content included in it.

REPRODUCIBILITY STATEMENT

The main manuscript defines BAM's outcome, belief, and decision components in Section 3 and reports the evaluation design and principal findings in Section 4. Appendix A consolidates notation; Appendix B explains every identifying assumption; Appendix C provides complete proofs of the identification, joint-inference, and boundary claims; Appendix D documents estimation and uncertainty calculation; Appendix E gives the rehabilitation, MIMIC-III, and NLP-assisted processing pipelines; Appendix F specifies the simulations, comparators, and metrics; and Appendix G reports complete supplementary tables, figures, ablations, and numerical diagnostics. The anonymous supplementary archive contains executable source code, fixed configurations and seeds, cohort-construction logs, replicate-level outputs, and machine-readable data dictionaries. Public MIMIC-III demo data are cited in the manuscript; access-controlled rehabilitation records are not redistributed, but their deterministic processing rules and audit summaries are supplied so the analysis can be reconstructed in an authorized environment.

REFERENCES

Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005. doi: 10.1111/j.1541-0420.2005.00377.x.

L. Battaglia et al. Inference for regression with variables generated from unstructured data, 2024. Working paper.

Ioana Bica, Ahmed M. Alaa, James Jordon, and Mihaela van der Schaar. Estimating counterfactual treatment outcomes over time through adversarially balanced representations. In *International Conference on Learning Representations*, 2020a. URL https://openreview.net/forum?id=BJg866NFvB.

Ioana Bica, Ahmed M. Alaa, and Mihaela van der Schaar. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 884–895, 2020b. URL https://proceedings.mlr.press/v119/bica20a.html.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. doi: 10.1023/A:1010933404324.

Amiram Catz, Malka Itzkovich, Eli Agranov, Haim Ring, and Arie Tamir. SCIM—spinal cord independence measure: A new disability scale for patients with spinal cord lesions. *Spinal Cord*, 35(12):850–856, 1997.

Kyunghyun Cho, Bart van Merriënboer, Çağlar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pp. 1724–1734, Doha, Qatar, 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL https://aclanthology.org/D14-1179/.

Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric J. Tchetgen Tchetgen. Semiparametric proximal causal inference. *Journal of the American Statistical Association*, 119(546):1348–1359, 2024. doi: 10.1080/01621459.2023.2191817.

Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Alicia Curth, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine learning for predicting treatment outcomes. *Nature Medicine*, 30(4):958–968, 2024. doi: 10.1038/s41591-024-02902-1.

Norine Foley, J. Andrew McClure, Matthew Meyer, Katherine Salter, Yves Bureau, and Robert Teasell. Inpatient rehabilitation following stroke: Amount of therapy received and associations with functional recovery. *Disability and Rehabilitation*, 34(25):2132–2138, 2012. doi: 10.3109/09638288.2012.676145.

Dennis Frauen, Konstantin Hess, and Stefan Feuerriegel. Model-agnostic meta-learners for estimating heterogeneous treatment effects over time. In *International Conference on Learning Representations*, 2025.

Jerome H. Friedman. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5):1189–1232, 2001. doi: 10.1214/aos/1013203451.

Robin Henderson, Peter Diggle, and Angela Dobson. Joint modelling of longitudinal measurements and event time data. *Biostatistics*, 1(4):465–480, 2000. doi: 10.1093/biostatistics/1.4.465.

Konstantin Hess and Stefan Feuerriegel. Stabilized neural prediction of potential outcomes in continuous time. In *International Conference on Learning Representations*, 2025.

Konstantin Hess, Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Bayesian neural controlled differential equations for treatment effect estimation. In *International Conference on Learning Representations*, 2024.

Jsabel Hodel, Cristina Ehrmann, Anke Scheel-Sailer, Gerold Stucki, Jerome E. Bickenbach, and Birgit Prodinger. Identification of classes of functioning trajectories and their predictors in individuals with spinal cord injury attending initial rehabilitation in switzerland. *Archives of Rehabilitation Research and Clinical Translation*, 3(2):100121, 2021. doi: 10.1016/j.arrct.2021.100121.

Alistair Johnson, Tom Pollard, and Roger Mark. MIMIC-III Clinical Database Demo. *PhysioNet*, April 2019. doi: 10.13026/C2HM2Q. URL https://doi.org/10.13026/C2HM2Q. Version 1.4.

Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3:160035, 2016.

Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.

Robert A. Keith, Carl V. Granger, Byron B. Hamilton, and F. S. Sherwin. The functional independence measure: A new tool for rehabilitation. *Advances in Clinical Rehabilitation*, 1:6–18, 1987.

Lara A. Kimmel, Anne E. Holland, Natasha Lannin, Elton R. Edwards, Richard S. Page, Andrew Bucknill, Raphael Hau, and Belinda J. Gabbe. Clinicians' perceptions of decision making regarding discharge from public hospitals to in-patient rehabilitation following trauma. *Australian Health Review*, 41(2):192–200, 2017. doi: 10.1071/AH16031.

Steven C. Kirshblum, Stephen P. Burns, Fin Biering-Sorensen, William Donovan, Daniel E. Graves, Amitabh Jha, Mark Johansen, Linda Jones, Andrei Krassioukov, M. J. Mulcahey, Mary Schmidt-Read, and William Waring. International standards for neurological classification of spinal cord injury (revised 2011). *The Journal of Spinal Cord Medicine*, 34(6):535–546, 2011.

Nan M. Laird and James H. Ware. Random-effects models for longitudinal data. *Biometrics*, 38(4): 963–974, 1982. doi: 10.2307/2529876.

Kuan Liu, Olli Saarela, George Tomlinson, Brian M. Feldman, and Eleanor Pullenayegum. A bayesian latent class approach to causal inference with longitudinal data. *Statistical Methods in Medical Research*, 34(1):55–68, 2025. doi: 10.1177/09622802241298704.

Mark P. McGlinchey, Lizz Paley, Alex Hoffman, Abdel Douiri, and Anthony G. Rudd. Physiotherapy provision to hospitalised stroke patients: Analysis from the UK Sentinel Stroke National Audit Programme. *European Stroke Journal*, 4(1):75–84, 2019. doi: 10.1177/2396987318800543.

Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Causal transformer for estimating counterfactual outcomes. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 15293–15329, 2022. URL https://proceedings.mlr.press/v162/melnychuk22a.html.

Wang Miao, Zhi Geng, and Eric J. Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018. doi: 10.1093/biomet/asy038.

Kevin M. Murphy and Robert H. Topel. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 1985.

John A. Nelder and Roger Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965. doi: 10.1093/comjnl/7.4.308.

Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, 2 edition, 2006. doi: 10.1007/978-0-387-40065-5.

Kenneth J. Ottenbacher, Pam M. Smith, Sandra B. Illig, Richard T. Linn, Glenn V. Ostir, and Carl V. Granger. Trends in length of stay, living setting, functional outcome, and mortality following medical rehabilitation. *JAMA*, 292(14):1687–1695, 2004. doi: 10.1001/jama.292.14.1687.

Tom J. Pollard, Alistair E. W. Johnson, Jesse D. Raffa, and Roger G. Mark. The MIMIC-III clinical database demo, 2019. Version 1.4, requires no credentialing.

