Mathematical Biosciences and Engineering

https://www.aimspress.com/journal/mbe

MBE, 23(8): 2462–2486.
DOI: 10.3934/mbe.2026089
Received: 02 May 2026
Revised: 06 June 2026
Accepted: 30 June 2026
Published: 11 August 2026

Research article

A latent burden–HVPG mathematical framework for threshold mapping and delay-differential scenario analysis in portal hypertension

M. Fakharany[1,*], Thoraya N. Alharthi[2] and M. A. Elfouly[3]

[1] Department of Mathematics and Statistics, College of Science in Yanbu, Taibah University, Yanbu Governorate, Saudi Arabia

[2] Department of Mathematics, College of Science, University of Bisha, P. O. Box 551, Bisha 61922, Saudi Arabia

[3] Department of Mathematics, Faculty of Science, Mansoura University, Mansoura, Egypt

* **Correspondence:** Email: fakharany@aucegypt.edu.

Abstract: Portal hypertension is clinically defined by invasive hepatic venous pressure gradient (HVPG) thresholds, whereas routine assessment relies on non-invasive markers that are biologically heterogeneous and only indirectly related to portal pressure. We have developed a computational framework designed as a threshold-translation and scenario-exploration tool rather than a predictive or mechanistic model with three explicitly separated components: a data-driven empirical core, a literature-informed delayed dynamic layer, and a synthetic scenario module. Using HVPG-linked datasets, the empirical core has constructed a constrained latent burden score from liver stiffness, inverse platelet count, spleen diameter, inverse albumin, and bilirubin and has compared it with an equal-weight score. The constrained score has improved threshold discrimination more than continuous prediction at clinically relevant HVPG levels. A monotone burden-to-HVPG mapping has translated the latent burden axis into the clinically interpretable HVPG scale and has yielded cohort-specific burden thresholds. Leave-one-source-out (LOSO) technique shows that pooled burden learning is informative but not source-invariant, indicating limited transportability across cohorts. To examine delayed separation between latent burden evolution and visible pressure manifestation, we have introduced a delay differential equation model with literature-informed parameters. The resulting temporal gaps are illustrative scenario outputs and not patient-specific predictions. Overall, the framework provides a transparent, threshold-oriented computational approach for translating heterogeneous biomarkers into interpretable HVPG thresholds while linking data-based burden learning with scenario-based dynamic simulation.

Keywords: portal hypertension; hepatic venous pressure gradient; latent burden; threshold discrimination; delay differential equations

1. Introduction

Portal hypertension is a central determinant of prognosis in chronic liver disease and is clinically defined using hepatic venous pressure gradient (HVPG) thresholds [1–3]. HVPG values ≥ 10 mmHg indicate clinically significant portal hypertension, while values ≥ 12 mmHg are associated with variceal bleeding risk [4, 5]. Despite its clinical importance, HVPG measurement remains invasive, technically demanding, and not routinely available outside specialized centers [3].

Consequently, a wide range of non-invasive markers including liver stiffness measurement (LSM), platelet count, spleen size, and biochemical indices have been proposed as surrogates for portal hypertension [6–8]. However, these markers are biologically heterogeneous and reflect different pathophysiological processes, such as fibrosis, portal flow resistance, and hypersplenism, rather than portal pressure itself [9–11]. This heterogeneity limits their direct interpretability and complicates their integration into a unified clinical framework.

Most existing approaches aim to predict HVPG or classify portal hypertension stages directly from these markers, implicitly treating them as interchangeable proxies. In contrast, we posit that these variables should be interpreted as partial, indirect observations of an underlying latent disease burden. This perspective shifts the modeling objective from direct prediction to structured translation between heterogeneous measurements and clinically meaningful pressure thresholds.

Previous work on portal hypertension has largely followed two complementary but still separated directions. One direction has focused on non-invasive clinical prediction, including liver stiffness, spleen-related measures, platelet-based rules, and composite criteria designed to rule in or rule out clinically significant portal hypertension and variceal risk [12]. A second direction has emphasized pathophysiological and hemodynamic interpretation, in which portal hypertension is viewed as the consequence of progressive intrahepatic resistance, splenic and splanchnic adaptation, and delayed clinical manifestation [13–15].

However, these lines of work are not usually integrated within a single framework that is simultaneously data-anchored, threshold-oriented, and explicit about the inferential limitations of cross-sectional pooled data.

Recent studies emphasize the movement toward non-invasive and threshold-oriented assessment of portal hypertension [16–18]. Recent reviews have highlighted the expanding role of liver stiffness, spleen stiffness, blood-based markers, and composite non-invasive tests for identifying clinically significant portal hypertension and monitoring treatment response, while also stressing that these tools require cautious interpretation in grey-zone cases [19, 20]. In parallel, multicenter and meta-analytic studies have refined Baveno VII-based strategies and incorporated spleen stiffness to improve classification of clinically significant portal hypertension [21, 22]. Imaging-based approaches have also advanced, including CT-derived spleen volume, subharmonic-aided pressure estimation, and multiparametric MRI/MR elastography models for HVPG-related threshold diagnosis [23–25]. These recent contributions support the relevance of threshold-based and multimarker portal-hypertension assessment [26]; however, the present work differs by constructing an explicit latent burden axis, an invertible monotone burden-to-HVPG map, and a delay-differential scenario layer for exploring latent-to-visible separation.

To address this gap, we propose a computational framework with three explicitly separated components: (i) a data-driven empirical layer that constructs a constrained latent burden score, (ii) a monotone

mapping that translates this latent axis into the HVPG scale and associated threshold geometry, and (iii) a delayed dynamic layer that explores temporal separation between latent progression and observed pressure.

Importantly, the dynamic component is not intended to provide patient-specific predictions but to generate interpretable scenario-based insights under explicitly defined assumptions.

By explicitly separating empirical calibration from dynamic simulation, the framework avoids overinterpretation and provides a transparent, threshold-oriented approach to portal hypertension modeling that complements existing clinical tools rather than replacing them.

The remainder of the paper is organized as follows. Section 2 presents the materials and methods, including the empirical data, preprocessing workflow, constrained burden construction, monotone burden-to-HVPG mapping, and mathematical properties of the framework. Section 3 presents the empirical results, including threshold discrimination, transportability analysis, the literature-anchored dynamic layer, and representative synthetic scenarios. Section 4 discusses interpretation and limitations, and Section 5 concludes.

2. Materials and methods

2.1. Data, symbols, and empirical workflow

The empirical core is built directly from three supplementary data Excel workbooks associated with Frankova et al. [9] and two workbooks from Tanaka et al. [10]. These sources have been harmonized into a common master table that has contained 309 rows at the initial pooling stage. Among them, 160 HVPG-positive complete-case rows have been retained for the primary five-marker empirical tier, while 49 rows with available spleen stiffness measurement (SSM) have formed a separate exploratory subset. All analyses in the primary tier are therefore based on complete-case observations, and no imputation is performed.

Table 1. Core symbols and roles.

Symbol	Definition	Unit	Role
B	Latent portal-hypertension burden score	—	Empirical core
$\hat{P}$	Visible pressure state on the HVPG scale	mmHg	Visible layer
$\tilde{P}(B)$	Burden-to-HVPG monotone map (static)	mmHg	Empirical core
X	Latent dynamic burden state	—	Dynamic layer
U	Synthetic burden-forcing input	—	Scenario module
LSM	Liver stiffness measurement	kPa	Observation
PLT	Platelet count	$\times 10^9$/L	Observation
SD	Spleen diameter	cm	Observation
ALB	Serum albumin	g/L	Observation
BILI	Total bilirubin	μmol/L	Observation
τ_1	Observation-to-latent delay	days	Dynamic parameter
τ_2	Latent-to-visible delay	days	Dynamic parameter
a_X	Latent relaxation rate	day^{-1}	Dynamic parameter
a_P	Visible relaxation rate	day^{-1}	Dynamic parameter

The dynamic layer is introduced only after empirical burden learning and burden-to-HVPG mapping has been completed, and it is not calibrated from patient-specific temporal trajectories but instead constructed as a scenario-based layer for illustrative analysis. The present study therefore operates with two empirical analysis tiers and one scenario tier. The primary tier uses five routinely available markers, namely LSM, platelet count (PLT), spleen diameter (SD), albumin (ALB), and bilirubin (BILI). The exploratory tier includes SSM within the smaller subset without redefining the main five-marker burden axis. The scenario tier does not incorporate longitudinal patient data and instead applies analytically defined synthetic forcing functions introduced in Section 3.2.

Core symbols and their roles have been summarized in Table 1, while source-level retention has been summarized in Table 2.

Table 2. Source package and retained cohorts.

Source	All rows	HVPG observed	Primary complete-case rows	SSM rows
[9]	109	109	104	0
[10] S1	130	43	41	36
[10] S2	70	16	15	13

Table 3. HVPG-stage distribution in the primary complete-case cohort.

Source	N	≤ 10	$(10, 12]$	$(12, 16]$	$(16, 20)$	≥ 20
[9]	104	9	12	19	27	37
[10] S1	41	20	5	8	6	2
[10] S2	15	12	0	2	1	0

Table 4. Learned constrained weights and monotone-map parameters.

Parameter	Estimate	95% CI low	95% CI high
w_{LSM}	0.600	0.519	0.600
w_{PLT}	0.108	0.000	0.198
w_{SD}	0.034	0.000	0.230
w_{ALB}	0.108	0.000	0.264
w_{BILI}	0.150	0.000	0.256
$P_{\min}$	4.889	3.001	4.999
A	16.442	14.359	21.995
k	11.096	6.666	13.996
b_0	0.468	0.422	0.552

The 160 HVPG-positive complete-case rows span two source publications with distinct clinical settings (Table 2). Frankova et al. contributes a higher-pressure block (104 rows, predominantly advanced cirrhosis candidates for liver transplantation), while Tanaka et al. contributes two lower-to-moderate pressure blocks (41 and 15 rows) obtained using transient elastography with a novel fusion device. This diversity introduces heterogeneity in disease-stage distribution, HVPG range, and measurement modality. Because the sample is modest for simultaneous learning of five biomarker weights and a four-parameter monotone map, the capped-simplex constraint in (2.3) serves as an

important regularizer but also permits non-unique weight allocations. Table 3 summarizes the source-level stage distribution, and the bootstrap intervals in Table 4 quantify sampling variability. The empirical outputs are therefore interpreted as cohort-specific calibration objects for threshold translation rather than universally generalizable biological constants.

