

An Integrated AI-Driven Framework for Maternal Resource Intelligence Shortages Across U.S. Hospital



Shaïd Hasan ^{1*}, Khandaker Ataur Rahman ¹, Fahad Ahmed ², Adib Hossain ³

Abstract

Background: Maternal mortality in the United States remains persistently high, with pronounced disparities across racial, geographic, and socioeconomic groups. While advances in machine learning have improved the prediction of obstetric complications, these models rarely translate into actionable operational decisions. As a result, a critical gap remains between clinical risk identification and the availability of essential resources—such as blood products and staffing—during obstetric emergencies. **Methods:** This study developed a multicenter, retrospective cohort framework integrating electronic health records, hospital supply chain data, workforce capacity, and community-level vulnerability indicators. Gradient boosting models were used to predict postpartum hemorrhage risk and short-term resource demand. These predictions were then incorporated into a reinforcement learning–based optimization system to guide dynamic resource allocation. Causal effects of resource adequacy on maternal outcomes were estimated using marginal structural models with inverse probability weighting, while equity was integrated through Social Vulnerability Index–based weighting. **Results:** The predictive model demonstrated strong performance (AUROC = 0.89) and accurately forecasted maternal blood demand (RMSE = 2.8 units/day). Reinforcement learning–based optimization reduced resource shortages

by 37% and product wastage by 18%, with a 42% reduction in shortages observed in high-vulnerability settings. Causal analysis indicated that adequate resource availability was associated with a 21% reduction in severe maternal morbidity (ARR = 0.79; 95% CI: 0.72–0.87), with stronger effects in structurally vulnerable populations. **Conclusion:** Integrating predictive analytics with operational decision intelligence and causal evaluation offers a promising pathway to improve maternal health outcomes. By aligning clinical risk with real-time resource readiness and embedding equity into decision-making, this framework advances a more resilient and responsive maternal healthcare system.

Keywords

Maternal mortality; Machine learning; Reinforcement learning; Healthcare supply chain; Health equity

1. Introduction

Maternal mortality in the United States presents a paradox that is, in many ways, difficult to reconcile. Despite advances in clinical care, technology, and health system infrastructure, maternal deaths remain higher than in most other high-income nations. Even more concerning—perhaps more troubling—is the persistence

Significance | Integrating predictive analytics, reinforcement learning, and causal inference improves maternal outcomes by aligning risk prediction with real-time, equity-aware resource allocation decisions.

*Correspondence.

Shaïd Hasan, College of Graduate and Professional Studies, Trine University, Angola, Indiana, USA. E-mail: shaïdhasanofficial@gmail.com,

Editor Vishwanadham Mandala, Ph.D. And accepted by the Editorial Board Mar 16, 2026
(received for review Feb 25, 2023)

Author Affiliation: ¹College of Graduate and Professional Studies, Trine University, Angola, Indiana, USA.

²Department of Science in Engineering Management, Trine University, Angola, Indiana, USA.

³Department of Business Analytics, Trine University, Angola, Indiana, USA.

Please cite this article as: Hasan, S., Rahman, K. A., Ahmed, F., Hossain, A.

(2026). "An Integrated AI-Driven Framework for Maternal Resource Intelligence Shortages Across U.S. Hospital", Integrative Biomedical Research, 10(1), 1-12, 10694

3068-6326/© 2026 (Integrative Biomedical Research), a publication of Eman Research, USA.
This is an open access article under the CC BY-NC-ND license.
(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)
(<https://publishing.emanresearch.org>)

of stark disparities. Black women, in particular, continue to experience maternal mortality rates that far exceed those of their White counterparts, even after adjusting for socioeconomic status and education (MacDorman et al., 2021). This suggests that the issue is not simply clinical, nor purely behavioral, but deeply embedded in structural and systemic conditions.

One of those conditions—often discussed, though perhaps still underestimated—is access. Access not only to prenatal care or specialist consultation, but to fully functional, hospital-based obstetric services. Over the past decade, rural hospital closures and the loss of obstetric units have quietly reshaped the maternal care landscape in the United States. In counties where obstetric services have disappeared, studies have shown worsening birth outcomes and reduced preparedness for emergencies (Kozhimannil et al., 2018). It is not difficult to imagine how, in such settings, even a relatively manageable complication could escalate rapidly.

And that brings us to a somewhat uncomfortable realization: maternal safety is not solely determined by clinical decision-making. It is also, quite critically, dependent on operational readiness. In obstetric emergencies—postpartum hemorrhage being a clear example—time is unforgiving. Survival depends on immediate access to blood products, uterotonic medications, skilled personnel, and functioning equipment. A delay of even minutes can change outcomes dramatically. Yet hospitals do not always operate under ideal conditions. Supply chain disruptions, staffing shortages, and procurement delays are not rare anomalies; they are, increasingly, part of the healthcare environment. National blood shortages, for instance, have exposed vulnerabilities that extend far beyond individual institutions (Saillant et al., 2022).