Janet M. Powell, Nancy R. Temkin, Joan E. Machamer, and Sureyya S. Dikmen. Nonrandomized studies of rehabilitation for traumatic brain injury: Can they determine effectiveness? *Archives of Physical Medicine and Rehabilitation*, 83(9):1235–1244, 2002. doi: 10.1053/apmr.2002.34556.

Tianchen Qian, Predrag Klasnja, and Susan A. Murphy. Linear mixed models with endogenous covariates: Modeling sequential treatment effects with application to a mobile health study. *Statistical Science*, 35(3):375–390, 2020. doi: 10.1214/19-STS720.

James M. Robins, Miguel A. Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000. doi: 10.1097/00001648-200009000-00011.

Holger Schielzeth, Niels J. Dingemanse, Shinichi Nakagawa, David F. Westneat, Hassen Allegue, Céline Teplitsky, Denis Réale, Ned A. Dochtermann, László Zsolt Garamszegi, and Yimen G. Araya-Ajoy. Robustness of linear mixed-effects models to violations of distributional assumptions. *Methods in Ecology and Evolution*, 11(9):1141–1152, 2020. doi: 10.1111/2041-210X.13434.

Nabeel Seedat, Fergus Imrie, Alexis Bellot, Zhaozhi Qian, and Mihaela van der Schaar. Continuous-time modeling of counterfactual outcomes using neural controlled differential equations. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 19497–19521, 2022. URL https://proceedings.mlr.press/v162/seedat22b.html.

Michelle Shardell and Luigi Ferrucci. Joint mixed-effects models for causal inference with longitudinal data. *Statistics in Medicine*, 37(5):829–846, 2018. doi: 10.1002/sim.7567.

Colleen M. Sitlani, Patrick J. Heagerty, Emily A. Blood, and Tor D. Tosteson. Longitudinal structural mixed models for the analysis of surgical trials with noncompliance. *Statistics in Medicine*, 31(16):1738–1760, 2012. doi: 10.1002/sim.4510.

Eric J. Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *Journal of the American Statistical Association*, 2020.

Eric J. Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal learning. *Statistical Science*, 39(3):375–390, 2024. doi: 10.1214/23-STS911.

Anastasios A. Tsiatis and Marie Davidian. Joint modeling of longitudinal and time-to-event data: An overview. *Statistica Sinica*, 14(3):809–834, 2004.

Stijn Vansteelandt and Marshall Joffe. Structural nested models and g-estimation: The partially realized promise. *Statistical Science*, 29(4):707–731, 2014. doi: 10.1214/14-STS493.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. doi: 10.1080/01621459.2017.1319839.

Carolee J. Winstein, Joel Stein, Ross Arena, Barbara Bates, Leora R. Cherney, Steven C. Cramer, Frank Deruyter, Janice J. Eng, Beth Fisher, Richard L. Harvey, Catherine E. Lang, Marilyn MacKay-Lyons, Kenneth J. Ottenbacher, Sue Pugh, Mathew J. Reeves, Lorie G. Richards, William Stiers, and Richard D. Zorowitz. Guidelines for adult stroke rehabilitation and recovery: A guideline for healthcare professionals from the American Heart Association/American Stroke Association. *Stroke*, 47(6):e98–e169, 2016. doi: 10.1161/STR.0000000000000098.

Yong Wu, Yanwei Fu, Shouyan Wang, and Xinwei Sun. Doubly robust proximal causal learning for continuous treatments. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=TjGJFkU3xL.

Michael S. Wulfsohn and Anastasios A. Tsiatis. A joint model for survival and longitudinal data measured with error. *Biometrics*, 53(1):330–339, 1997. URL https://pubmed.ncbi.nlm.nih.gov/9147598/.

A NOTATION

Table 3 consolidates the objects that connect the daily treatment process, intermittent outcomes, the posterior belief, and the causal estimands. This separation makes the two BAM interfaces explicit: treatment history enters the outcome model through $X_i(t)$, while outcome history enters the decision model through $\widehat{m}_{2i}(t)$.

Table 3: Notation used by BAM.

Symbol	Definition	Role in BAM
i, d, j	Patient, calendar day, and assessment index	Separates dense daily treatment records from sparse outcome assessments.
T, t_{ij}	Analysis horizon and normalized assessment time $t_{ij} = d_{ij}/T$	$T = 60$ for rehabilitation; cohort-specific for external analyses.
$D_i(d), D_{ij}$	Daily treatment and assessment-day treatment $D_i(d_{ij})$	$D_i(d) = 1$ indicates treatment on day d.
Y_{ij}, a_i	Standardized functional outcome and admission age	Intermittent response and baseline adjustment.
$\bar{\mathcal{H}}_i(t)$	Outcomes and treatments recorded strictly before t	Information available before the current decision.
$b_i = (b_{1i}, b_{2i})^\top$	Latent untreated level and recovery-rate deviations	Patient-specific heterogeneity, with $b_i \sim \mathcal{N}(0, \Sigma)$.
$\mu(t), \gamma$	Population untreated trajectory and age coefficient	Fixed part of the untreated outcome trajectory.
$\varepsilon_{ij}, \sigma^2$	Outcome error and its variance	$\varepsilon_{ij} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$.
$B(s), \beta$	Prespecified spline basis and coefficient vector	$\beta(s) = B(s)^\top \beta$ is the single-day blip curve.
$X_i(t)$	$\sum_{d \le Tt} D_i(d) B(d/T)$	Carries the complete daily treatment history into the outcome model.
$r_i(t), G_i(t)$	Historical outcome residuals and random-effects design	Inputs to Gaussian posterior updating before time t.
$\widehat{m}_i(t)$	$\mathbb{E}[b_i \mid \bar{\mathcal{H}}_i(t)]$	Posterior belief about the untreated patient trajectory.
$\widehat{m}_{2i}(t)$	Slope component of $\widehat{m}_i(t)$	Decision-relevant belief about untreated recovery speed.
$C_i(t)$	$\sum_{d < Tt} D_i(d)$	Accumulated exposure strictly before the current day.
α	$(\alpha_0, \alpha_2, \alpha_3, \alpha_4, \alpha_5)^\top$	Intercept, belief, time, exposure, and continuity channels.
q_{ij}	$\Pr(D_{ij} = 1 \mid \bar{\mathcal{H}}_i(t_{ij}))$	Assessment-day treatment probability.
η	$(\mu, \gamma, \beta, \sigma^2, \Sigma)$	Outcome-side parameters that also generate the belief.
θ	$(\eta^\top, \alpha^\top)^\top$	Full jointly estimated parameter vector.
$\Gamma(t)$	$\sum_{d \le Tt} \beta(d/T)$	Cumulative contrast for daily treatment through time t.
I, S_η, S_α	Fisher information and Schur complements	Joint uncertainty propagation across the two likelihood components.

B INTERPRETATION OF ASSUMPTIONS

A1: causal consistency, no interference, and no anticipation. The observed outcome is the potential outcome indexed by the realized treatment history; treatment of one patient does not alter another patient's outcome; and treatment cannot change outcomes recorded before administration. A1 aligns the potential-outcome estimand with the recorded longitudinal process. The within-day order $\bar{\mathcal{H}}_i(t_{ij}) \to D_{ij} \to Y_{ij}$ additionally states which same-day event is causally prior.