All raw markers have been harmonized at source level before modeling. Robust lower and upper anchors have then been estimated once from the pooled HVPG-positive master rows using the 5^{th} and 95^{th} percentiles, as has been reported in Table 5. These anchors have been fixed across all subsequent analyses to preserve a common latent scale and avoid data leakage between model components.

Burden-increasing variables in the primary empirical tier, namely LSM, SD, and BILI, are normalized so that larger values correspond to worse burden. Burden-decreasing variables, namely PLT and ALB, are inverted so that lower values also correspond to worse burden. In the exploratory SSM subset, spleen stiffness is handled as an additional burden-increasing marker without redefining the main five-marker burden axis. This normalization imposes a unified monotone ordering across heterogeneous biomarkers while preserving interpretability at the level of individual components.

Table 5. Normalization anchors for the five-marker empirical core.

Variable	Lower anchor	Upper anchor	Direction
LSM (kPa)	8.100	55.480	increasing burden
Platelets ($\times 10^9$/L)	47.350	280.150	decreasing burden
Spleen diameter (cm)	8.121	20.000	increasing burden
Albumin (g/L)	22.105	44.000	decreasing burden
Bilirubin (μmol/L)	9.023	160.200	increasing burden

For any burden marker Z_j, the normalized burden component for increasing and decreasing cases is defined by

$$b_j = \begin{cases} \max\left(0, \min\left(1, \frac{Z_j - L_j}{U_j - L_j}\right)\right), & \text{when burden-increase,} \\ \max\left(0, \min\left(1, \frac{U_j - Z_j}{U_j - L_j}\right)\right), & \text{when burden-decrease,} \end{cases} \tag{2.1}$$

where L_j and U_j denote the lower and upper anchors for marker j.

Under this construction, larger values of b_j always indicate greater latent burden, regardless of whether the original raw marker increases or decreases clinically with worsening portal hypertension. Equation (2.1) therefore places all retained routine markers on a common burden-oriented scale.

This representation ensures boundedness ($b_j \in [0, 1]$), monotonicity within anchor intervals, and comparability across markers, which are required for subsequent convex aggregation in the latent burden model.

The empirical burden layer is subsequently built on this normalized representation and carried forward into later sections for constrained burden learning, monotone burden-to-HVPG mapping, and scenario construction.

2.2. Constrained burden construction and burden HVPG mapping

The empirical layer constructs a one-dimensional *latent burden axis* on $[0, 1]$ from five routine markers: LSM, PLT, SD, ALB, and BILI. Because the markers have different physical scales and do not all worsen in the same numerical direction, each raw observation is first converted by (2.1)

into a direction-aware normalized burden component on the unit interval. Their constrained convex combination provides an ordered burden ranking rather than a full probabilistic latent-state model or a generative biological decomposition.

For each retained row i, let

$$\mathbf{b}^{(i)} = (b^{(i)}_{\mathrm{LSM}}, b^{(i)}_{\mathrm{PLT}}, b^{(i)}_{\mathrm{SD}}, b^{(i)}_{\mathrm{ALB}}, b^{(i)}_{\mathrm{BILI}})$$

denote the five normalized burden components. The learned empirical burden score has then been defined as the constrained convex combination

$$B_i = w_{\mathrm{LSM}}\, b^{(i)}_{\mathrm{LSM}} + w_{\mathrm{PLT}}\, b^{(i)}_{\mathrm{PLT}} + w_{\mathrm{SD}}\, b^{(i)}_{\mathrm{SD}} + w_{\mathrm{ALB}}\, b^{(i)}_{\mathrm{ALB}} + w_{\mathrm{BILI}}\, b^{(i)}_{\mathrm{BILI}}, \tag{2.2}$$

subject to

$$w_j \geq 0, \qquad \sum_j w_j = 1, \qquad w_j \leq 0.6. \tag{2.3}$$

The constraint set in (2.3) has been chosen to balance interpretability and stability. The simplex condition preserves interpretability of the burden score as a weighted average on a common normalized scale, whereas the upper cap prevents collapse into a single-marker solution. This cap acts as a regularization mechanism in the presence of correlated biomarkers, although it does not guarantee full identifiability of the weight vector.

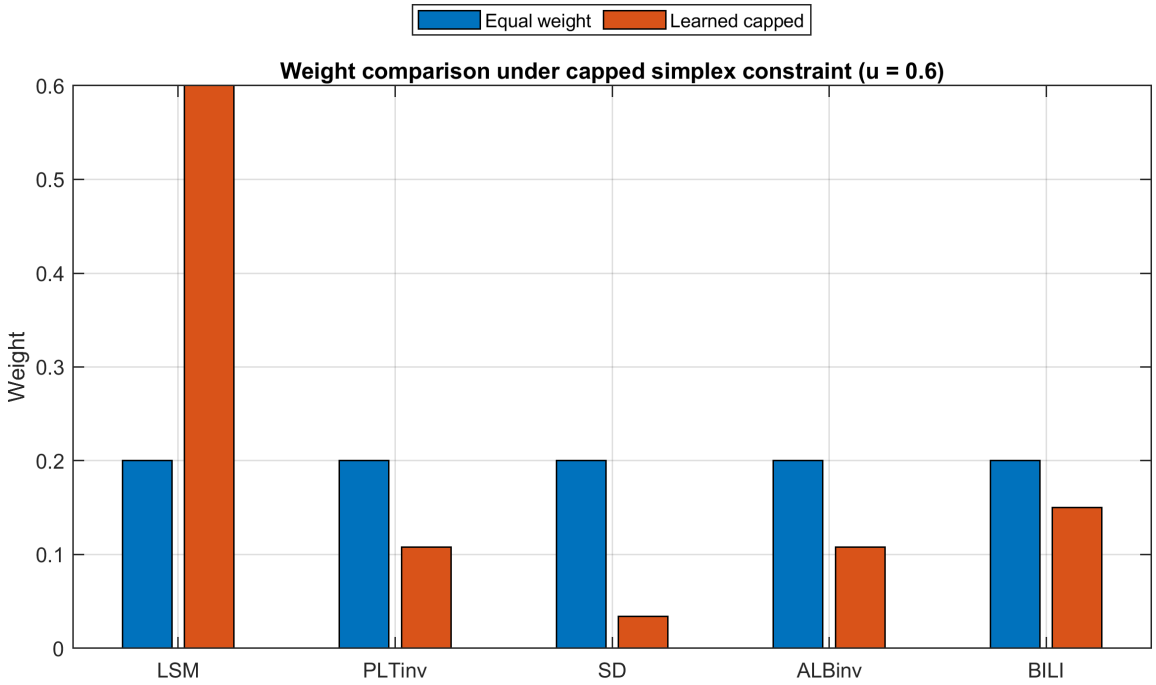

Figure 1. Learned versus equal weights. The constrained empirical fit shifts mass toward LSM while preserving nonzero secondary contributions under the cap 0.6 constraint.

Accordingly, the learned weights should be interpreted as one admissible solution within a constrained class rather than as a unique biological decomposition of portal hypertension burden.

For direct comparison, an equal-weight burden score has also been defined on the same normalized variables:

$$B_i^{\mathrm{eq}} = \frac{1}{5}\sum_{j=1}^{5} b_j^{(i)}. \tag{2.4}$$

The comparison between the learned constrained burden and the equal-weight burden is one of the central empirical tests of the present work because it distinguishes a fully transparent shared-burden construction from a data-informed constrained weighting strategy under the same pooled data framework. The final fitted weights and their bootstrap intervals are reported in Table 4, while the weight distribution relative to the equal-weight comparator is displayed in Figure 1.

Bootstrap intervals have been computed to quantify variability under resampling; however, these intervals reflect sampling variability and do not propagate uncertainty through the full modeling pipeline.

Table 6. Weight-identifiability diagnostics from capped-simplex resampling.

Marker	VIF	Near-optimal median	2.5%	97.5%
LSM	1.311	0.566	0.509	0.598
PLT	1.261	0.080	0.007	0.190
SD	1.295	0.072	0.007	0.182
ALB	1.361	0.132	0.012	0.254
BILI	1.133	0.143	0.038	0.256

Note: Near-optimal weights denote capped-simplex samples within 0.01 mean threshold area under the curve (AUC) of the best sampled vector. The resampling set contained 20,000 random admissible vectors; 272 vectors (1.36%) satisfied this near-optimal criterion.

Identifiability of the weight vector. The capped-simplex constraint set (2.3) defines a compact polytope in $\mathbb{R}^5$. Even though the variance-inflation factors reported in Table 6 do not indicate severe linear multicollinearity, the threshold-based performance objective remains relatively flat over a family of feasible weight vectors. The bootstrap intervals in Table 4 and the near-optimal family in Table 6 and Figure 2 show that alternative allocations among the secondary markers can achieve similar performance. The LSM weight is consistently dominant and frequently reaches the imposed cap, but this should be interpreted as stability under the present regularization and data rather than as formal biological identifiability. Accordingly, the learned weight vector is treated as *one admissible solution* rather than a unique or biologically canonical decomposition. A formal profile-objective or Hessian-based analysis could further characterize the weakly identified directions in future work.

A key observation regarding robustness of the threshold geometry is that the inverse threshold values B^* (Table 7) depend on the map parameters $P_{\min}, A, k, b_0$ through the inversion formula (2.6), and are therefore more sensitive to the monotone-map fit than to the individual weight allocation. Since $P_{\min}$ and A have relatively narrow bootstrap confidence intervals (Table 4), the threshold geometry is more stable than the weight vector, which supports the claim that threshold discrimination is the stronger empirical output of the framework even under the present sample size.

After burden construction, the learned burden axis has been translated to the visible HVPG scale through a monotone logistic map. This map has been used as an empirical burden-to-HVPG translation layer rather than as a fully mechanistic pressure law.