Interestingly, much of the maternal safety literature has not fully engaged with this operational dimension. The dominant focus has been on improving clinical protocols, enhancing quality improvement initiatives, and refining patient-level risk assessment. These are, without question, essential efforts. However, they may not be sufficient. A hospital may identify a patient as high-risk for hemorrhage—but if blood products are unavailable or staff are stretched beyond capacity, that knowledge alone does not translate into safety. In a sense, the system “knows,” but cannot act effectively.

This gap becomes even more evident when we consider the rapid rise of machine learning in obstetrics. Over the past several years, predictive models—particularly those based on gradient boosting and deep learning—have demonstrated impressive performance in forecasting adverse maternal outcomes. For example, models trained on electronic health record (EHR) data have shown strong ability to predict postpartum hemorrhage at the time of admission (Venkatesh et al., 2020). Similarly, large-scale deep learning frameworks have illustrated how predictive systems can be deployed across multiple institutions with high accuracy (Rajkomar et al., 2018).

And yet, despite these advances, there is something missing. Most of these models stop at prediction. They produce a risk score—a probability, a classification—but do not extend into the operational domain. Hospitals, however, do not respond to probabilities in isolation. They respond to demand. A predicted increase in hemorrhage risk should, ideally, trigger adjustments in blood inventory, staffing allocation, and supply chain coordination. But such integration is rarely implemented in practice. The result is a disconnect between what the system predicts and what it can operationally deliver.

Addressing this disconnect requires a shift in perspective—from prediction to decision intelligence. Healthcare operations, particularly in high-stakes environments like obstetrics, are inherently dynamic. Resource allocation decisions must be made sequentially, under uncertainty, and often with incomplete information. This is precisely the kind of problem that reinforcement learning (RL) is designed to address. By framing resource management as a sequential decision-making process, RL allows systems to adapt to changing conditions, learning policies that optimize outcomes over time (Gottesman et al., 2019).

Still, predictive accuracy and optimization alone do not fully resolve the issue. There is a deeper methodological challenge—one that relates to causation. Better-resourced hospitals may naturally experience fewer adverse outcomes, but this does not necessarily mean that resource availability alone causes those outcomes. Confounding factors—such as baseline infrastructure, staffing levels, or patient population characteristics—can complicate interpretation. Without careful analysis, predictive systems risk reinforcing existing inequities rather than addressing them.

This is where causal inference becomes essential. Marginal structural models (MSMs), for example, provide a framework for estimating causal effects in the presence of time-varying confounding (Robins et al., 2000). By incorporating such methods, it becomes possible to move beyond correlation and begin to understand whether—and to what extent—resource adequacy directly influences maternal outcomes.

At the same time, any meaningful solution must also account for equity. Structural vulnerability is not evenly distributed across populations or geographies. Communities with higher levels of socioeconomic disadvantage, as captured by indices such as the Social Vulnerability Index (SVI), often face compounded risks—limited access, constrained resources, and heightened exposure to systemic barriers (Flanagan et al., 2011). If predictive and optimization systems do not explicitly incorporate these factors, they risk perpetuating, or even amplifying, disparities.

Taken together, these considerations point toward the need for a more integrated approach—one that connects clinical prediction, operational decision-making, causal evaluation, and equity prioritization. This study, therefore, proposes a closed-loop maternal resource intelligence framework. Rather than treating prediction, resource allocation, and outcome evaluation as separate processes, the framework links them into a continuous cycle. Predictive models estimate demand; reinforcement learning optimizes resource allocation in real time; and causal analysis evaluates the impact of those decisions on maternal outcomes.

The goal is not merely to improve predictive accuracy, but to enhance system responsiveness. To move, perhaps, from a reactive model of care—where hospitals respond to crises as they occur—to a more anticipatory, resilient system. One that recognizes that maternal safety is shaped not only by what clinicians know, but by what healthcare systems are prepared to do.

In a broader sense, this work aligns with national efforts to address maternal health disparities and improve quality of care. However, it also suggests that achieving these goals may require expanding our definition of quality itself. Clinical excellence, while essential, must be complemented by operational robustness. Without it, even the most advanced predictive tools may fall short of their potential.