A2: daily history-conditional ignorability. After conditioning on all recorded information available before day d, treatment is independent of the latent untreated trajectory and future potential

outcomes. The assumption is daily, including dates without a directly observed score. It does not require a score every day: past daily treatments remain recorded, and past intermittent outcomes update the belief. A2 rules out an additional direct path from unrecorded recovery information to treatment after recorded history has been conditioned upon. Its empirical boundary is deliberately probed by the $\delta b_{2i}/\tau_2$ stress test in Appendix G.2.

A3: belief-state sufficiency. On assessment days, the recorded history can be compressed for decision modeling into the posterior slope belief, normalized time, accumulated exposure, and previous-day treatment. A3 is a structural dimension-reduction statement about the treatment process. It is distinct from the logistic-linear implementation: Theorem 1 allows the decision function g to be any measurable function of the stated summaries, whereas the logistic form makes the decision component estimable and interpretable in finite samples.

A4: overlap and full rank. For histories represented in the target population, both treatment states must occur with non-negligible probability. In addition, the supported spline vectors and fixed-effect design must span their parameter spaces. Overlap prevents the causal contrast from being based on deterministic extrapolation; full rank converts supported conditional contrasts into a unique coefficient vector β.

A5: regularity and conditionally non-informative assessment. Variance parameters remain away from singular boundaries, the true parameter is interior, and standard smoothness and moment conditions hold. Conditional on recorded history, the assessment process supplies no additional parameter-relevant selection information for the main likelihood. This condition enables a clean sequential likelihood while preserving the actual irregular assessment schedule.

What the assumptions jointly deliver. A1 connects observed and potential outcomes, A2 makes assessment-day treated–untreated contrasts causal conditional on history, A3 supplies BAM's interpretable belief-mediated decision representation, A4 turns those contrasts into a unique time-varying curve, and A5 supports likelihood-based inference. Together they identify the full spline coefficient vector rather than a single landmark effect; day 42 is only a common reporting point in the experiments.

C Complete Mathematical Proofs

C.1 Preliminary Lemmas

Lemma 2 (History-conditional balance of the latent trajectory)**.** *Under A2, for every supported assessment day d_{ij} and every integrable function h,*

$$\mathbb{E}\{h(b_i) \mid \bar{\mathcal{H}}_i(t_{ij}), D_{ij} = 1\} = \mathbb{E}\{h(b_i) \mid \bar{\mathcal{H}}_i(t_{ij}), D_{ij} = 0\} = \mathbb{E}\{h(b_i) \mid \bar{\mathcal{H}}_i(t_{ij})\}.$$

Proof. A2 implies $D_{ij} \perp\!\!\!\perp b_i \mid \bar{\mathcal{H}}_i(t_{ij})$. For any bounded measurable h, conditional independence gives

$$\mathbb{E}\{h(b_i)\mathbf{1}(D_{ij} = a) \mid \bar{\mathcal{H}}_i\} = \mathbb{E}\{h(b_i) \mid \bar{\mathcal{H}}_i\} \Pr(D_{ij} = a \mid \bar{\mathcal{H}}_i),$$

for $a \in \{0, 1\}$. A4 makes both conditional probabilities positive on the target support, so division proves the claim. Extension to every integrable h follows by truncation and dominated convergence. □

Lemma 3 (Assessment-day conditional contrast)**.** *Under the outcome model, A1, and A2, the supported assessment-day contrast satisfies*

$$\mathbb{E}(Y_{ij} \mid \bar{\mathcal{H}}_i(t_{ij}), D_{ij} = 1) - \mathbb{E}(Y_{ij} \mid \bar{\mathcal{H}}_i(t_{ij}), D_{ij} = 0) = B(t_{ij})^\top \beta.$$

Proof. Condition on $\bar{\mathcal{H}}_i(t_{ij})$. Every treatment before d_{ij}, admission age, normalized time, and the corresponding historical part of $X_i(t_{ij})$ are then fixed. Changing only D_{ij} from zero to one adds $B(t_{ij})$ to $X_i(t_{ij})$. By A1 and the no-anticipation ordering, the current treatment changes the conditional mean by $B(t_{ij})^\top \beta$. By Lemma 2, the conditional distribution of b_i is identical across the two current-treatment groups. The mean-zero error distribution is also unchanged. Hence the

untreated fixed component, posterior latent component, and error cancel in the difference, leaving $B(t_{ij})^\top \beta$. □

Lemma 4 (Uniqueness of the spline coefficients)**.** *If $B(t)^\top \beta = B(t)^\top \beta'$ for every supported assessment time and the supported spline design has rank K, then $\beta = \beta'$.*

Proof. Let $\mathcal{B}$ be a matrix containing K linearly independent supported row vectors $B(t)^\top$. The assumed equality implies $\mathcal{B}(\beta - \beta') = 0$. Since $\mathcal{B}$ is nonsingular, $\beta - \beta' = 0$. □

C.2 Proof of Theorem 1

Proof. Consider two parameterizations (η, g) and (η', g') that induce the same observed-data law. Equality of the observed law implies equality of every regular conditional distribution derived from it, including the conditional distribution of Y_{ij} given a supported pre-decision history and $D_{ij} = a$, $a \in \{0, 1\}$. Consequently, their observed conditional mean contrasts are equal.

By Lemma 3, the first parameterization yields $B(t_{ij})^\top \beta$ and the second yields $B(t_{ij})^\top \beta'$. Therefore

$$B(t)^\top \beta = B(t)^\top \beta'$$

for every supported assessment time. A4 and Lemma 4 imply $\beta = \beta'$. It follows immediately that $B(d/T)^\top \beta$ is identified at every day represented by the prespecified basis, including days without a directly observed score. Those non-assessment-day values are basis-implied evaluations of the identified curve, not separately estimated day-specific contrasts. Finally, $\Gamma(t) = \sum_{d \le Tt} B(d/T)^\top \beta$ is a known linear functional of the identified vector and is therefore identified.

The argument uses A3 only to define the decision-relevant representation adopted by BAM. It never invokes the logistic-linear form of g, so identification is form-free within the class of measurable decision functions satisfying A3. □

C.3 Proof of Theorem 2

Proof. Write $\ell_n(\theta) = \sum_{i=1}^n \ell_i(\theta)$, $U_i(\theta) = \partial \ell_i(\theta)/\partial\theta$, and $H_i(\theta) = \partial U_i(\theta)/\partial\theta^\top$. Patients are independent clusters. Under the stated identifiability, compact-neighborhood, continuity, and domination conditions, a uniform law of large numbers gives

$$n^{-1}\ell_n(\theta) \longrightarrow_p M(\theta) := \mathbb{E}_{\theta_0}\{\ell_i(\theta)\}$$

uniformly on a neighborhood of θ_0. Correct specification and identifiability make $M(\theta)$ uniquely maximized at θ_0, so the argmax theorem yields $\widehat{\theta} \to_p \theta_0$.