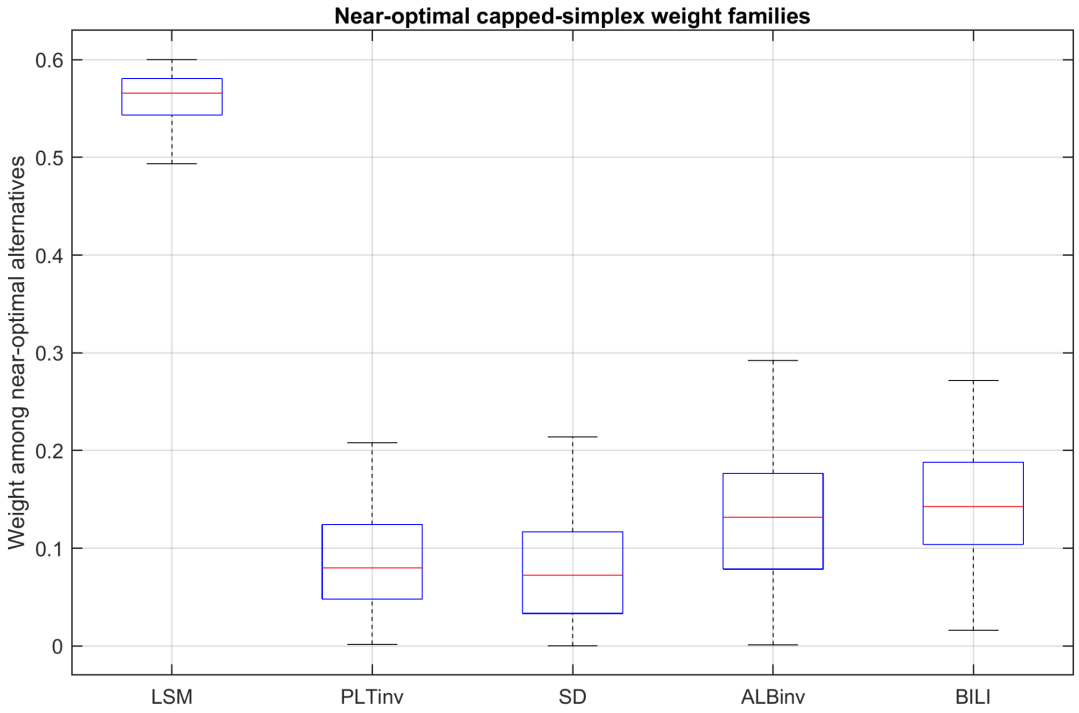

Figure 2. Near-optimal capped-simplex weight families. The LSM contribution remains dominant across the near-optimal family, while secondary-marker allocations vary over wider intervals.

Table 7. Inverse threshold geometry of the learned burden-to-HVPG map.

HVPG threshold (mmHg)	Burden threshold	Local slope	Inside domain
10	0.396	39.084	Yes
12	0.444	44.781	Yes
16	0.534	39.978	Yes
20	0.687	13.579	Yes

The logistic form is chosen for its monotonicity, boundedness, and interpretability as a population-level ranking device; it should not be interpreted as a mechanistic pressure law or as a tool for individual HVPG prediction.

The fitted visible map $\tilde{P}(B)$ has been defined by

$$\tilde{P}(B) = P_{\min} + \frac{A}{1 + \exp(-k(B - b_0))}, \tag{2.5}$$

where $P_{\min}$ denotes the lower visible anchor, A is the map amplitude, k is the transition steepness, and b_0 is the transition location on the burden axis.

Because the map $\tilde{P}(B)$ in (2.5) is strictly monotone in B, it guarantees invertibility and supports consistent translation between latent burden and clinically defined HVPG thresholds.

In the final empirical fit, the map has remained physiologically anchored and has avoided the nonphysical negative baseline that has appeared in earlier unrestricted versions. The fitted parameter

estimates and their bootstrap intervals are summarized in Table 4, and the resulting burden-to-HVPG translation is illustrated in Figure 3.

For any clinically relevant HVPG threshold $\mathcal{P}^*$ satisfying

$$P_{\min} < \mathcal{P}^* < P_{\min} + A,$$

the corresponding burden threshold B^* is obtained by inversion of (2.5):

$$B^* = b_0 - \frac{1}{k}\log\left(\frac{A}{\mathcal{P}^* - P_{\min}} - 1\right). \tag{2.6}$$

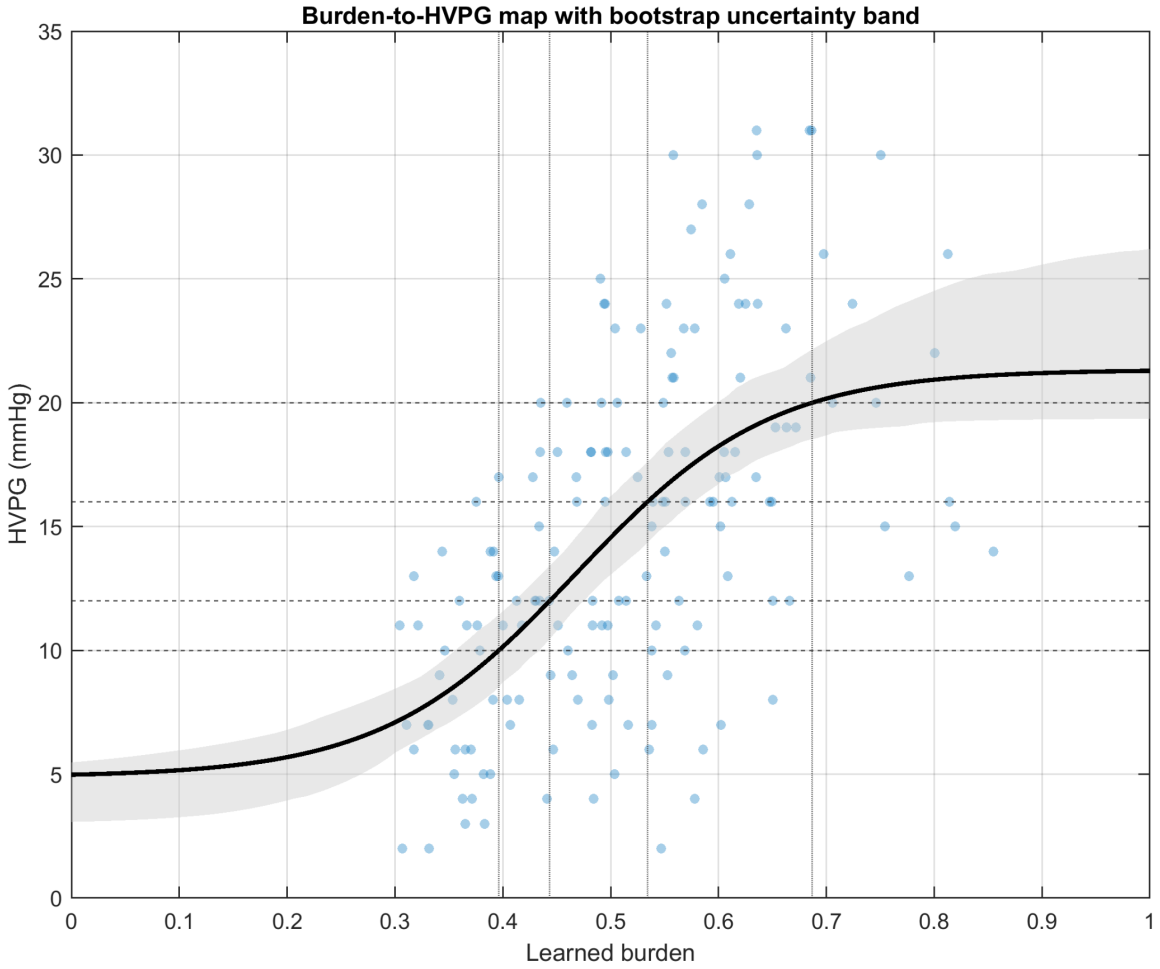

Figure 3. Burden-to-HVPG map $\tilde{P}(B)$. The learned monotone sigmoid map $\tilde{P}(B)$ translates the empirical burden axis into the visible HVPG scale and defines the cohort-specific inverse thresholds used in the scenario module. The smoothness of the fitted curve relative to the data scatter reflects its role as a *population-level smoothing and threshold-ranking device* rather than as a quantitatively predictive physiological relationship; individual predictions carry substantial uncertainty not captured by the fitted curve.

This inverse mapping establishes a one-to-one correspondence between visible pressure thresholds and latent burden levels, which forms the basis for threshold-aligned scenario construction in later sections.

The inverse-threshold geometry is central to the framework because it converts visible HVPG thresholds into latent burden thresholds on a common normalized burden axis. In the empirical core, this translation has been evaluated for the clinically relevant thresholds of 10, 12, 16, and 20 mmHg, yielding cohort-specific burden thresholds rather than universal biological constants. These thresholds are therefore empirical calibration outputs and should not be interpreted as fixed physiological constants across populations.

The resulting inverse thresholds, their local slopes, and their location within the fitted burden span have been summarized in Table 7. These quantities are later reused by the delayed scenario module, where synthetic forcing profiles are defined relative to B^* rather than by arbitrary unlabeled burden targets.

The role of this section is deliberately empirical and structural. It defines a pooled burden axis from heterogeneous routine observations, compares constrained learning against equal weighting, and translates the resulting burden into the visible HVPG scale through the monotone fitted map in (2.5).

Importantly, this section does not establish causality, nor does it provide patient-specific predictions. Instead, it produces calibrated transformation objects for threshold interpretation that are subsequently used in the scenario-based dynamic layer.

2.3. Mathematical analysis

This section formalizes the main structural properties of the framework. The first group of results concerns the empirical burden layer and the fitted burden–HVPG map introduced earlier. These results justify the bounded burden axis, its monotone translation to the HVPG scale, and the inverse-threshold geometry used later in the present study. The second group of results concerns the literature-anchored delayed system used in the representative scenario module. These results justify well-posedness, positive invariance, and convergence under constant forcing, thereby supporting the numerical delayed simulations reported in later tables and figures.

All results in this section are structural and do not rely on patient-specific temporal calibration. They provide theoretical support for the scenario-based dynamic layer rather than predictive claims at the individual level.

Throughout this section, we impose the following standing assumptions:

(A1) The dynamic parameters satisfy $\tau_1 > 0$, $\tau_2 > 0$, $a_X > 0$, and $a_P > 0$.
(A2) The synthetic burden-forcing input U is continuous on $[-\tau_1, \infty)$ and satisfies $0 \le U(t) \le 1$ for all t.
(A3) The initial histories are continuous and satisfy $X(t) \in [0, 1]$ and $\hat{P}(t) \in [P_{\min}, P_{\min} + A]$ for $t \in [-\max\{\tau_1, \tau_2\}, 0]$.
(A4) The fitted burden–HVPG map parameters satisfy $A > 0$ and $k > 0$.