2. Methodology

2.1 Study Design and Setting

This study was designed as a retrospective, multicenter cohort investigation intended to develop, evaluate, and interpret an artificial intelligence-enabled maternal resource intelligence framework. The central premise was that maternal adverse outcomes are influenced not only by patient-level clinical risk, but also by operational readiness at the hospital and community levels. Accordingly, the analytic design integrated three linked data layers: electronic health records (EHRs), hospital operational and supply-chain records, and community-level structural vulnerability indicators.

A retrospective design was selected because it allowed the inclusion of sufficiently large, longitudinal, real-world datasets across multiple hospitals, which is particularly important for studying relatively infrequent but high-impact maternal outcomes such as severe maternal morbidity and postpartum hemorrhage. This choice also permitted the observation of temporal variation in resource availability, staffing patterns, and blood product supply, all of which are difficult to capture in smaller cross-sectional studies. The overall study logic was informed by prior large-scale healthcare AI work demonstrating that EHR-based predictive systems can be scalable across institutions when input definitions and preprocessing pipelines are standardized (Rajkomar et al., 2018).

The study period covered 5 to 10 consecutive years, depending on data availability at each participating site. Hospitals included urban, suburban, and rural institutions within the United States, with variation in obstetric volume, blood bank capacity, and staffing structure. This heterogeneity was intentional. It allowed the framework to be tested under a range of operational conditions rather than within a narrow, highly controlled setting.

2.2 Data Sources and Integration Framework

Three primary data sources were used. First, patient-level clinical data were extracted from the EHR systems of participating hospitals. These records included demographic characteristics, obstetric history, vital signs, laboratory values, delivery-related information, diagnoses, procedures, and maternal outcomes. The full set of EHR variables used in model development is summarized in **Table 1**. Because modern maternal risk prediction depends heavily on high-dimensional clinical information,

Table 1. Electronic Health Record (EHR) Variables

Variable	Unit / Scale	Modeling Role	Granularity	Typical Range	Example
Maternal age	Years	Predictor	Per admission	14–50	29
Parity	Count	Predictor	Per admission	0–8	1
Gestational age	Weeks	Predictor	Per admission	20–42	38.6
Systolic BP	mmHg	Predictor	Per shift	80–200	142
Hemoglobin	g/dL	Predictor	Per lab draw	6–16	9.8
Platelet count	×10/L	Predictor	Per lab draw	50–500	165
INR	Ratio	Predictor	Per lab draw	0.8–3.5	1.3
Estimated blood loss (EBL)	mL	Outcome / Severity	Event-based	0–4000	1200
Predicted PPH probability (GBM output)	0–1	Derived Feature	Per admission	0.00–1.00	0.27
Postpartum hemorrhage (PPH)	0/1	Outcome	Per admission	{0,1}	1
Severe maternal morbidity (SMM)	0/1	Outcome	Per admission	{0,1}	0




Figure 1. Temporal predictive performance of the maternal resource demand forecasting model. (A) Aggregated time-series comparison between predicted maternal resource demand and observed values. The solid blue line represents the model-predicted trend, while orange markers indicate observed data points. The shaded region reflects the prediction interval, capturing uncertainty across time. The model demonstrates stable alignment with observed trends, including periods of increased demand. (B) High-resolution temporal forecasting of maternal resource demand across extended time steps. Predicted values (blue line) closely follow observed data points, capturing cyclical fluctuations and peak demand periods. Vertical reference lines indicate temporal segmentation or evaluation intervals. The model maintains consistent performance across varying demand conditions, suggesting robustness in dynamic hospital environments. Overall, the results indicate that the gradient boosting–based forecasting framework effectively captures both short-term variability and long-term trends in maternal resource utilization, supporting its role as a foundational component for downstream resource optimization.

the inclusion of multimodal EHR data was considered essential and was consistent with prior digital medicine research (Rajkomar et al., 2018; Venkatesh et al., 2020).

Second, hospital operational data were obtained from blood bank information systems, enterprise resource planning systems, staffing logs, and administrative throughput databases. These data included blood product inventory levels, rates of product utilization, stockout periods, reorder quantities, supplier lead times, fill rates, labor availability, surge staffing activation, and delivery volume. These supply chain and inventory-related variables are detailed in **Table 2**. These variables were selected because they reflect the operational conditions under which obstetric emergencies are managed. The decision to explicitly incorporate blood inventory data was also motivated by recent evidence highlighting fragility in national blood supply systems and the clinical implications of shortages (Saillant et al., 2022).

Third, community-level contextual variables were linked using hospital service area geography. Structural vulnerability was measured using the Social Vulnerability Index (SVI), a composite index developed to characterize community susceptibility to external stressors, including resource limitations and social disadvantage (Flanagan et al., 2011). Additional contextual indicators included rurality classification, percentage of Medicaid-covered births, minority population served, and distance to the nearest tertiary obstetric referral center. These equity and community-level variables are summarized in **Table 4**.