Order each assessment-day contribution as decision followed by outcome. The decision score is measurable after the pre-decision history and current D_{ij}; the subsequent outcome score has conditional mean zero given that enlarged history. Correctly specified conditional likelihood factors therefore generate martingale-difference score increments. Cross-products between distinct increments have expectation zero by iterated expectation. The variation-independent non-assessment-day laws do not contain θ and hence contribute no score. Thus

$$\mathbb{E}\{U_i(\theta_0)\} = 0, \qquad \mathbb{E}\{U_i(\theta_0)U_i(\theta_0)^\top\} = -\mathbb{E}\{H_i(\theta_0)\} = I.$$

The patient-level multivariate central limit theorem and Hessian law of large numbers give

$$n^{-1/2}\sum_{i=1}^n U_i(\theta_0) \xrightarrow{d} \mathcal{N}(0, I), \qquad -n^{-1}\sum_{i=1}^n H_i(\widetilde{\theta}) \to_p I$$

for every intermediate $\widetilde{\theta} \to_p \theta_0$. Taylor expansion of the likelihood equation $U_n(\widehat{\theta}) = 0$ around θ_0 now yields

$$\sqrt{n}(\widehat{\theta} - \theta_0) = I^{-1} n^{-1/2} \sum_{i=1}^n U_i(\theta_0) + o_p(1) \xrightarrow{d} \mathcal{N}(0, I^{-1}).$$

Partition I conformably with $\theta = (\eta^\top, \alpha^\top)^\top$. The block inverse formula gives

$$(I^{-1})_{\alpha\alpha} = \{I_{\alpha\alpha} - I_{\alpha\eta} I_{\eta\eta}^{-1} I_{\eta\alpha}\}^{-1} = S_\alpha^{-1}$$

and, symmetrically, $(I^{-1})_{\eta\eta} = S_\eta^{-1}$. Positive definiteness of I implies positive definiteness of both Schur complements. Moreover,

$$I_{\alpha\alpha} - S_\alpha = I_{\alpha\eta} I_{\eta\eta}^{-1} I_{\eta\alpha} \succeq 0, \qquad I_{\eta\eta} - S_\eta = I_{\eta\alpha} I_{\alpha\alpha}^{-1} I_{\alpha\eta} \succeq 0.$$

Because inversion reverses the Loewner order on positive-definite matrices, joint estimation has variances $S_\alpha^{-1} \succeq I_{\alpha\alpha}^{-1}$ and $S_\eta^{-1} \succeq I_{\eta\eta}^{-1}$ relative to treating the other block as known.

For a single Bernoulli decision, let $p_{ij} = q_{ij}(\eta_0, \alpha_0)$ and $z_{ij} = z_{ij}(\eta_0)$. Its log likelihood has scores

$$U_{\alpha,ij}^D = (D_{ij} - p_{ij}) z_{ij}, \qquad U_{\eta,ij}^D = (D_{ij} - p_{ij}) \alpha_2 \frac{\partial \widehat{m}_{2i}(t_{ij}; \eta_0)}{\partial \eta}.$$

Conditional on pre-decision history, $\mathbb{E}(D_{ij} - p_{ij})^2 = p_{ij}(1 - p_{ij})$. The same-step cross outer product is therefore

$$\alpha_2 p_{ij}(1 - p_{ij}) z_{ij} \frac{\partial \widehat{m}_{2i}(t_{ij}; \eta_0)}{\partial \eta^\top}.$$

All cross-time terms vanish by the martingale-difference argument, and the outcome likelihood has no α derivative. Summing over j and averaging over independent patient clusters establishes Equation equation 9. The expression is proportional to α_2 and hence vanishes when $\alpha_2 = 0$.

For any day s, define $h_s(\theta) = B(s)^\top \beta$. Its gradient is zero outside the β coordinates and equals $B(s)$ inside them. The delta method gives

$$\sqrt{n}\{\widehat{\beta}(s) - \beta_0(s)\} \xrightarrow{d} \mathcal{N}\big(0, \nabla h_s(\theta_0)^\top I^{-1} \nabla h_s(\theta_0)\big),$$

which justifies pointwise intervals obtained from the full joint covariance. A fixed-weight smoothness penalty contributes $O(1)$ to an $O(n)$ objective and therefore does not change this first-order limit. □

C.4 Proof of Proposition 1

Proof. Fix a supported pre-decision history H and suppress patient and time indices. Let the latent untreated slope satisfy $b_2 \mid H \sim \mathcal{N}(m, v)$, write $Z = (b_2 - m)/\sqrt{v}$, and consider the threshold decision

$$D = \mathbf{1}\{a(H) + \delta Z + U > 0\}, \qquad U \sim \mathcal{N}(0,1), \quad U \perp Z \mid H.$$

When $\delta \neq 0$, treatment depends directly on the latent slope after conditioning on H, violating A2. Put $c = a(H)/\sqrt{1+\delta^2}$ and $\rho = \delta/\sqrt{1+\delta^2}$. Standard truncated-normal calculations give

$$\mathbb{E}(Z \mid H, D=1) = \rho \frac{\phi(c)}{\Phi(c)}, \qquad \mathbb{E}(Z \mid H, D=0) = -\rho \frac{\phi(c)}{1 - \Phi(c)}.$$

At an assessment time t, the outcome contains the latent component tb_2 and the current causal blip $B(t)^\top \beta D$. Consequently,

$$\begin{aligned} &\mathbb{E}(Y \mid H, D=1) - \mathbb{E}(Y \mid H, D=0) \\ &\quad = B(t)^\top \beta + t\sqrt{v}\, \rho \phi(c) \left\{ \frac{1}{\Phi(c)} + \frac{1}{1 - \Phi(c)} \right\}. \end{aligned}$$

The second term is a Mills-ratio selection contrast. It is zero when $\delta = 0$ and generally nonzero otherwise.

If the latent-selection law is unrestricted after A2 is removed, the observed law determines the sum of the causal and selection components but supplies no observed variable that labels either component. For any candidate causal blip, an alternative latent-selection distribution can shift the second component by the opposite amount while preserving the observed conditional distribution of (D, Y); this is the standard observational equivalence underlying unmeasured-confounding nonidentification. The Gaussian threshold calculation displays the ambiguity explicitly at the conditional-mean level. Therefore $B(t)^\top \beta$, and hence β, is not identified in general without A2 or an additional source of identification such as a valid proxy, instrument, negative control, or a justified parametric selection restriction (Miao et al., 2018; Tchetgen Tchetgen et al., 2024; Cui et al., 2024). □

D Estimation and Implementation Details

Belief update. For every assessment time, BAM builds $r_i(t)$ only from outcomes observed strictly before the current decision. With $V_i(t) = \{\Sigma^{-1} + \sigma^{-2} G_i(t)^\top G_i(t)\}^{-1}$, the posterior mean and covariance are

$$\widehat{m}_i(t) = V_i(t)\sigma^{-2} G_i(t)^\top r_i(t), \qquad \mathrm{Var}\{b_i \mid \bar{\mathcal{H}}_i(t)\} = V_i(t).$$

The closed form makes the belief differentiable with respect to every outcome-side parameter and exposes the exact interface through which sparse outcomes inform treatment decisions.

Joint objective. The optimization target is $\ell(\eta, \alpha) = \sum_i \{\ell_{Y,i}(\eta) + \ell_{D,i}(\eta, \alpha)\}$. Daily treatment records enter every later $X_i(t)$ and the decision summaries even when no outcome is observed that day. Assessment-day treatment decisions contribute Bernoulli likelihood terms. This preserves the full exposure history without pretending that an unobserved daily outcome exists.