Assumption (A2) reflects the synthetic scenario design rather than observed longitudinal data and constrains the system to operate within the empirically defined burden domain.

We first state the properties of the empirical burden layer.

Proposition 2.1. *For every retained primary-cohort row, each normalized burden component satisfies* $b_j \in [0, 1]$*, and the empirical burden satisfies* $B \in [0, 1]$.

Proof. By (2.1), each normalized component lies in $[0, 1]$. Since B is a convex combination (2.2) and (2.3), it also lies in $[0, 1]$.

Proposition 2.2. *The empirical burden B is monotone nondecreasing in each component b_j. If $w_j > 0$, then B is strictly increasing in b_j.*

Proof. From (2.2),

$$\frac{\partial B}{\partial b_j} = w_j. \tag{2.7}$$

Formally, this derivative corresponds to the continuous extension of the burden function; equivalently, the result can be interpreted using finite differences. Since $w_j \geq 0$, monotonicity follows.

Proposition 2.3. *The fitted map $\tilde{P}(B)$ in* (2.5) *is strictly increasing in B and invertible.*

Proof. Differentiating (2.5),

$$\frac{d\tilde{P}}{dB} = \frac{Ak\, e^{-k(B-b_0)}}{\left(1 + e^{-k(B-b_0)}\right)^2} > 0. \tag{2.8}$$

The monotonicity ensures a one-to-one correspondence between latent burden and visible HVPG values, supporting threshold translation rather than mechanistic interpretation. The following system is not calibrated from patient trajectories but serves as a synthetic scenario module driven by externally specified forcing profiles.

$$X'(t) = a_X(U(t - \tau_1) - X(t)), \tag{2.9}$$

$$\hat{P}'(t) = a_P(\tilde{P}(X(t - \tau_2)) - \hat{P}(t)), \tag{2.10}$$

$$X(t) = \phi_X(t), \qquad \hat{P}(t) = \phi_{\hat{P}}(t), \qquad t \in [-\max\{\tau_1, \tau_2\}, 0]. \tag{2.11}$$

Theorem 2.1. *The delayed system admits a unique continuous solution on any finite interval.*

Proof. The first right-hand side is affine in X, and the second is affine in $\hat{P}$ and contains the smooth logistic function $\tilde{P}(X(t - \tau_2))$. Since $\tilde{P}$ has a bounded derivative, the complete delayed vector field is globally Lipschitz in the present state and delayed-state arguments. Existence and uniqueness on successive finite intervals therefore follow by the method of steps.

Proposition 2.4. *The region*

$$0 \leq X(t) \leq 1, \qquad P_{\min} \leq \hat{P}(t) \leq P_{\min} + A$$

is invariant.

Proof. We show that the rectangular region

$$\mathcal{R} = \{(X, \hat{P}) \in \mathbb{R}^2 : 0 \leq X \leq 1,\ P_{\min} \leq \hat{P} \leq P_{\min} + A\}$$

is positively invariant for the system (2.9) and (2.10) under Assumptions (A1)–(A4).

Invariance of the X-component. From (2.9),

$$X'(t) = a_X(U(t - \tau_1) - X(t)).$$

- If $X(t) = 0$, then $X'(t) = a_X\, U(t - \tau_1) \geq 0$ by (A2), so the boundary $X = 0$ is not crossed downward.
- If $X(t) = 1$, then $X'(t) = a_X\big(U(t - \tau_1) - 1\big) \leq 0$ by (A2), so the boundary $X = 1$ is not crossed upward.

Hence $X(t) \in [0, 1]$ is maintained for all $t \geq 0$ by continuity and the method-of-steps construction.

Invariance of the $\hat{P}$-component. From (2.10),

$$\hat{P}'(t) = a_P(\tilde{P}(X(t - \tau_2)) - \hat{P}(t)).$$

By the invariance of X established above, $X(t - \tau_2) \in [0, 1]$ for all t. Since $\tilde{P}(B)$ is strictly increasing (Proposition 2.3) and $B \in [0, 1]$ (Proposition 2.1), we have

$$P_{\min} \leq \tilde{P}(X(t - \tau_2)) \leq P_{\min} + A, \quad \forall\, t.$$

- If $\hat{P}(t) = P_{\min}$, then $\hat{P}'(t) = a_P\big(\tilde{P}(X(t - \tau_2)) - P_{\min}\big) \geq 0$, so $\hat{P}$ cannot decrease below $P_{\min}$.
- If $\hat{P}(t) = P_{\min} + A$, then $\hat{P}'(t) = a_P\big(\tilde{P}(X(t - \tau_2)) - (P_{\min} + A)\big) \leq 0$, so $\hat{P}$ cannot increase above $P_{\min} + A$.

Hence $\hat{P}(t) \in [P_{\min}, P_{\min} + A]$ is preserved for all $t \geq 0$.

Since both components remain within their respective bounds at the boundary, $\mathcal{R}$ is positively invariant.

Proposition 2.5. *If $U(t) \equiv U_0$ is constant, then $X(t) \to U_0$ and $\hat{P}(t) \to \tilde{P}(U_0)$.*

Proof.

$$X(t) = U_0 + (X(0) - U_0)e^{-a_X t}. \tag{2.12}$$

Define $Y(t) = \hat{P}(t) - \tilde{P}(U_0)$:

$$Y'(t) + a_P Y(t) = a_P\big(\tilde{P}(X(t - \tau_2)) - \tilde{P}(U_0)\big), \tag{2.13}$$

$$Y(t) = e^{-a_P t}Y(0) + a_P \int_0^t e^{-a_P(t-s)}\big(\tilde{P}(X(s - \tau_2)) - \tilde{P}(U_0)\big)\, ds. \tag{2.14}$$

Since $X(t) \to U_0$ and $\tilde{P}$ is continuous, the integrand converges pointwise to zero. Boundedness and exponential decay imply convergence via dominated convergence.

These results establish well-posedness and stability under bounded forcing but do not imply identifiability of dynamic parameters or clinical predictive validity of simulated time scales.

3. Results

3.1. Empirical results

The final empirical core has retained 160 HVPG-positive complete-case rows for the primary five-marker analysis and 49 rows for the exploratory spleen-stiffness subset. Frankova et al. [9] has represented the higher-pressure source block, whereas Tanaka et al. [10] has represented lower-pressure

source environments, making cross-source evaluation a substantive part of the empirical analysis rather than a secondary sensitivity check. All reported results are based on cross-sectional data and therefore reflect statistical associations within the pooled dataset rather than longitudinal or causal relationships.

Within this pooled empirical core, the learned constrained burden has improved on the equal-weight comparator. In the full fit, the root mean square error (RMSE) has decreased from 6.245 to 5.664 mmHg, while the fitted relation has captured the main ordering and threshold structure more clearly than tail-extreme continuous reconstruction.

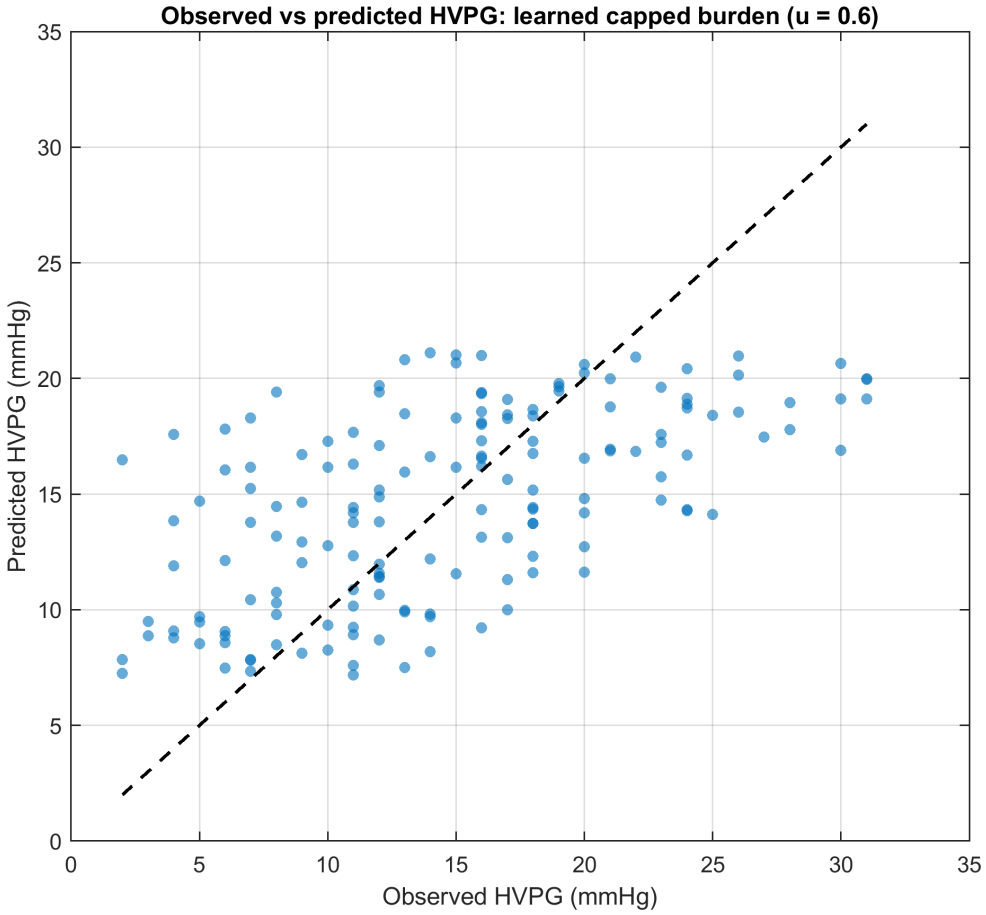

Figure 4. Observed versus predicted HVPG under the learned constrained burden. The monotone empirical core captures ordering and threshold structure more clearly than tail-extreme continuous reconstruction. Note the regression compression toward intermediate HVPG values: The sigmoid map is not a precise continuous HVPG estimator but a monotone threshold-ranking device.

However, this improvement should be interpreted as moderate and primarily structural rather than as evidence of precise continuous HVPG prediction.

Under repeated cross-validation on the same burden representation, the learned constrained burden has remained superior to the equal-weight burden, supporting the claim that the data are informative for constrained weight learning under the imposed cap of 0.6. The corresponding continuous empirical performance measures are summarized in Table 8, and the observed-versus-predicted relation under the learned constrained burden is shown in Figure 4.