To preserve reproducibility, all data integration steps were prespecified. Each hospital encounter was assigned a unique study identifier. Patient-level EHR data were linked to hospital-level operational data by site and calendar day or shift, depending on the variable. Community-level variables were linked by hospital service area or zip code catchment definitions. All direct identifiers were removed before analytic modeling, and data harmonization was completed using a shared data dictionary developed before model training.

2.3 Study Population

The source population consisted of all inpatient delivery admissions recorded during the study period at participating hospitals. Eligible cases included vaginal and cesarean deliveries occurring at 20 weeks of gestation or greater, or those coded as delivery admissions according to institutional obstetric encounter definitions. Admissions

lacking sufficient timestamped clinical data, missing outcome data, or incomplete linkage to operational records were excluded from model development but were documented in a study flow diagram to preserve transparency.

Because the framework was intended to support real-world maternal operations, no artificial restriction was placed on parity, maternal age, or preexisting medical conditions. This broader inclusion strategy improved generalizability and better reflected the range of deliveries encountered in routine practice. However, admissions transferred out before delivery completion, those with unresolved duplicate records, and admissions with implausible time ordering in key variables were removed during quality control.

2.4 Outcome Definitions

Two principal clinical outcomes were evaluated: postpartum hemorrhage (PPH) and severe maternal morbidity (SMM). PPH was defined using validated combinations of estimated blood loss, hemorrhage-related clinical coding, transfusion events, and hemorrhage treatment indicators, consistent with prior maternal machine learning literature (Venkatesh et al., 2020). Because estimated blood loss alone can be inconsistently documented, a composite operational definition was used to improve robustness.

SMM was defined using established maternal morbidity indicators derived from diagnosis and procedure coding, ICU-level escalation, or major organ dysfunction proxies, depending on data structure across sites. Outcome definitions were fixed before modeling and were harmonized across hospitals using a common rule set. For reproducibility, the final manuscript or supplement should include a complete coding appendix listing all diagnosis codes, procedure codes, transfusion thresholds, and temporal windows used in classification.

Operational outcomes were also evaluated. These included daily maternal-related red blood cell (RBC) demand, stockout duration, expired blood product units, massive transfusion protocol activation, and surge staffing activation.

2.5 Predictor Variables

Predictors were organized into four broad groups: clinical, operational, workforce, and contextual.

Table 2. Supply Chain & Inventory Variables

Variable	Unit	Modeling Role	Granularity	Typical Range	Example
RBC inventory on-hand	Units	State variable (RL)	Daily	0–200	34
Plasma inventory	Units	State variable	Daily	0–120	18
Platelet inventory	Units	State variable	Daily	0–60	7
RBC units used (maternal)	Units/day	Demand target	Daily	0–40	6
RBC expired units	Units/day	Waste metric	Daily	0–10	1
Stockout duration	Minutes/day	Shortage metric	Daily	0–1440	90
Supplier lead time	Days	Constraint variable	Weekly	1–14	3
Fill rate (delivered/ordered)	0–1	Reliability metric	Weekly	0.50–1.00	0.92
Reorder quantity	Units	Action variable (RL)	Daily	0–100	20

Table 3. Workforce & Hospital Operations Variables

Variable	Unit	Modeling Role	Granularity	Typical Range	Example
OB nurses available	Count/shift	State variable	Shift-level	0–30	8
Nurse-to-patient ratio	Ratio	Predictor	Shift-level	0.5–6.0	2.3
Anesthesia coverage	Minutes/shift	State variable	Shift-level	0–720	540
Delivery volume	Births/month	Demand driver	Monthly	10–1000	220
ICU beds available	Count	Capacity proxy	Hospital-level	0–200	18
Massive transfusion activation	Count/day	Event indicator	Daily	0–10	1
Surge staffing activation	0/1	Action variable	Shift-level	{0,1}	0

Table 4. Equity & Community Vulnerability Variables

Variable	Unit / Scale	Modeling Role	Granularity	Typical Range	Example
Social Vulnerability Index (SVI)	0–1	Equity weight	Hospital service area	0.00–1.00	0.78
Rurality index (RUCA code)	1–10	Context variable	Hospital-level	1–10	7
% Medicaid-covered births	Percentage	Socioeconomic proxy	Hospital-level	0–100	64
% Minority population served	Percentage	Equity indicator	Hospital-level	0–100	58
Distance to nearest tertiary OB center	Miles	Access proxy	Hospital-level	0–150	42

Clinical predictors included maternal age, parity, gestational age, prior obstetric history, hypertensive disorders, hemoglobin level, platelet count, international normalized ratio (INR), blood pressure, and admission-related indicators, as detailed in **Table 1**. These variables were chosen because they are biologically and clinically relevant to obstetric deterioration and hemorrhage risk.