Numerical optimization. Variance parameters are represented on unconstrained scales, with the random-effects covariance parameterized through a Cholesky factor. Fits are initialized from the longitudinal outcome component and a treatment logistic regression, followed by joint optimization from multiple dispersed starts. Candidate solutions are ranked by the same unpenalized joint likelihood; convergence is checked using objective agreement, score norms, parameter-boundary checks, and the observed-information spectrum. When a smoothness penalty is used for spline stabilization, reported first-order covariance is computed from the unpenalized likelihood at the selected estimate.

Uncertainty. The reported covariance is the inverse of the full observed information for $\theta = (\eta^\top, \alpha^\top)^\top$, rather than a covariance obtained after freezing either likelihood block. Pointwise uncertainty for $\beta(s)$ is computed by the delta method using the β rows and columns of that full covariance. The matched experiment in Appendix G.5 verifies the predicted direction of the Schur-complement correction and shows that BAM retains precise effect estimation at the empirically observed coupling strength.

Optimization and audit trail. The supplementary archive records, for every simulation replicate and observed-data fit, the initialization identifier, final objective, convergence code, score norm, observed-information eigenvalue check, and retained estimate. It also contains replicate-level bias/RMSE inputs, the complete treatment and outcome construction logs, and NLP evidence tables with source timing. This audit trail makes every reported aggregate traceable to patient-level construction rules and replicate-level numerical output.

E Data Construction and Processing

E.1 Rehabilitation Cohort

The analysis begins from encounter-linked rehabilitation records and applies the following deterministic pipeline.

1. Define admission as day 0 and retain follow-up through day 60. Assign a stable anonymous patient index and remove records that cannot be linked to a valid study encounter.
2. Harmonize neurological and functional assessment fields to a common direction in which larger values indicate better function, using the documented scoring directions of ISNCSCI, FIM, and SCIM (Kirshblum et al., 2011; Keith et al., 1987; Catz et al., 1997). Standardize the analysis outcome using the transformation fixed for the cohort.
3. Collapse valid assessments to their calendar day and order same-day information as prior history, current-day treatment, then assessment. The resulting primary cohort contains 419 patients and 3,674 observed outcome assessments.
4. Construct $D_i(d) = 1$ when any qualifying rehabilitation therapy is recorded on calendar day d, and $D_i(d) = 0$ otherwise. Preserve all daily indicators from day 0 through the end of observed follow-up.

5. At each assessment, build $X_i(t)$ from the complete prior and current daily treatment sequence, $C_i(t)$ from treatment strictly before the current day, and the previous-day indicator from $D_i(d-1)$. Build the belief only from outcomes observed strictly before the decision.

6. Retain the empirical assessment schedule. BAM does not impute unobserved daily outcomes; it uses the longitudinal model to update the posterior belief whenever an assessment becomes available.

E.2 NLP-assisted Construction of Assessment Outcomes and Clinical Indicators

A rule-based natural language processing pipeline was used to transform clinical narratives and semi-structured records into standardized variables for longitudinal analysis. The extracted information included total and component scores from rehabilitation assessments, neurological, motor, and sensory indicators, assessment dates, and clinical phenotypes such as injury grade, injury completeness, sacral-function evidence, and neurological level. Definitions of functional independence, spinal-cord independence, and neurological status followed the corresponding assessment instruments and clinical standards (Keith et al., 1987; Catz et al., 1997; Kirshblum et al., 2011). Each extracted item retained its anonymous encounter identifier, source document, source field, documentation time, variable name, normalized value, and supporting text span, allowing every analysis variable to be traced to its original record.

The extracted records underwent terminology normalization, range validation, temporal alignment, duplicate reconciliation, and conflict detection before being organized into patient-level longitudinal data. Standardized functional scores obtained through this pipeline formed the BAM outcome Y_{ij}, and their assessment dates defined d_{ij}; the remaining clinical indicators supported cohort characterization, data-quality assessment, and subsequent heterogeneity analyses. Distinct instruments and grading systems retained their original definitions. For example, Frankel grades were not automatically converted into AIS/ASIA grades, and records with conflicting or insufficient evidence within an encounter were flagged for review rather than resolved algorithmically. This procedure organized scores and clinical indicators distributed across multiple records while preserving their provenance, timing, and clinical meaning.

E.3 MIMIC-III Demo Analyses

The external decision-structure analysis uses the public 100-patient MIMIC-III Clinical Database Demo (Johnson et al., 2016; Pollard et al., 2019). For each medication–outcome pair, prescription orders are converted into daily covered spans, overlapping spans are merged, and $D_i(d)$ indicates medication coverage on day d. Laboratory measurements are restricted to the corresponding outcome, aligned to calendar day, and ordered after that day's treatment indicator. Time is normalized to the 14-day analysis horizon. The four prespecified pairs are two furosemide–creatinine extractions, vasopressin–creatinine, and insulin–glucose. These analyses test portability of the decision channels; they do not reuse or reinterpret the rehabilitation effect curve.

E.4 Missingness and Temporal Integrity

Missing outcome days remain missing by design. No forward filling is used to create artificial assessments. Daily treatments are retained because they are directly recorded and form the exposure history. All history-derived covariates are constructed with strict temporal cutoffs, preventing the current or future assessment from entering the belief used for the same day's decision.

The rehabilitation cohort contains sensitive clinical records and cannot be made publicly available because of patient-privacy and institutional data-governance requirements. To support reproducibility, we provide anonymized analysis code, simulation programs, data schemas, cohort-construction rules, and complete preprocessing and evaluation procedures in the supplementary material. The publicly available MIMIC-III demo data can be accessed under its original data-use conditions.

F Simulation Design, Comparators, and Metrics

F.1 Cohort-Calibrated Data-Generating Process

Each replicate contains 419 patients and reuses the empirical rehabilitation assessment schedules, giving approximately 3,674 intermittent assessments per replicate. Patient-level untreated heterogeneity is generated from the fitted two-dimensional random-effects distribution. Outcomes are generated from the fitted population trajectory, admission-age adjustment, complete daily treatment history, and a treatment-effect curve calibrated so that $\beta(42/60) = 0.088$. Daily treatment is generated from the fitted belief, normalized-time, accumulated-exposure, and previous-day channels. The main benchmark uses $B = 40$ independent replicates; matched robustness and component experiments use $B = 20$.

Trajectory misspecification. The reference generator uses Gaussian errors and linear patient-specific trajectories. Robustness scenarios replace the errors by variance-matched t_3 errors or add $b_{3i}t^2$, with $\kappa = \mathrm{SD}(b_3)/\mathrm{SD}(b_2) \in \{0.3, 0.6, 1.2\}$. BAM continues to fit its prespecified linear-Gaussian belief model, making these genuine out-of-model tests.

Identification-boundary stress test. The treatment linear predictor receives the additional term $\delta b_{2i}/\tau_2$, where $\tau_2 = \mathrm{SD}(b_{2i})$ and $\delta \in \{0.15, 0.30, 0.60\}$. This creates a direct unrecorded selection channel after conditioning on recorded history and therefore crosses A2.

F.2 Comparator Definitions

Table 4: Comparators used in the cohort-calibrated experiments.