The observed-versus-predicted scatter in Figure 4 shows the characteristic *regression compression* of a monotone sigmoid smoother: Predictions are pulled toward the center of the burden range, and both high-HVPG (> 16 mmHg) and low-HVPG (< 8 mmHg) observations are systematically underestimated and overestimated, respectively. This compression is a direct consequence of the ℓ_2-optimal monotone fit and is expected given the logistic functional form. It confirms that the framework is not suited for precise continuous HVPG estimation; its empirical value lies in threshold discrimination and burden ranking, tasks for which the monotone ordering is preserved even when absolute values are compressed.

Table 8. Continuous empirical performance.

Model	RMSE	MAE	R^2	Pearson r
Equal-weight	6.245	4.991	0.202	0.449
Full-fit	5.664	4.623	0.343	0.586
Cross-validation	5.853	4.781	0.299	0.550
LOSO technique	4.839	3.955	0.154	0.621

In Table 8, the equal-weight burden denotes the untrained comparator in which all five normalized burden components contribute equally, with $w_j = 1/5$ for each marker. The full-fit denotes evaluation after fitting the model on all retained complete-case observations. Cross-validation evaluates performance on held-out folds rather than only on the full-fitted dataset. LOSO technique summarizes the average held-out performance across source-level splits. The learned capped burden has been fitted under the capped-simplex constraint $w_j \leq 0.6$.

Table 9. Threshold discrimination.

Endpoint	Model	AUC	Average precision	Sensitivity	Specificity	Best cutoff
HVPG10	Equal burden	0.736	0.885	0.756	0.634	0.469
HVPG10	Learned capped burden	0.771	0.906	0.689	0.707	0.491
HVPG12	Equal burden	0.745	0.832	0.784	0.603	0.476
HVPG12	Learned capped burden	0.805	0.883	0.578	0.879	0.548
HVPG16	Equal burden	0.724	0.623	0.808	0.540	0.483
HVPG16	Learned capped burden	0.800	0.683	0.849	0.632	0.491
HVPG20	Equal burden	0.662	0.338	0.846	0.455	0.483
HVPG20	Learned capped burden	0.789	0.457	0.769	0.711	0.549

The strongest empirical role of the learned burden is threshold-oriented rather than tail-precise continuous reconstruction. Across the clinically relevant endpoints of 10, 12, 16, and 20 mmHg, the learned constrained burden improves on the equal-weight burden in both the area under the receiver operating characteristic curve and average precision, with the clearest gains beyond 12 mmHg. The framework should therefore be interpreted as a threshold-aware latent burden translator and ranking device, not as a precise continuous HVPG replacement. The corresponding results are summarized in Table 9 and Figure 5.

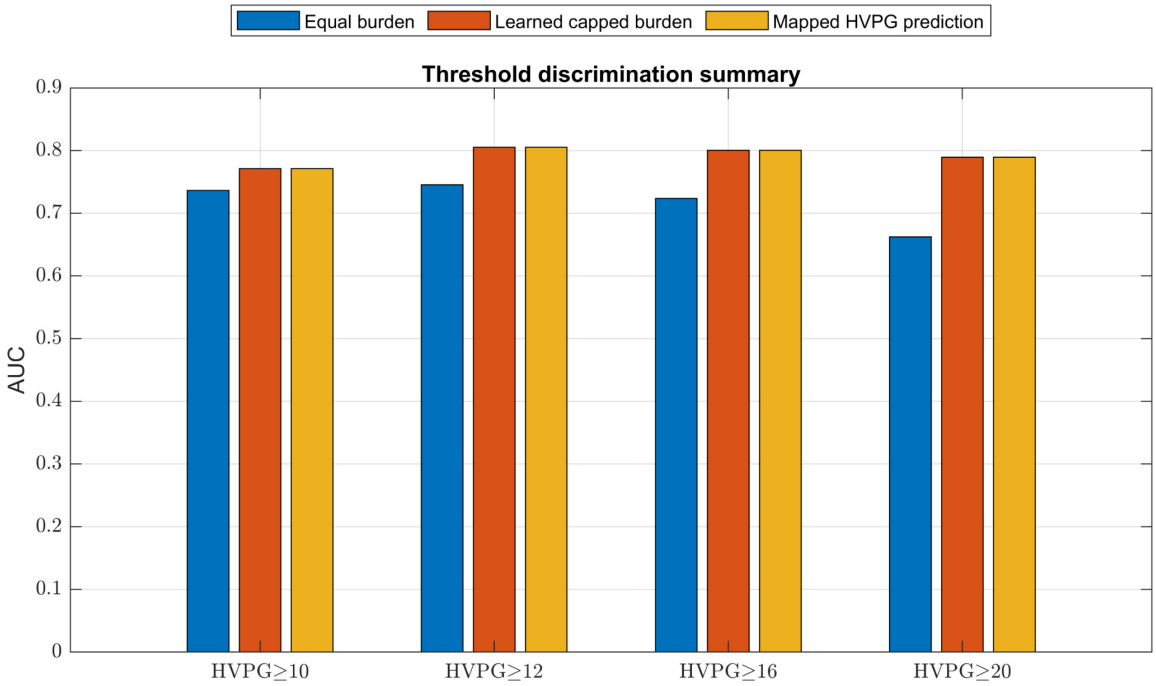

Figure 5. Threshold discrimination across clinically relevant HVPG levels. The learned constrained burden improves over equal weights across all four endpoints, with the clearest gains beyond 12 mmHg.

LOSO technique has shown that pooled burden learning has been informative across sources but has remained heterogeneous in calibration quality. The strongest held-out performance has occurred in [10] S1, whereas the weakest held-out performance has remained in [9], indicating that the pooled burden axis is useful but not source-invariant in its present form.

This heterogeneity reflects underlying differences in cohort composition, disease stage, or measurement protocols, and highlights the need for source-aware interpretation. This heterogeneity is one of the principal findings of the empirical core rather than a peripheral caveat. The results therefore support cautious transport claims, motivate source-aware interpretation, and suggest that future deployment would likely benefit from source-specific recalibration or hierarchical extension rather than direct assumption of invariance. The corresponding held-out continuous transport results are summarized in Table 10, and the source-wise transport comparison is shown in Figure 6.

Table 10. LOSO continuous transport.

Left-out source	Train n	Test n	RMSE	MAE	R^2	Pearson r
[9]	56	104	6.078	4.811	0.030	0.475
[10] S1	119	41	4.174	3.452	0.374	0.769
[10] S2	145	15	4.265	3.602	0.058	0.618

Taken together, the empirical results support a narrower but stronger interpretation of the framework. The learned constrained burden improves on equal weighting within the pooled data setting, especially for threshold discrimination, but it does not erase source heterogeneity and should not be presented as source-invariant.

Accordingly, the primary contribution of this section is the demonstration of threshold-oriented burden learning with explicit transportability auditing, rather than the construction of a universally generalizable predictive model.

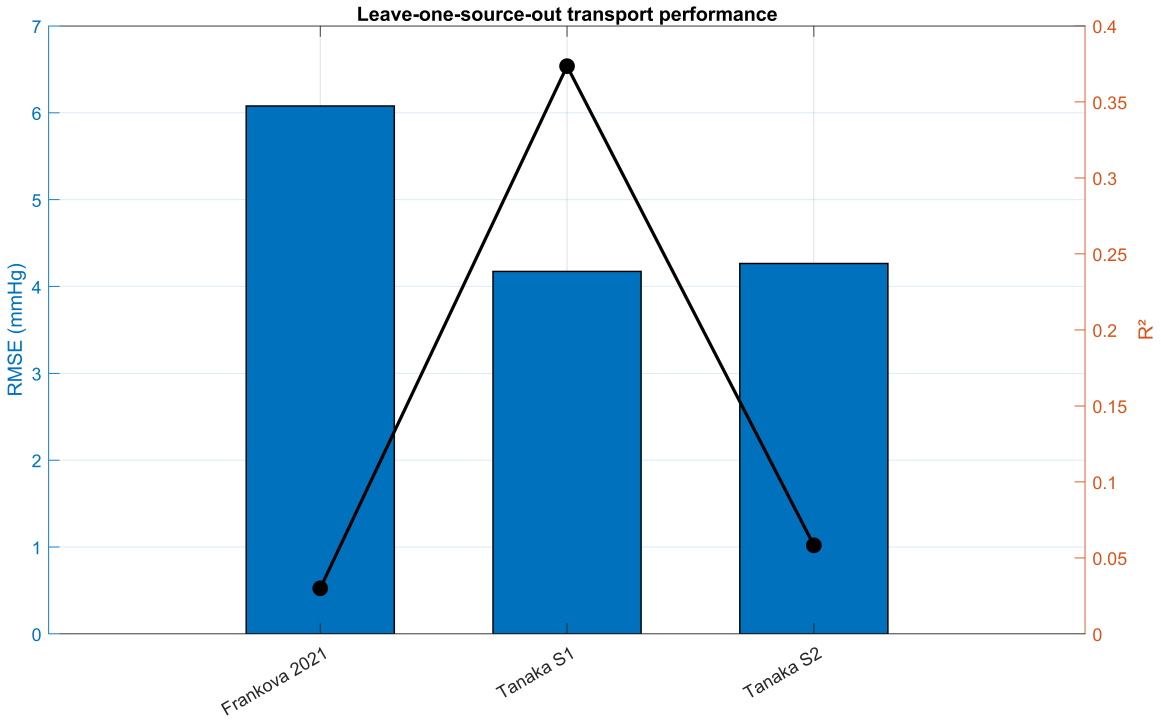

Figure 6. Cross-source transportability. Held-out performance varies materially by source, confirming that the pooled burden should not be presented as source-invariant.

3.2. Literature-anchored dynamic layer and representative synthetic scenarios

The delayed layer is not calibrated from repeated within-patient trajectories. Delay and relaxation parameters are assigned through literature-anchored profiles and explicit scenario rules, while the burden weights and monotone-map parameters remain empirically calibrated quantities from the data core. This deliberate separation avoids assigning patient-specific temporal meaning to quantities that are not identifiable from pooled cross-sectional observations and keeps the framework auditable by distinguishing empirical, literature-anchored, and scenario-imposed parameter blocks. The corresponding provenance structure and dynamic profiles are summarized in Tables 11 and 12.