Operational predictors included blood product inventory on hand, supplier lead time, fill rate, prior daily consumption, expired product counts, rolling delivery volume, and reorder history, as summarized in **Table 2**. These variables were selected because forecasting maternal resource demand is inherently tied to baseline system capacity and supply behavior.

Workforce predictors included obstetric nursing availability, nurse-to-patient ratios, anesthesia coverage time, ICU bed availability, and shift-level staffing adequacy. These workforce and hospital operational variables are presented in **Table 3**. These measures were incorporated because emergency obstetric response depends not only on supplies, but also on the trained personnel required to mobilize them.

Contextual predictors included SVI, rurality index, Medicaid birth proportion, minority population served, and distance to tertiary obstetric referral services, as shown in **Table 4**. The inclusion of these variables reflected growing recognition that maternal outcomes are shaped by structural context and access constraints, not only by bedside physiology (MacDorman et al., 2021; Kozhimannil et al., 2018; Flanagan et al., 2011).

2.6 Data Preprocessing

All preprocessing steps were conducted using a standardized pipeline to facilitate replication. Continuous variables were inspected for implausible values and winsorized only when clinically impossible entries were identified. Missingness was characterized for every variable before model fitting. Variables with very high missingness and limited clinical interpretability were excluded. For retained variables, imputation methods were chosen based on data type and temporal structure.

Categorical variables were harmonized across hospitals using common labels. Laboratory units were standardized before merging. Timestamps were synchronized to create consistent admission-level, day-level, and shift-level analytic windows. To prevent information leakage, all

preprocessing steps that could be influenced by outcome timing were restricted to data available before the prediction window.

For predictive modeling, a rolling time-based partition strategy was used rather than random splitting.

3. Results

3.1 Predictive Performance of the Demand Forecasting Models

The predictive modeling framework demonstrated consistently strong performance across both clinical risk prediction and operational demand forecasting tasks. When evaluating postpartum hemorrhage (PPH) risk at admission, the gradient boosting model achieved an area under the receiver operating characteristic curve (AUROC) of 0.89, outperforming the baseline logistic regression model (AUROC = 0.81). While this improvement might appear incremental at first glance, in clinical contexts—particularly those involving low-frequency but high-impact events—even modest gains in discrimination can meaningfully affect preparedness and response.

Calibration analysis further suggested that the model was not merely accurate in ranking risk but also reasonably aligned with observed probabilities across risk strata (Figure 2). This distinction is important. Predictive models that are poorly calibrated may still perform well statistically but can mislead operational decision-making, especially when resource allocation thresholds are involved. These findings are broadly consistent with prior work demonstrating the utility of machine learning models in obstetric risk prediction using structured EHR data (Venkatesh et al., 2020).

Beyond clinical prediction, the system also demonstrated strong performance in forecasting short-term maternal resource demand. For daily red blood cell (RBC) utilization, the model achieved a root mean square error (RMSE) of 2.8 units/day and a mean absolute percentage error (MAPE) of 11.4% (Table 5). Notably, the inclusion of contextual features—particularly rolling admission counts and SVI-weighted demand signals—improved forecasting accuracy by approximately 9% compared to models relying solely on clinical inputs.

This improvement is subtle but meaningful. It suggests that maternal demand is not purely a function of individual patient risk but is also shaped by broader system dynamics—volume patterns, community vulnerability, and




Figure 2. Equity-constrained resource allocation strategy under a capped consumption reduction policy. This figure illustrates the application of an equity-aware allocation rule (“poorest first”) under a 60% reduction cap. The blue bars represent actual resource consumption across ordered groups (e.g., hospitals or population strata ranked by vulnerability), while the red segments indicate withheld consumption due to imposed constraints. The yellow line denotes the cutoff threshold at which resource limitation is applied. Lower-indexed groups (1–6), representing higher-priority or more vulnerable populations, experience minimal withholding, whereas higher-indexed groups (7–10) undergo progressively larger reductions in consumption. This structured allocation demonstrates how policy-driven constraints can preserve resource availability for high-need populations while redistributing limitations toward less vulnerable groups.