Method	Role in the comparison
LMM	Outcome-only version retaining BAM's 16-parameter longitudinal outcome component but omitting the treatment-decision likelihood (Qian et al., 2020).
IPTW	Inverse-probability weighting based on the fitted recorded-history treatment model (Robins et al., 2000).
O-IPTW	Oracle weighting using the data-generating treatment probabilities; separates propensity-estimation error from weighting variability.
RF	Random-forest g-computation using recorded longitudinal summaries (Breiman, 2001).
GBT	Gradient-boosted-tree g-computation using the same comparison information set (Friedman, 2001).
DR	Doubly robust learner combining fitted treatment and outcome regressions (Bang & Robins, 2005).
CF	Causal forest fitted to the prespecified longitudinal summaries (Wager & Athey, 2018).
GRU	Gated recurrent sequence model receiving the longitudinal treatment/outcome sequence and observation masks (Cho et al., 2014).

All methods are evaluated on the same replicate and target the same day-42 single-day effect. The fixed preprocessing, random seeds, method configurations, and replicate-level outputs are included in the anonymous supplementary archive.

F.3 Evaluation Metrics

For replicate estimates $\widehat{\psi}_b$ of $\psi = \beta(42/60)$,

$$\mathrm{Bias} = B^{-1}\sum_{b=1}^{B}(\widehat{\psi}_b - \psi), \qquad \mathrm{RMSE} = \left\{B^{-1}\sum_{b=1}^{B}(\widehat{\psi}_b - \psi)^2\right\}^{1/2}.$$

Matched method differences use the within-replicate contrast before computing the Monte Carlo standard error. Observed-cohort decision channels are evaluated by likelihood-ratio statistics comparing the full fit with the corresponding coefficient fixed to zero.

G Complete supplementary results

G.1 Simulation benchmark and robustness

BAM achieves the lowest RMSE among standard comparators in the main benchmark and throughout all heavy-tailed and trajectory-shape perturbations. The result is strongest precisely where the data structure is hardest: daily treatment, sparse outcomes, and recovery-dependent decisions. The δ rows then produce the predicted transition at the identification boundary, turning A2 into an empirically visible scope condition rather than an abstract assumption.

G.2 Identification-boundary response

As δ increases from 0 to 0.15, 0.30, and 0.60, BAM's bias moves from 0.009 to 0.031, 0.044, and 0.062, while $\widehat{\alpha}_2$ moves from -0.142 to -0.072, 0.020, and 0.151. The simultaneous increase in causal-effect bias and reversal of the belief-channel estimate is the model's predicted signature of an unrecorded direct selection path. The experiment therefore anchors BAM's reliable operating region and demonstrates that the belief coefficient is informative about how treatment selection is represented.

G.3 Belief-feature ablation

Table 5: Belief-feature ablation in the cohort-calibrated benchmark.

Estimator	Original RMSE	With $\widehat{m}_2(42/60)$	BAM RMSE
IPTW	0.220	0.222	0.036
DR	0.240	0.231	0.036

Appending the estimated belief to an existing learner leaves IPTW unchanged and only modestly improves DR; both remain more than six times BAM's RMSE. BAM's gain therefore comes from the structured belief interface and joint likelihood, not from adding one more predictive feature. This is consistent with the broader generated-regressor literature: when a learned state links estimating equations, its role is defined by the joint statistical structure rather than by its marginal predictive content alone (Murphy & Topel, 1985; Battaglia et al., 2024).

G.4 Outcome-only LMM comparison

Table 6: Matched $B = 20$ comparison with the outcome-only LMM; the paired-difference Monte Carlo SE is 0.0004.

Method	Bias	Empirical SD	RMSE	Paired difference
BAM	0.0092	0.0300	0.0307	-0.0001
Outcome-only LMM	0.0093	0.0305	0.0312	reference

BAM preserves the high precision of the correctly specified longitudinal outcome model while adding estimable belief, time, accumulated-exposure, and continuity channels. The comparison isolates BAM's distinctive contribution: richer information about the treatment process is obtained without sacrificing effect-estimation accuracy.

G.5 JOINT-UNCERTAINTY EXPERIMENT

Table 7: Finite-sample comparison of conditional and full-joint uncertainty for $\widehat{\beta}(42/60)$ ($B = 20$ per scenario).

Scenario	True α_2	Mean $\widehat{\alpha}_2$	η-known SE	Full-joint SE	Relative increase
Active belief channel	-0.15085	-0.14195	0.022304	0.022318	0.060%
Inactive belief channel	0	0.00670	0.023605	0.023612	0.030%

The full-joint standard error is larger in the active-channel experiment, exactly as the Schur-complement result predicts. Its small magnitude at the cohort-calibrated channel strength is also valuable: BAM propagates the generated-belief uncertainty correctly while retaining essentially the same precision as the outcome-side calculation. When the channel is set to zero, the remaining difference contracts toward the numerical finite-sample scale predicted by Equation equation 9.

G.6 REHABILITATION DECISION CHANNELS

Table 8: Decision-component estimates in the rehabilitation cohort.

Channel	Estimate	SE	LR statistic
Intercept α_0	-2.252	0.090	–
Belief α_2	-0.151	0.075	4.17
Normalized time α_3	-1.092	0.258	19.55
Accumulated exposure α_4	12.040	1.098	113.94
Previous-day treatment α_5	0.737	0.168	20.93

The negative belief coefficient shows that a lower inferred untreated recovery slope is associated with a higher treatment probability. The time, exposure, and continuity channels distinguish where the patient is in the rehabilitation course from what the model currently believes about untreated recovery. Together with $\widehat{\beta}(42/60) = 0.088$ (SE 0.033), these estimates turn the treatment process into an interpretable part of the analysis rather than an unreported nuisance mechanism.

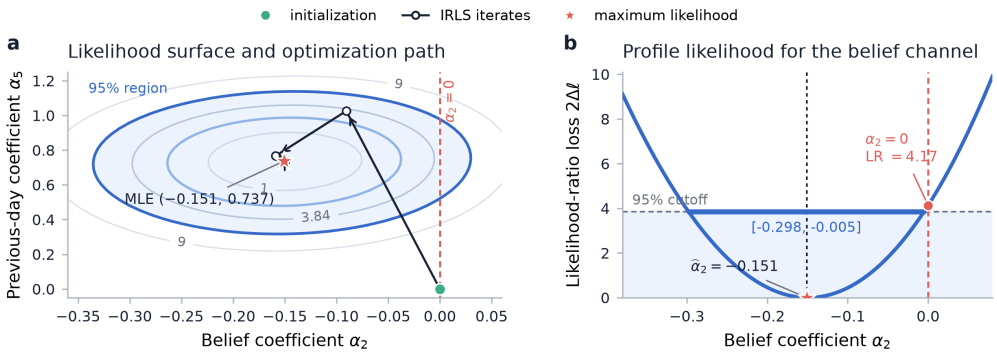


Figure 5: Joint-likelihood geometry for the belief and treatment-continuity channels. The surface and profile views locate the fitted belief coefficient and show how the decision likelihood distinguishes an active belief channel from $\alpha_2 = 0$.

G.7 NEURAL VARIANTS

VI-BAM's learned belief aligns closely with the closed-form belief (Pearson $r = 0.962$; Spearman $r = 0.953$) and yields $\widehat{\beta}(42/60) = 0.075$. A neural outcome-mean variant retains an active belief

channel ($\widehat{\alpha}_2 = -0.163$, LR $p = 0.033$) and yields $\widehat{\beta}(42/60) = 0.049$. These variants demonstrate that BAM's central interface is portable: the belief can be learned and the outcome mean can be nonlinear while the decision interpretation remains intact.