The short delay τ_1 has been anchored to rapid spleen-stiffness dynamics reported after recompensation-oriented clinical change, where marked decline has already been visible by day three with little additional decline by day five. The slower delay τ_2 has been anchored to the multi-week follow-up logic used in short-term monitoring studies [11, 13, 14].

These delay values should be interpreted as representative ranges rather than as patient-level estimates.

The rates a_X and a_P have not been copied directly from a single study. Instead, they have been obtained analytically from literature-anchored response half-lives, with faster latent adjustment and slower visible adjustment by design.

Table 11. Parameter provenance.

Block	Values / rule	Source or basis
τ_1	2, 3, 5 days	[13, 14]
τ_2	28, 35, 56 days	[11, 13]
a_X	0.3466, 0.2310, 0.1386	[14]
a_P	0.0495, 0.0248, 0.0165	[11, 13]
Burden offsets	Near-threshold, severe, recovery, latent-only, and full-crossing offsets	Scenario design
w_j	$\sum_j w_j = 1$ and $w_j \leq 0.6$	data-empirical core
$P_{\min}, A, k, b_0$	Parameters of $\tilde{P}(B)$: monotone burden-to-HVPG map	data-empirical core

Table 12. Literature-anchored dynamic profiles.

Profile	τ_1 (days)	τ_2 (days)	a_X (day^{-1})	a_P (day^{-1})	Horizon (days)
conservative	2	28	0.3466	0.0495	140
moderate	3	35	0.2310	0.0248	180
high separation	5	56	0.1386	0.0165	220

Accordingly, these parameters define internally consistent dynamic profiles rather than uniquely identified physiological constants.

The scenario module has been built around clinically interpretable burden thresholds rather than around patient labels. The reported scenarios comprise representative near-threshold progression, representative severe progression, representative recovery, representative slow visible normalization, representative latent-only crossing, and representative full crossing.

Each scenario is defined through a prescribed forcing profile rather than fitted to observed longitudinal trajectories.

To summarize these scenarios, the forcing input has been implemented as a profile-dependent function

$$U(t) = U_{\text{base}} + \Delta U\, s(t), \tag{3.1}$$

where $U_{\text{base}} \in [0, 1]$ denotes the baseline burden level, ΔU is a scenario-specific burden offset relative to the inverse threshold B^*, and $s(t)$ is a prescribed synthetic forcing profile.

In the present framework, $s(t)$ is implemented using simple analytic forms (e.g., step, ramp, or smooth sigmoid transitions) depending on the scenario type. These functions are chosen to represent qualitatively distinct progression or recovery patterns rather than to reproduce specific patient trajectories.

In the present framework, $s(t)$ is a bounded scenario function satisfying

$$0 \leq s(t) \leq 1, \tag{3.2}$$

so that the latent forcing remains inside the normalized burden range and can be interpreted relative to the inverse threshold B^* defined by (2.6).

The resulting outputs are therefore scenario-dependent numerical timing summaries under explicitly defined assumptions, not estimates of clinical time-to-event.

The representative synthetic scenario summary is given in Table 13.

Table 13. Representative synthetic scenario summary.

Scenario representative	Mode	HVPG threshold	Profile	Latent crossing (days)	Visible crossing (days)	Gap (days)	Regime
near-threshold progression	progression	12	conservative	24.56	64.67	40.12	2
near-threshold progression	progression	12	moderate	26.33	84.75	58.42	2
near-threshold progression	progression	12	high separation	29.89	121.43	91.54	2
severe progression	progression	20	conservative	24.17	66.95	42.78	2
severe progression	progression	20	moderate	25.76	89.82	64.06	2
severe progression	progression	20	high separation	28.93	128.77	99.84	2
recovery	recovery	12	conservative	25.88	74.80	48.91	2
recovery	recovery	12	moderate	28.32	104.99	76.67	2
recovery	recovery	12	high separation	33.20	151.78	118.57	2
slow visible normalization	recovery	16	conservative	28.17	85.98	57.81	2
slow visible normalization	recovery	16	moderate	31.76	127.05	95.29	2
slow visible normalization	recovery	16	high separation	38.93	185.02	146.09	2
latent-only crossing probe	progression	12	conservative	24.17	62.09	37.92	2
latent-only crossing probe	progression	12	moderate	25.76	—	—	1
latent-only crossing probe	progression	12	high separation	28.93	—	—	1
full crossing	progression	12	conservative	24.17	61.91	37.74	2
full crossing	progression	12	moderate	25.76	79.33	53.58	2
full crossing	progression	12	high separation	28.93	113.25	84.32	2

The dynamic outputs show that the delayed layer behaves coherently under the three profile families. Near-threshold progression and severe progression produced full crossing across all profiles, while the hidden-to-visible gap widened systematically as the profile moved from conservative to high-separation.

This widening gap reflects the imposed delay structure rather than directly measured biological lag.

The clearest qualitative discriminator is the latent-only probe. Under the conservative profile the probe still produced full crossing, whereas under the moderate and high-separation profiles it produced latent threshold crossing without visible threshold crossing inside the simulation horizon.

This result illustrates a possible separation regime between latent burden progression and visible pressure manifestation, rather than establishing its prevalence in clinical populations.

Recovery and slow visible normalization produced the longest separation windows, especially under the high-separation profile.

These longer windows should be interpreted as upper-bound illustrative scenarios under slow visible dynamics, not as expected recovery times in clinical practice.

The same interpretation applies to the delay- and rate-sensitivity surfaces, which summarize how the separation window responds to changes in the delayed layer rather than measuring clinical duration itself.

In particular, the sensitivity analyses highlight the dominant role of τ_2 and a_P in shaping visible delay, which reflects the structural design of the model rather than independent empirical identification.

4. Discussion and limitations

The present study supports a narrower but stronger claim than earlier versions. The available data consist of patient-level measurements pooled across individuals and source cohorts rather than repeated longitudinal measurements from the same patient. This structure is well suited to learning a cross-patient

burden axis, a burden-to-HVPG translation, and clinically relevant threshold geometry, but it does not by itself identify within-patient temporal delay kinetics.

Accordingly, the framework should be interpreted as a structured threshold-translation and scenario-analysis approach rather than as a predictive or mechanistic model of portal hypertension progression.

The empirical and dynamic components should therefore be interpreted in a layered way, with the data core carrying the burden of learning and threshold translation and the delayed layer carrying only literature-anchored scenario structure.

First, the learned burden is informative but not biologically balanced in a literal decomposition sense. Even under the imposed cap of 0.6, LSM remains the dominant contributor, while PLT, ALB, BILI, and SD play smaller and less stable roles.

This dominance reflects both the statistical strength of LSM in the available datasets and the imposed regularization constraints, and it limits interpretation of the burden as a fully distributed physiological construct.

The empirical core should therefore be described as a constrained LSM-led burden rather than as a fully balanced multimarker partition. This interpretation follows directly from the constrained burden weights reported earlier in Table 4 and from the weight distribution shown in Figure 1.

The stability of the learned burden axis under resampling has been assessed via bootstrap (Table 4). The dominant weight w_{LSM} is stable (CI: $[0.519, 0.600]$), while the four secondary weights have wide intervals that include zero. The LOSO technique in Table 10 confirms that calibration quality varies across sources. Together, the bootstrap and near-optimal-family analyses suggest that the *ordering* induced by the learned burden axis is stable, with the framework correctly ranking higher-burden patients near the clinically relevant thresholds, even though the specific weight allocation among secondary markers is not uniquely identified. Claims about the relative importance of PLT, SD, ALB, and BILI should therefore be treated with caution.

The correlation structure among the five normalized burden components creates a non-unique optimization landscape: Multiple weight vectors within the capped-simplex polytope can achieve similar area under the receiver operating characteristic curve and average precision at the clinically relevant HVPG thresholds. This is empirically visible in the wide bootstrap intervals for the secondary weights (Table 4). It implies that the specific numerical values of w_{PLT}, w_{SD}, w_{ALB}, and w_{BILI} should not be read as quantitative estimates of the biological importance of each marker; they represent one admissible solution to the constrained optimization problem under the available data and the imposed regularization. The latent burden axis should therefore be presented as a structured threshold-ranking tool, not as a unique biological decomposition of portal hypertension burden.

Second, the clearest empirical value of the framework lies in threshold interpretation more than in exact continuous HVPG reconstruction. The threshold analyses are stronger and more clinically aligned than the continuous-fit results, and the inverse-threshold map provides an interpretable bridge from burden space to visible HVPG cut points.

In particular, the framework is better suited for classification and discrimination around clinically meaningful thresholds than for precise estimation of absolute HVPG values and should be positioned accordingly in clinical interpretation. This interpretation is supported by the inverse threshold geometry in Table 7, the continuous performance results in Table 8, the threshold discrimination results in Table 9, and the burden-to-HVPG map shown in Figure 3 together with the cross-endpoint comparison in Figure 5.

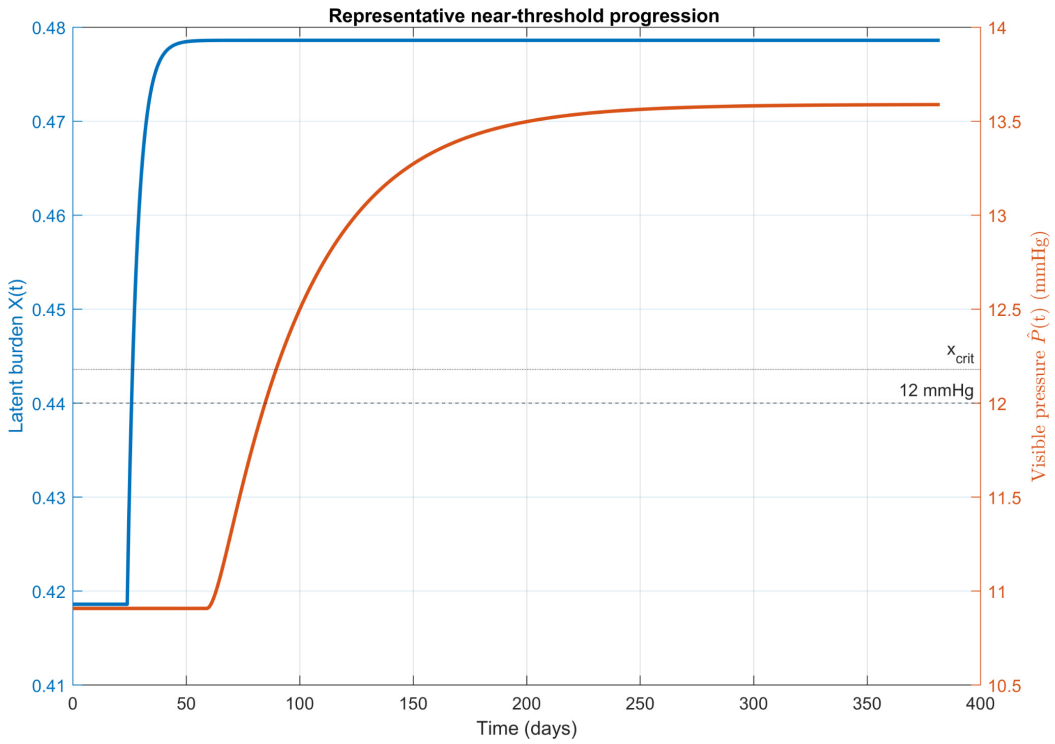

Figure 7. Representative near-threshold progression. The latent burden crosses the burden threshold before the visible HVPG-scale trajectory reaches the corresponding pressure threshold, yielding a profile-dependent hidden-to-visible gap.