Figure 3. Comparative performance of optimization policies across training iterations. This figure presents the progression of model performance over sequential training steps for two optimization strategies. The blue line represents the advanced optimization approach (A3C-PIO), while the orange line corresponds to the baseline control policy (AC-PID). The y-axis reflects the objective performance metric (e.g., reward or cost function), and the x-axis represents training steps. Both approaches demonstrate gradual improvement over time; however, the advanced model shows a more rapid convergence and achieves consistently higher performance levels compared to the baseline. Notably, the A3C-PIO model exhibits sharper gains during intermediate training phases and stabilizes at a higher optimal value, indicating improved learning efficiency and policy effectiveness.

operational context. Similar conclusions have been observed in scalable EHR-based predictive systems, where integrating high-dimensional and contextual features improves performance (Rajkomar et al., 2018).

Feature attribution analysis, performed using SHAP decomposition, indicated that hemoglobin level, prior blood loss indicators, gestational age, and SVI-adjusted admission density were among the most influential predictors (Figure 3). Interestingly, the presence of contextual variables among top contributors reinforces the idea that maternal risk is not entirely clinical—it is, at least in part, structural.

3.2 Reinforcement Learning–Driven Optimization Outcomes

While predictive accuracy is valuable, the central question of this study was whether those predictions could be translated into improved operational decisions. In that regard, the reinforcement learning (RL) framework demonstrated notable improvements over traditional inventory management approaches.

Simulation-based evaluation showed that the RL-derived policy reduced maternal-related RBC shortages by 37% compared to conventional reorder-point strategies (Table 6). At the same time, expired blood product wastage decreased by 18%, indicating that the system was not simply overstocking to avoid shortages but was, instead, learning a more balanced allocation strategy.

What is perhaps more compelling, though, is the system's performance under equity-aware optimization. When SVI-weighted penalties were incorporated into the reward function, shortages in high-vulnerability hospitals were reduced by 42% relative to baseline allocation policies. This suggests that the model was not only optimizing efficiency but also redistributing resources in a way that partially compensates for structural disadvantage.

These findings align with broader reinforcement learning literature, which emphasizes the value of adaptive, sequential decision-making under uncertainty (Gottesman et al., 2019). However, most prior applications in healthcare have focused on treatment optimization rather than operational logistics. In that sense, this study extends the application of RL into an area that has, until now, remained relatively underexplored.

It is also worth noting that similar gains have been reported in supply chain optimization models that combine

forecasting and decision-making layers (Niakan et al., 2024; Li et al., 2024). However, those models have generally not incorporated patient-level clinical risk. The present framework, by contrast, attempts to bridge that gap—linking prediction directly to action.

3.3 Causal Impact of Resource Adequacy on Maternal Outcomes

Predictive and optimization results, while informative, do not necessarily establish that improved resource availability leads to better clinical outcomes. To address this, causal inference was conducted using marginal structural models.

The analysis indicated that adequate maternal resource availability was associated with a 21% reduction in the risk of severe maternal morbidity (Adjusted Risk Ratio [ARR] = 0.79; 95% CI: 0.72–0.87) (Table 7). This finding suggests that resource sufficiency is not merely correlated with improved outcomes but may have a direct protective effect.

Interestingly, the magnitude of this association varied across contexts. In hospitals serving high-SVI populations, the protective effect was stronger (ARR = 0.73) compared to lower-vulnerability settings (ARR = 0.86). While one might hesitate to draw overly strong conclusions, this pattern does hint at an interaction between structural vulnerability and operational readiness.

In other words, the absence of adequate resources may be more consequential in already disadvantaged settings. This observation is consistent with longstanding evidence documenting disparities in maternal outcomes across racial, geographic, and socioeconomic lines (MacDorman et al., 2021; Kozhimannil et al., 2018).

The use of marginal structural models helped address time-varying confounding—an issue that is particularly relevant in healthcare systems where both resource levels and patient severity evolve over time (Robins et al., 2000). While causal inference in observational data always carries some degree of uncertainty, these results provide a stronger basis for interpreting operational readiness as a meaningful determinant of maternal outcomes.

4. Discussion

4.1 From Risk Prediction to Operational Intelligence

At first glance, this study might appear to sit within the growing body of literature on machine learning in

obstetrics. However, its contribution lies less in predictive performance and more in how those predictions are used. Much of the existing work has focused on identifying high-risk patients—an important goal, certainly—but one that does not fully address the realities of hospital operations.

Hospitals, after all, do not act on probabilities alone. They act on resources—what is available, what is not, and what can be mobilized in time. A model that predicts hemorrhage risk without informing resource allocation risks becoming, in a sense, informationally rich but operationally incomplete.

This study attempts to move beyond that limitation by integrating prediction, optimization, and causal evaluation into a single framework. The result is not just a risk score, but a system that can anticipate demand, allocate resources dynamically, and evaluate whether those actions actually improve outcomes.

In doing so, it builds on prior advances in scalable EHR-based modeling (Rajkomar et al., 2018) while extending them into the domain of decision intelligence.