G.8 Cross-cohort MIMIC-III results

Table 9: Previous-day treatment channel in four MIMIC-III demo analyses.

Medication–outcome analysis	$\widehat{\alpha}_5$	SE	LR statistic
Furosemide–creatinine (1)	3.44	0.40	98.4
Furosemide–creatinine (2)	3.58	0.43	93.5
Vasopressin–creatinine	4.18	0.42	156.4
Insulin–glucose	4.82	0.84	54.0

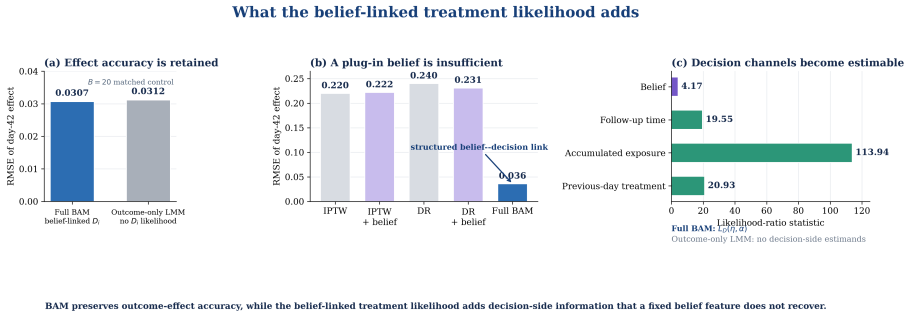


Figure 6: Summary of the belief interface across accuracy, ablation, and decision-channel evidence. The figure separates the value of joint structure from the weaker strategy of treating belief as an ordinary added feature.

The positive continuity coefficient recurs in all four intensive-care analyses. This cross-cohort result supports the portability of BAM's decision-side decomposition: persistence can be measured consistently even when the clinical outcome, medication, and time horizon differ from rehabilitation.

H Estimand Construction and Expanded Derivations

This section makes explicit the calculations that connect the daily treatment record, the intermittently observed outcome process, and the estimands reported in the paper. It is intended to make the statistical interface implementable without changing the model stated in Section 3.

H.1 From Daily Blips to the Observed Outcome Mean

Let $B(s) \in \mathbb{R}^K$ denote the prespecified spline basis and let $\beta \in \mathbb{R}^K$ be its coefficient vector. The causal contribution of treatment administered on day d is $D_i(d)B(d/T)^\top\beta$. Hence the complete treatment-history design at an assessment time t is

$$X_i(t) = \sum_{d \le tT} D_i(d)B(d/T), \qquad X_i(t)^\top\beta = \sum_{d \le tT} D_i(d)\beta(d/T).$$

The first equality is a design construction; the second shows its causal interpretation. A treatment-free patient has $X_i(t) = 0$, whereas two patients assessed on the same day may have different $X_i(t)$ because their prior daily treatment paths differ. No daily outcome is imputed in this construction.

For a one-day intervention at normalized time s, the reported curve is $\beta(s) = B(s)^\top\beta$. For a treatment schedule $\bar{d} = (d_0, \ldots, d_T)$, the corresponding cumulative contrast through time t is

$$\Gamma_{\bar{d}}(t) = \sum_{u \le tT} d_u\, \beta(u/T).$$

Thus the full fitted curve determines any supported regimen contrast. The day-42 quantity $\beta(42/60)$ is a common reporting coordinate, not a second estimand and not a replacement for the full curve.

H.2 Posterior belief without future leakage

At a decision time t, subtract the fixed untreated trajectory, age term, and treatment-history contribution from outcomes observed strictly before t. Stacking these residuals gives $r_i(t)$ and the associated random-effects design gives $G_i(t)$. Under the Gaussian outcome component,

$$r_i(t) = G_i(t)b_i + \varepsilon_i(t), \qquad b_i \sim \mathcal{N}(0, \Sigma), \qquad \varepsilon_i(t) \sim \mathcal{N}(0, \sigma^2 I).$$

Completing the square yields

$$V_i(t) = \left\{\Sigma^{-1} + \sigma^{-2} G_i(t)^\top G_i(t)\right\}^{-1}, \qquad \widehat{m}_i(t) = V_i(t)\sigma^{-2} G_i(t)^\top r_i(t).$$

The slope coordinate $\widehat{m}_{2i}(t)$ is the belief used by the treatment model. The construction is predictable: an outcome recorded at t is excluded when forming the belief for the treatment decision at t and becomes available only for later decisions. A patient with no prior assessment receives the population prior mean, while posterior uncertainty contracts as informative assessments accumulate.

H.3 Why the belief is a joint-model interface

Write η for all outcome-side parameters and α for the decision-side parameters. The patient contribution to the likelihood is

$$L_i(\eta, \alpha) = L_{Y,i}(\eta) \prod_{j \in \mathcal{J}_i} p_{ij}(\eta, \alpha)^{D_{ij}} \{1 - p_{ij}(\eta, \alpha)\}^{1 - D_{ij}},$$

where $\mathcal{J}_i$ indexes assessment-day decisions and

$$\operatorname{logit} p_{ij} = \alpha_0 + \alpha_2 \widehat{m}_{2i}(t_{ij}; \eta) + \alpha_3 t_{ij} + \alpha_4 C_i(t_{ij})/60 + \alpha_5 D_i(d_{ij} - 1).$$

The dependence of $\widehat{m}_{2i}$ on η is the key distinction between BAM and a two-stage regression with a frozen generated covariate. Differentiating the decision log likelihood with respect to an outcome parameter gives

$$\frac{\partial \ell_{D,i}}{\partial \eta} = \sum_{j \in \mathcal{J}_i} (D_{ij} - p_{ij}) \alpha_2 \frac{\partial \widehat{m}_{2i}(t_{ij}; \eta)}{\partial \eta}.$$

This score is exactly zero through the belief path when $\alpha_2 = 0$. When the channel is active, the decision records contribute information about outcome-side parameters through the evolving posterior state, producing the cross-information term used in Theorem 2.

I Fitting Algorithm and Numerical Verification

I.1 Deterministic Fitting Sequence

The implementation follows the same ordered operations in every observed-data fit and simulation replicate.

1. Construct each patient's daily treatment vector, irregular assessment index, normalized time, accumulated exposure, and previous-day treatment indicator.
2. Fit the longitudinal outcome component to obtain stable initial values for fixed effects, spline coefficients, residual variance, and the random-effects covariance.
3. Using only outcomes strictly prior to each decision, compute $r_i(t)$, $G_i(t)$, $V_i(t)$, and $\widehat{m}_{2i}(t)$.
4. Initialize α from the assessment-day logistic regression based on the current belief and recorded decision summaries.
5. Jointly maximize $\ell_Y(\eta) + \ell_D(\eta, \alpha)$ from dispersed starts. Recompute the belief and its derivatives whenever η changes.
6. Retain the converged solution with the largest common unpenalized objective after score, boundary, and information-matrix checks.
7. Form the full observed-information matrix, invert it jointly, and use the delta method for pointwise or cumulative treatment-effect contrasts.