Third, transportability heterogeneity is a central empirical finding. The LOSO technique (Table 10 and Figure 6) indicates that the pooled burden axis remains informative across sources but is not source-invariant in calibration quality, most likely because of differences in disease-stage mix, etiology, measurement protocols, and reference HVPG ranges. Application in a new clinical environment would therefore require cohort-specific recalibration of the normalization anchors and monotone-map parameters, and it may benefit from a hierarchical extension. Moreover, LOSO is only an internal portability audit: The present study does not include a prospectively collected independent external cohort. Such external validation, ideally spanning different elastography platforms, etiologies, and disease stages, is required before clinical use can be considered.

Fourth, the literature-anchored delayed layer is useful precisely because it is labeled correctly. It provides representative delayed scenarios for hidden-to-visible separation, latent-only crossing, and delayed recovery, while remaining explicit that these outputs are profile-dependent scenario quantities rather than patient-fitted biological delay estimates.

The simulated time gaps therefore represent internally consistent consequences of the imposed delay structure and forcing profiles, not direct measurements of clinical delay durations or time-to-event.

With richer longitudinal repeated-measurement data from the same patients, future work could move toward more direct dynamic calibration and reduce the need for the present separation between empirical calibration, literature ranges, and scenario-imposed forcing rules.

In the current framework, the most appropriate reading of the dynamic results is therefore as an auditable scenario module anchored to the empirical threshold structure and summarized in Tables 11–Table 13, and Figures 7–9.

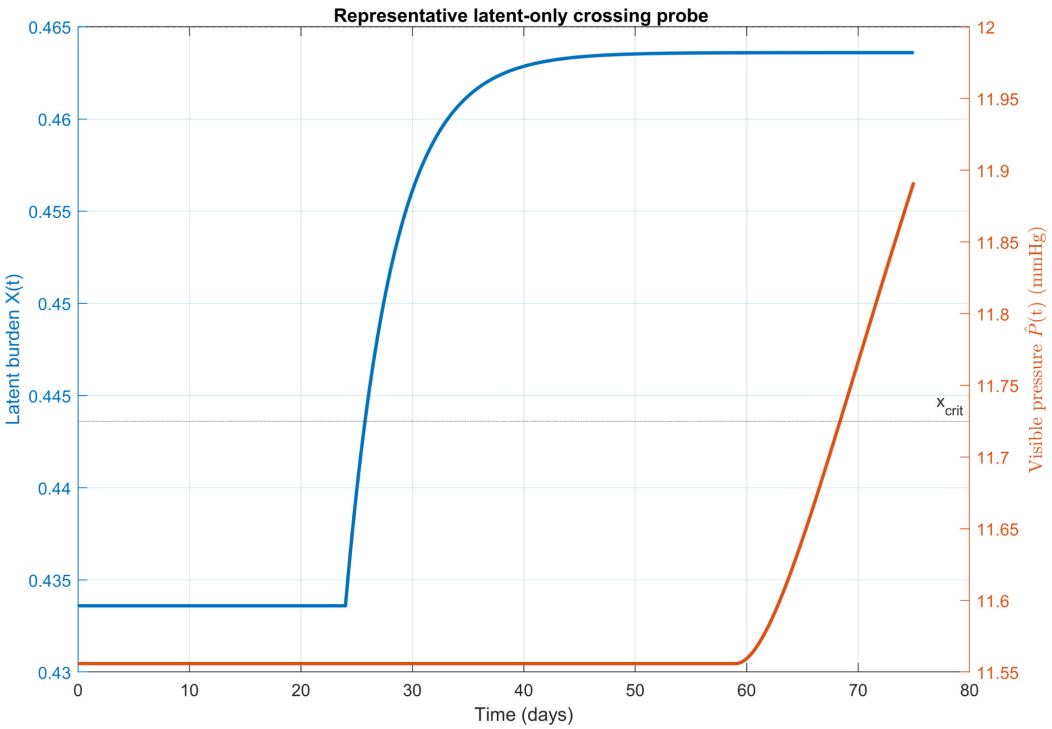

Figure 8. Representative latent-only crossing probe. In the moderate and high-separation profiles, latent threshold crossing occurs without visible threshold crossing inside the simulation horizon.

Regarding the latent burden module: The 160-row dataset is sufficient to demonstrate that constrained weight learning improves over equal weighting in threshold discrimination, but it is too small and heterogeneous to yield a stable unique weight vector with narrow bootstrap intervals. The dominant role of LSM ($w_{\mathrm{LSM}} = 0.6$, at the imposed cap) reflects both its statistical dominance in the available data and the regularization design, not necessarily a biologically universal decomposition.

For the delay-scenario module: Because the dynamic layer is *not* calibrated from patient-specific longitudinal trajectories, dataset limitations do not directly affect the structural correctness of the delay differential equation well-posedness and invariance results (Theorem 2.1, Propositions 2.4 and 2.5). However, the literature-anchored parameter ranges for τ_1, τ_2, a_X, a_P are drawn from small short-term monitoring studies (Tables 11 and 12), and the scenario outputs should accordingly be interpreted as internally consistent illustrative ranges rather than patient-validated timing estimates.

Additional data-handling considerations also apply. The dataset has been constructed by pooling supplementary cohorts, retaining HVPG-positive complete-case rows, and applying fixed normalization anchors prior to burden learning. While this multi-stage approach improves transparency and reproducibility, it may introduce selection bias and limits direct comparability across heterogeneous cohorts.

Transportability limitations further indicate that the framework should not be applied across clinical environments without recalibration. Similarly, the spleen-stiffness subset remains exploratory due to its limited size and should not redefine the main empirical conclusions.

Despite these limitations, the framework provides a transparent and modular structure for linking heterogeneous non-invasive markers to clinically meaningful HVPG thresholds, while offering a controlled environment for exploring latent-to-visible separation under explicitly defined assumptions.

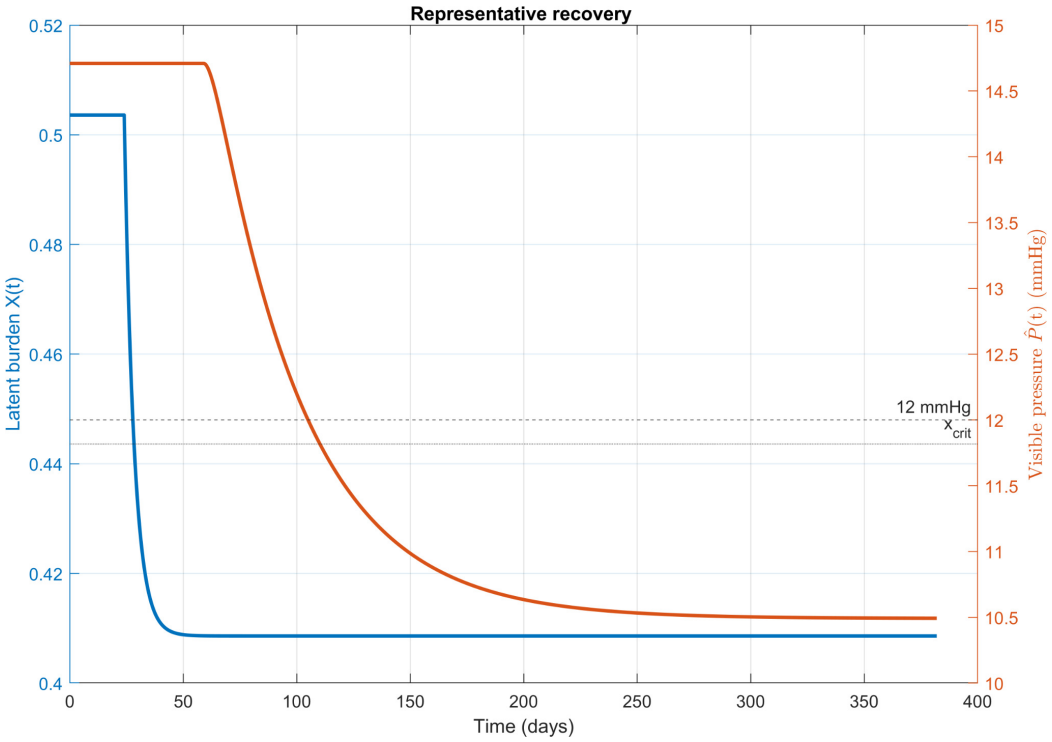

Figure 9. Representative recovery. The latent burden decreases earlier than the visible pressure state normalizes, illustrating delayed visible recovery rather than immediate hemodynamic reversal.

5. Conclusions

The present work supports a data-anchored and methodologically disciplined framework for portal hypertension modeling. The data-based empirical core learns a constrained burden, improves on equal weights, and maps that burden to clinically meaningful HVPG thresholds, while also revealing that cross-source transport remains informative but heterogeneous rather than invariant.

Within this empirical layer, the principal contribution is the construction of a unified latent burden axis and its translation into clinically interpretable threshold geometry, rather than precise continuous prediction of HVPG.

The delayed layer is retained in a clearly limited but useful role: a literature-anchored scenario module that has been solved numerically and has been driven by representative synthetic forcing profiles. It is explicitly not calibrated to patient-level longitudinal data and should be interpreted as a scenario-based exploration of latent-to-visible separation rather than a predictive timing model.