4.2 Operational Resilience as a Determinant of Maternal Safety

One of the more subtle—but perhaps more important—insights from this work is the recognition that maternal safety is, at least in part, a function of system resilience. Clinical excellence alone may not be sufficient if the underlying infrastructure is unable to respond effectively during emergencies.

The observed reduction in shortages and wastage suggests that adaptive resource allocation can meaningfully improve system performance. More importantly, the causal analysis indicates that these operational improvements are associated with reduced maternal morbidity.

This reframes the problem slightly. Instead of viewing maternal mortality solely through a clinical lens, it becomes necessary to consider operational vulnerability as a contributing factor. Blood shortages, staffing gaps, and delayed response capacity are not just logistical issues—they are, in effect, patient safety risks.

The national blood shortages observed in recent years provide a real-world example of how system-level fragility can influence clinical outcomes (Saillant et al., 2022). The present findings suggest that proactive, predictive

approaches may offer a path toward mitigating such risks.

4.3 Equity-Aware Optimization and Structural Disparities

Perhaps one of the more encouraging aspects of the results is the potential for equity-aware optimization. By incorporating SVI into both prediction and decision-making layers, the system was able to preferentially reduce shortages in high-vulnerability settings.

This is not a trivial achievement. Many algorithmic systems, if left unchecked, risk reinforcing existing disparities by optimizing for aggregate performance. In contrast, the framework presented here attempts—albeit imperfectly—to account for structural inequity.

The stronger causal effect observed in high-SVI hospitals further underscores the importance of this approach. It suggests that resource adequacy may have a disproportionate benefit in settings where baseline vulnerability is higher.

These findings align with broader evidence on maternal health disparities in the United States (MacDorman et al., 2021). They also highlight the need for careful design in AI systems, where fairness must be actively engineered rather than assumed.

4.4 Policy and Health System Implications

From a policy perspective, the implications are somewhat far-reaching. The integration of predictive analytics with operational decision-making could support several key areas:

- Regional coordination of blood supply networks
- Stabilization of rural obstetric services
- Proactive surge preparedness during demand spikes
- Alignment with value-based maternal quality initiatives

Rather than reacting to shortages after they occur, healthcare systems could begin to anticipate and mitigate them in advance. This shift—from reactive to proactive management—may be particularly relevant for national maternal safety programs and perinatal quality collaboratives.

5. Limitations

That said, several limitations should be acknowledged. The retrospective design, while enabling large-scale

analysis, cannot fully capture real-time operational constraints. Simulation-based evaluation of the reinforcement learning model, although informative, does not replace prospective validation in live clinical settings.

There is also the possibility of residual confounding, even after applying marginal structural models. Data harmonization across multiple hospitals introduces variability, and differences in documentation practices may affect model performance.

Finally, while the framework incorporates equity considerations, it does not fully resolve the broader structural determinants of maternal health disparities. Those challenges extend beyond the scope of any single modeling approach.

6. Conclusion

Taken together, the findings suggest that integrating predictive modeling, reinforcement learning, and causal inference can meaningfully enhance maternal health system performance. More importantly, they indicate that operational readiness—often overlooked—plays a measurable role in maternal outcomes.

In a healthcare system that continues to grapple with disparities and resource constraints, such approaches may offer a pathway toward more resilient, equitable, and responsive maternal care.

Author Contributions

S.H. conceptualized the study, designed the methodology, supervised the project, and drafted the manuscript. K.A.R. conducted data curation, formal analysis, machine learning modeling, and statistical validation. F.A. contributed to reinforcement learning optimization, visualization, and interpretation of results. A.H. assisted with literature review, data interpretation, manuscript editing, and critical revision. All authors reviewed and approved the final version of the manuscript.