I.2 Checks Applied to Every Retained Fit

Table 10: Numerical and temporal checks for the joint BAM fit.

Check	Interpretation
Objective agreement	Independent starts that reach the same optimum should agree to numerical tolerance.
Score norm	A small gradient verifies first-order stationarity rather than termination caused only by an iteration limit.
Variance boundaries	Residual and random-effects variances must remain inside the admissible parameter space.
Information spectrum	Near-zero or negative eigenvalues flag weak local identification or an unsuitable stationary point.
Belief reconstruction	Recomputed posterior beliefs must match the values used in the retained decision likelihood.
History audit	No outcome at or after a decision time may enter the corresponding belief state.

Smoothness stabilization, when used, ranks candidate curves during fitting but does not redefine the scientific target. The reported covariance is evaluated from the unpenalized likelihood at the retained estimate. This separates numerical regularization from inferential uncertainty.

I.3 Computational Scaling

Posterior updating requires inversion of only the patient-level random-effects matrix, which is 2×2 in the primary implementation. The main cost therefore grows approximately linearly with the number of assessment records once the daily history summaries have been constructed. Daily treatment information is preserved in $X_i(t)$ and $C_i(t)$, but likelihood evaluation does not create artificial daily outcomes. The compact posterior state is therefore computationally useful as well as scientifically interpretable.

J Assumption Diagnostics and Evidence Map

The identifying assumptions are scientific restrictions, not hypotheses that can all be certified from the observed data. Nevertheless, each assumption has observable implications or deliberately designed stress tests. Table 11 separates supporting diagnostics from logical identification claims.

Table 11: Relationship between assumptions and supplementary diagnostics.

Assumption	Empirical check reported	What the check establishes
A1	Verified event ordering and exclusion of post-decision outcomes	Aligns recorded variables with the stated within-day causal order.
A2	Direct latent-selection stress test with $\delta b_{2i}/\tau_2$	Shows the direction and magnitude of deterioration when the boundary is crossed; it does not prove absence of residual confounding in the cohort.
A3	Likelihood-ratio tests, neural variants, and the belief-feature ablation	Tests whether the belief channel carries decision information and whether a fixed feature can reproduce the joint construction.
A4	Treatment frequencies, fitted propensity distribution, and spline-design rank	Detects deterministic treatment regions and unsupported curve directions.
A5	Reuse of the empirical assessment schedule and temporal audit	Preserves the observed irregularity and checks that assessment timing is not accidentally converted into future information.

J.1 Interpreting the A2 Boundary Experiment

At $\delta = 0$, treatment depends on the recorded history summaries and the history-derived belief. Increasing δ adds access to the latent slope itself after conditioning on those quantities. The resulting monotone bias increase and reversal of $\widehat{\alpha}_2$ demonstrate that the restriction in A2 is doing substantive identifying work. This anchors the boundary of BAM: the method addresses sparse measurement through an explicit belief state, but it does not claim to recover effects when clinicians use additional unrecorded prognostic information that bypasses that state.

J.2 Interpreting A3 Separately from the Logistic Implementation

A3 states that the selected summaries are sufficient for the treatment decision among recorded histories. It does not require the decision probability to be logistic-linear. The logistic form is a finite-dimensional working representation that yields channel-specific coefficients and likelihood-ratio comparisons. The nonparametric identification argument uses conditional contrasts and overlap, whereas estimation uses the prespecified logistic index. Consequently, functional-form diagnostics and flexible neural variants assess the implementation without changing the logical content of A3.

K Expanded Reproducibility Map

K.1 Units of Observation and Linkage Keys

The analysis separates three resolutions: patient, calendar day, and assessment. Patient identifiers link baseline age and random-effects structure; calendar-day records define treatment histories; assessment identifiers link the sparse standardized outcomes to the history available immediately before each same-day decision. This separation prevents a common error in which an assessment row is treated as if it represented the entire interval since the previous assessment.

K.2 Required Intermediate Files

Table 12: Intermediate artifacts that connect raw records to reported results.

Artifact	Minimum contents
Patient index	Anonymous patient key, cohort entry, follow-up horizon, admission age.
Daily treatment panel	Patient key, calendar day, binary treatment, previous-day treatment, cumulative exposure.
Assessment panel	Patient key, assessment day, normalized time, standardized outcome, instrument provenance.
Belief audit table	Decision time, number of prior outcomes, posterior mean, posterior variance, latest permissible outcome time.
Fit log	Initialization, objective, convergence code, score norm, information check, retained-solution flag.
Simulation ledger	Seed, scenario, replicate, target effect, estimate, bias contribution, squared-error contribution.
NLP evidence table	Source workbook, row key, text field, matched evidence span, timing, rule/version identifier.

K.3 NLP Processing Safeguards

Text extraction is auxiliary to the primary BAM effect analysis. The processing pipeline retains the source row, normalized text, evidence span, temporal field, and rule version so that every extracted item can be traced back to its record. Negation and historical mentions are handled before aggregation, duplicated exports are reconciled by stable keys, and evidence after the relevant decision time is excluded from decision-history construction. Aggregate counts are produced only after the row-level audit table is frozen. These safeguards prevent the NLP layer from silently creating post-treatment information or changing the primary numerical endpoint.

K.4 How to Reproduce Each Headline Result

The main benchmark is reconstructed from the simulation ledger by grouping on method and computing the signed mean error and root mean squared error at $42/60$. The trajectory and boundary columns use the same code with the scenario identifier changed. The ablation table uses matched replicate identifiers before comparing estimators. The LMM comparison uses the same outcome draws and reports paired differences. Decision-channel statistics are obtained by refitting the joint model with the selected coefficient fixed to zero. The likelihood figure evaluates the retained objective on a fixed grid around the optimum, and the belief summary figure draws only from archived table entries. This one-to-one mapping keeps every figure and table downstream of machine-readable numerical results.

L Scope, Limitations, and Future Directions

BAM is designed for longitudinal settings in which treatment is recorded frequently but outcomes are observed intermittently. Its causal interpretation relies on the assumptions detailed in Appendix B, particularly that the recorded pre-decision history adequately accounts for treatment selection related to the latent untreated trajectory. The current linear-Gaussian belief model and logistic-linear decision model provide an analytically tractable and interpretable implementation of this structure. The primary clinical analysis is based on one rehabilitation cohort; therefore, the estimated effect curve and other numerical results should not be assumed to remain unchanged across populations or clinical settings.

Future methodological work may replace the current components with nonlinear mixed-effects models, Gaussian processes, neural state-space models, or semiparametric decision models while retaining the belief-mediated interface between outcomes and treatment decisions. Sensitivity analyses for residual latent selection could further quantify departures from A2, and extensions could jointly represent assessment timing, treatment intensity, multiple interventions, and heterogeneous patient-level effects.

Future empirical work should evaluate BAM in larger multicenter cohorts and additional treatment settings. The NLP-derived clinical indicators may support analyses of effect heterogeneity across recovery phenotypes and baseline severity, while prospective studies could examine whether belief-linked decision summaries improve treatment monitoring and policy evaluation. These extensions build on BAM's central contribution: using intermittent outcomes to update a belief about untreated recovery and incorporating that belief into joint estimation of treatment decisions and time-varying causal effects.