Within that role, it generates near-threshold progression, severe progression, recovery, slow visible normalization, latent-only crossing, and full crossing in a way that is interpretable and internally consistent. These scenarios provide a structured illustration of how delays and relaxation rates can shape observable dynamics without asserting direct correspondence to individual patient trajectories.

Better longitudinal repeated-measurement data from the same patients could support more direct dynamic calibration in future work, but the present study already provides a transparent burden-to-HVPG bridge together with a cautious hidden-to-visible timing framework.

Overall, the framework offers a modular and auditable approach for integrating heterogeneous noninvasive markers into a common threshold-oriented representation while maintaining clear boundaries between empirical calibration, literature-anchored assumptions, and scenario-based interpretation.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article.

Author contributions

Conceptualization, Mohamed El Fakharany; methodology, Mohamed El Fakharany and Mohamed A. Elfouly; formal analysis, Mohamed El Fakharany; investigation, Mohamed El Fakharany, Thoraya N. Alharthi and Mohamed A. Elfouly; writing original draft preparation, Mohamed El Fakharany; writing, review and editing, Mohamed El Fakharany, Thoraya N. Alharthi and Mohamed A. Elfouly. All authors have read and agreed to the published version of the present work.

Data and code availability statement

The MATLAB code is publicly available in the Zenodo repository under DOI: 10.5281/zenodo.19819048. The empirical datasets used in this study were obtained from the published data sources and their supplementary materials described in Section 2.1 and are not redistributed in the Zenodo repository. No identifiable patient-level data were collected by the authors or included in the repository.

Acknowledgments

The authors are thankful to the Deanship of Graduate Studies and Scientific Research at the University of Bisha for supporting this work through the Fast-Track Research Support Program.

Conflict of interest

The authors declare no conflicts of interest.

References

1. Y. Iwakiri, J. Trebicka, Portal hypertension in cirrhosis: Pathophysiological mechanisms and therapy, *JHEP Rep.*, **3** (2021), 100316. https://doi.org/10.1016/j.jhepr.2021.100316
2. R. de Franchis, J. Bosch, G. Garcia-Tsao, T. Reiberger, C. Ripoll, J. G. Abraldes, et al., Baveno VII - renewing consensus in portal hypertension, *J. Hepatol.*, **76** (2022), 959–974. https://doi.org/10.1016/j.jhep.2021.12.022
3. R. Paternostro, W. J. Kwanten, B. S. Hofer, G. Semmler, A. Bagdadi, I. Luzko, et al., Hepatic venous pressure gradient predicts risk of hepatic decompensation and liver-related mortality in patients with MASLD, *J. Hepatol.*, **81** (2024), 827–836. https://doi.org/10.1016/j.jhep.2024.05.033

4. D. E. Kaplan, C. Ripoll, M. Thiele, B. E. Fortune, D. A. Simonetto, G. Garcia-Tsao, et al., AASLD practice guidance on risk stratification and management of portal hypertension and varices in cirrhosis, *Hepatology*, **79** (2024), 1180–1211. https://doi.org/10.1097/HEP.0000000000000647

5. C. Ripoll, R. Groszmann, G. Garcia-Tsao, N. Grace, A. Burroughs, R. Planas, et al., Hepatic venous pressure gradient predicts clinical decompensation in patients with compensated cirrhosis, *Gastroenterology*, **133** (2007), 481–488. https://doi.org/10.1053/j.gastro.2007.05.024

6. A. Berzigotti, S. Seijo, U. Arena, J. G. Abraldes, F. Vizzutti, J. C. García–Pagán, et al., Elastography, spleen size, and platelet count identify portal hypertension in patients with compensated cirrhosis, *Gastroenterology*, **144** (2013), 102–111.e1. https://doi.org/10.1053/j.gastro.2012.10.001

7. J. G. Abraldes, C. Bureau, H. Stefanescu, S. Augustin, M. Ney, H. Blasco, et al., Noninvasive tools and risk of clinically significant portal hypertension and varices in compensated cirrhosis: The ANTICIPATE study, *Hepatology*, **64** (2016), 2173–2184. https://doi.org/10.1002/hep.28824

8. M. Pons, S. Augustin, B. Scheiner, M. Guillaume, M. Rosselli, S. G. Rodrigues, et al., Noninvasive diagnosis of portal hypertension in patients with compensated advanced chronic liver disease, *Am. J. Gastroenterol.*, **116** (2021), 723–732. https://doi.org/10.14309/ajg.0000000000000994

9. S. Frankova, M. Lunova, H. Gottfriedova, R. Senkerikova, M. Neroldova, J. Kovac, et al., Liver stiffness measured by two-dimensional shear-wave elastography predicts hepatic vein pressure gradient at high values in liver transplant candidates with advanced liver cirrhosis, *PLoS One*, **16** (2021), e0244934. https://doi.org/10.1371/journal.pone.0244934

10. T. Tanaka, M. Hirooka, Y. Koizumi, T. Watanabe, O. Yoshida, Y. Tokumoto, et al., Development of a method for measuring spleen stiffness by transient elastography using a new device and ultrasound-fusion method, *PLoS One*, **16** (2021), e0246315. https://doi.org/10.1371/journal.pone.0246315

11. T. Reiberger, The value of liver and spleen stiffness for evaluation of portal hypertension in compensated cirrhosis, *Hepatol. Commun.*, **6** (2022), 950–964. https://doi.org/10.1002/hep4.1855

12. E. Dajti, F. Ravaioli, R. Zykus, P. E. Rautou, L. Elkrief, I. Grgurevic, et al., Accuracy of spleen stiffness measurement for the diagnosis of clinically significant portal hypertension in patients with compensated advanced chronic liver disease: A systematic review and individual patient data meta-analysis, *Lancet Gastroenterol. Hepatol.*, **8** (2023), 816–828. https://doi.org/10.1016/S2468-1253(23)00150-4

13. A. Binzberger, M. Hänle, M. Pfahler, W. Kratzer, T. Seufferlein, E. Zizer, Spleen and liver stiffness evaluation by ARFI imaging: A reliable tool for a short-term monitoring of portal hypertension, *Int. J. Hepatol.*, **2022** (2022), 7384144. https://doi.org/10.1155/2022/7384144

14. D. E. Bozal, J. Hannemann, M. Bürger, G. Allo, E. Dittmann, A. Martin, et al., Spleen stiffness measurements during recompensation in patients with acutely decompensated liver cirrhosis: Preliminary findings of a pilot study, *Front. Med.*, **11** (2024), 1475997. https://doi.org/10.3389/fmed.2024.1475997

15. X. Xu, J. Liu, Y. Zhu, F. Rui, C. Wu, J. Li, Spleen stiffness measurement as a non-invasive assessment in patients with portal hypertension, *Egastroenterology*, **2** (2024), e100031. https://doi.org/10.1136/egastro-2023-100031

16. W. Nasomsong, P. Laukongtam, S. Chirapongsathorn, Clinically significant portal hypertension assessed by spleen stiffness measurement and its associated outcomes among people living with HIV/AIDS, *Sci. Rep.*, **15** (2025), 39952. https://doi.org/10.1038/s41598-025-23635-2
17. W. A. Hassan, R. M. Kemaly, M. A. Sayed, H. M. Nafeh, Liver stiffness measurement, doppler ultrasound, and serum non-invasive markers as predictors of varices and variceal bleeding in cirrhotic patients, *Egypt. Liver J.*, **15** (2025), 78. https://doi.org/10.1186/s43066-025-00483-3
18. W. Lv, L. Yu, L. Wang, S. Tang, J. Yuan, Z. Wu, et al., Development and evaluation of a predictive model based on multi-frequency magnetic resonance elastography for high-risk esophagogastric varices in patients with cirrhotic portal hypertension, *Insights Imaging*, **17** (2026), 69. https://doi.org/10.1186/s13244-026-02246-z
19. M. Mandorfer, J. G. Abraldes, A. Berzigotti, Non-invasive assessment of portal hypertension: Liver stiffness and beyond, *JHEP Rep.*, **7** (2025), 101300. https://doi.org/10.1016/j.jhepr.2024.101300
20. O. S. Rudra, A. Jindal, S. K. Sarin, Use of noninvasive tests to assess clinically significant portal hypertension and response to beta-blockers, *Gastro Hep Adv.*, **4** (2025), 100761. https://doi.org/10.1016/j.gastha.2025.100761
21. C. Liu, J. Li, R. He, I. Grgurevic, Y. Guo, S. Liu, et al., A combined model of spleen stiffness and baveno vii criteria for clinically significant portal hypertension in compensated advanced chronic liver disease: An international multicenter study, *Portal Hypertens. Cirrhosis*, **4** (2025), 4–12. https://doi.org/10.1002/poh2.70004
22. J. Bañares, M. Pons, T. Reiberger, M. Mandorfer, M. Jachs, G. Semmler, et al., Refining the baveno vii criteria for clinically significant portal hypertension: An individual-patient data meta-analysis, *J. Hepatol.*, **84** (2026), 543–553. https://doi.org/10.1016/j.jhep.2025.10.014
23. X. Chen, J. Zhou, Y. Lin, K. Jia, J. Tian, F. Li, et al., CT-derived spleen volume accurately diagnoses severe portal hypertension at HVPG thresholds of 16 mmHg, *Sci. Rep.*, **15** (2025), 16850. https://doi.org/10.1038/s41598-025-02143-3
24. Y. Nakamura, M. Hirooka, R. Yano, N. Kamiyama, T. Oguri, T. Shimamoto, et al., Predictive value of subharmonic-aided pressure estimation for diagnosing clinically significant portal hypertension, *Sci. Rep.*, **15** (2025), 42181. https://doi.org/10.1038/s41598-025-26100-2
25. C. de Margerie-Mellon, F. Joly, A. Tenenhaus, R. Sartoris, V. Paradis, M. Ronot, et al., Preliminary quantitative assessment of liver and spleen MR imaging and elastography in clinically significant portal hypertension, *Sci. Rep.*, **16** (2026), 5430. https://doi.org/10.1038/s41598-025-34761-2
26. L. B. Yang, X. Gao, M. Xu, Y. Li, L. Dong, X. D. Huang, et al., Noninvasive model based on liver and spleen stiffness for predicting clinical decompensation in patients with cirrhosis, *World J. Gastroenterol.*, **31** (2025), 107408. https://doi.org/10.3748/wjg.v31.i33.107408

AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)