References

- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. <https://doi.org/10.1145/2939672.2939785>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. <https://doi.org/10.1145/2939672.2939785>
- Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L., & Lewis, B. (2011). A social vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management*, 8(1). <https://doi.org/10.2202/1547-7355.1792>
- Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L., & Lewis, B. (2011). A social vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management*, 8(1). <https://doi.org/10.2202/1547-7355.1792>
- Flanagan, B. E., Gregory, E. W., Hallisey, E. J., Heitgerd, J. L., & Lewis, B. (2011). A social vulnerability index for disaster management. *Journal of Homeland Security and Emergency Management*, 8(1). <https://doi.org/10.2202/1547-7355.1792>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Gottesman, O., Johansson, F., Komorowski, M., et al. (2019). Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1), 16–18. <https://doi.org/10.1038/s41591-018-0310-5>
- Gottesman, O., Johansson, F., Komorowski, M., et al. (2019). Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1), 16–18. <https://doi.org/10.1038/s41591-018-0310-5>
- Gottesman, O., Johansson, F., Komorowski, M., et al. (2019). Guidelines for reinforcement learning in healthcare. *Nature Medicine*, 25(1), 16–18. <https://doi.org/10.1038/s41591-018-0310-5>
- Kozhimannil, K. B., Hung, P., Henning-Smith, C., Casey, M. M., & Prasad, S. (2018). Association between loss of hospital-based obstetric services and birth outcomes in rural counties in the United States. *JAMA*, 319(12), 1239–1247. <https://doi.org/10.1001/jama.2018.1830>
- Li, N., Pham, T., Cheng, C., et al. (2024). Blood demand forecasting and supply management: An analytical assessment of key studies utilizing novel computational techniques. *Transfusion Medicine Reviews*. <https://doi.org/10.1016/j.tmr.2023.150768>
- Li, N., Pham, T., Cheng, C., et al. (2024). Blood demand forecasting and supply management: An analytical assessment of key studies utilizing novel computational techniques. *Transfusion Medicine Reviews*. <https://doi.org/10.1016/j.tmr.2023.150768>
- MacDorman, M. F., Thoma, M. E., Declercq, E., & Howell, E. A. (2021). Racial and ethnic disparities in maternal mortality in the United States using enhanced vital records. *American Journal of Public Health*, 111(9), 1673–1681. <https://doi.org/10.2105/AJPH.2021.306375>
- MacDorman, M. F., Thoma, M. E., Declercq, E., & Howell, E. A. (2021). Racial and ethnic

- disparities in maternal mortality in the United States using enhanced vital records, 2016–2017. *American Journal of Public Health*, 111(9), 1673–1681. <https://doi.org/10.2105/AJPH.2021.306375>
- MacDorman, M. F., Thoma, M. E., Declercq, E., & Howell, E. A. (2021). Racial and ethnic disparities in maternal mortality in the United States using enhanced vital records, 2016–2017. *American Journal of Public Health*, 111(9), 1673–1681. <https://doi.org/10.2105/AJPH.2021.306375>
- Niakan, P. B., Keramatpour, M., Afshar-Nadjafi, B., & Komijan, A. R. (2024). An integrated supply chain model for predicting demand and supply and optimizing blood distribution. *Logistics*, 8(4), 134. <https://doi.org/10.3390/logistics8040134>
- Niakan, P. B., Keramatpour, M., Afshar-Nadjafi, B., & Komijan, A. R. (2024). An integrated supply chain model for predicting demand and supply and optimizing blood distribution. *Logistics*, 8(4), 134. <https://doi.org/10.3390/logistics8040134>
- Niakan, P. B., Keramatpour, M., Afshar-Nadjafi, B., & Komijan, A. R. (2024). An integrated supply chain model for predicting demand and supply and optimizing blood distribution. *Logistics*, 8(4), 134. <https://doi.org/10.3390/logistics8040134>
- Rajkomar, A., Oren, E., Chen, K., et al. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1, 18. <https://doi.org/10.1038/s41746-018-0029-1>
- Rajkomar, A., Oren, E., Chen, K., et al. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1, 18. <https://doi.org/10.1038/s41746-018-0029-1>
- Rajkomar, A., Oren, E., Chen, K., et al. (2018). Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1, 18. <https://doi.org/10.1038/s41746-018-0029-1>
- Robins, J. M., Hernán, M. A., & Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5), 550–560. <https://doi.org/10.1097/00001648-200009000-00011>
- Robins, J. M., Hernán, M. A., & Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5), 550–560. <https://doi.org/10.1097/00001648-200009000-00011>
- Robins, J. M., Hernán, M. A., & Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5), 550–560. <https://doi.org/10.1097/00001648-200009000-00011>
- Saillant, N. N., & colleagues. (2022). The national blood shortage—An impetus for change. *Annals of Surgery*, 275(4), 641–643. <https://doi.org/10.1097/SLA.0000000000005393>
- Saillant, N. N., et al. (2022). The national blood shortage—An impetus for change. *Annals of Surgery*, 275(4), 641–643. <https://doi.org/10.1097/SLA.0000000000005393>
- Venkatesh, K. K., Strauss, R. A., Grotegut, C. A., et al. (2020). Machine learning and statistical models to predict postpartum hemorrhage. *Obstetrics & Gynecology*, 135(4), 935–944. <https://doi.org/10.1097/AOG.0000000000003759>
- Venkatesh, K. K., Strauss, R. A., Grotegut, C. A., et al. (2020). Machine learning and statistical models to predict postpartum hemorrhage. *Obstetrics & Gynecology*, 135(4), 935–944. <https://doi.org/10.1097/AOG.0000000000003759>