Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

SUPPORTING STRATEGIC HEALTHCARE DECISIONS WITH SIMULATION: A DIGITAL TWIN FOR REDESIGNING TRAUMATOLOGY SERVICES

Marta Cildoz[1,2], Miguel Baigorri[1,2], Isabel Rodrigo-Rincón[3], and Fermin Mallor[1,2]

[1]Institute of Smart Cities, Public University of Navarre, Pamplona, SPAIN
[2]Dept. Statistics, Informatics and Mathematics, Public University of Navarre, Pamplona, SPAIN
[3]University Hospital of Navarre, Pamplona, SPAIN

ABSTRACT

Reducing waiting times in specialized healthcare has become a pressing concern in many countries, particularly in high-demand services such as traumatology. This study introduces a simulation-based approach to support strategic decision-making for redesigning the referral interface between Primary Care and specialized care, as well as reorganizing internal pathways in the Traumatology Service of the University Hospital of Navarre (Spain). A discrete-event simulation model, developed using real patient data and designed to capture the system's transient behavior from its current state, is employed to evaluate the effects of these changes on key performance indicators such as number of consultations per patient, physician workload, and waiting list reduction. The model also evaluates how different referral behaviors among Primary Care physicians influence system performance. Results demonstrate the model's capacity to provide evidence-based guidance for strategic healthcare decisions and highlight its potential to evolve into a digital twin for continuous improvement and operational planning.

1 INTRODUCTION

Waiting lists in specialized healthcare represent a significant challenge in many developed countries, often leading to adverse patient outcomes and increased overall costs. For instance, according to a 2020 OECD study (OECD 2020), 76% of the 34 countries analyzed considered reducing waiting lists a high or medium-high priority. In the region of Navarra (Spain), 32% of the population is currently waiting to receive some form of specialized healthcare (as of 8/03/2025). Various factors contribute to these prolonged waiting times, including high patient demand, limited capacity, inefficient resource allocation, and suboptimal referral and scheduling practices (OECD 2020). Addressing these issues requires a deep understanding of how patients flow through the healthcare system and the underlying causes of delays at different levels of care.

In recent years, a range of interventions has been reported in the medical literature to address these challenges and improve patient referrals from Primary to Specialty Care. These include centralized referral systems, triage protocols, and integrated scheduling platforms (Greenwood-Lee et al. 2018). While some of these interventions have led to measurable improvements at the local level, their impact on the overall performance of the healthcare system has often been limited. One of the main reasons is that healthcare processes are complex and interconnected; interventions designed in isolation may produce unintended effects when interacting with other parts of the system. This organizational complexity makes it difficult to anticipate the impact of changes in any single component. As a result, local improvements may not translate into meaningful reductions in overall waiting times or resource use.

To overcome these limitations, simulation models have been widely adopted in healthcare to support decision-making and improve service delivery. Simulation allows researchers and healthcare managers to build comprehensive representations of healthcare systems and to test proposed interventions under controlled conditions. These models make it possible to observe the systemic effects of changes in resource allocation, scheduling policies, or referral protocols before implementing them in practice. The use of simulation in healthcare has been extensively documented in the literature for many years (Fone et al. 2003; Jun, Jacobson, and Swisher 1999; Brailsford et al. 2009) and it continues to be an essential tool for analyzing complex systems where real-life experimentation is costly or unfeasible (see also recent reviews by Rachuba et al. (2024) and by Wang and Demeulemeester (2023))

Most simulation studies in healthcare have focused on operational or tactical decisions. At the operational level, models are often used to optimize patient flow through clinics or hospitals, or to assign available resources. At the tactical level, they support decisions such as appointment scheduling or capacity planning. However, far fewer studies have addressed strategic-level problems, particularly in the context of integrating different levels of care. Rachuba et al. (2024) emphasize the importance of vertical integration of hospital resources, noting that integrated planning across multiple resources holds even greater potential for improvement than optimizing isolated departments. However, out of 319 reviewed papers on departmental integration, only two addressed strategic decision-making in the context of admissions. One of these papers deals with an appointment system for surgery planning that considers downstream resources (Kianfar and Atighehchian 2023), while the other addresses medical waste planning during the COVID-19 pandemic (Rahiminia et al. 2025). More broadly, strategic decision-making represents a small minority among simulation-based healthcare studies-only 16 out of 125 in the aforementioned review.

An additional challenge in simulation-based healthcare analysis arises when the objective is to study the transient behavior of the system (Garcia-Vicuña et al. 2022). This is often the case when evaluating interventions aimed at improving an ongoing situation, which requires the simulation to start from an accurate representation of the current system state. Building such models demands exhaustive and precise use of healthcare information systems to reconstruct the status of each patient—both those currently in care and those still awaiting attention. When simulation models are linked to real-time data sources and continuously updated, they begin to resemble a digital twin—a concept gaining popularity in healthcare (Elkefi and Asan 2022)—although very few digital twins are currently applied to the management of clinical and medical resources.

In this study, we develop a simulation model to address a strategic-level problem in healthcare management: the redesign of the interface between Primary Care and a hospital's Traumatology Service (TS), as well as the reorganization of internal patient pathways within the department. The proposed intervention involves substantial changes at multiple levels of the healthcare system. These include modifications to how Primary Care physicians refer patients, the creation of new referral and scheduling protocols, changes to the hospital's information systems, and a reallocation of consultation time across different physician roles.

Given the scale and complexity of the proposed redesign, such a decision must be supported by robust evidence demonstrating its effectiveness and efficiency. However, since the system has not yet been implemented, no observational data are available to assess its impact. For this reason, a simulation model was developed and used as a virtual testbed to evaluate the proposed alternative against the current configuration. To do so, the model incorporates real data extracted from hospital information systems and is initialized from the current state of the TS. This allows for the analysis of the system's transient behavior, which is essential in understanding how long it would take to reduce current waiting lists and reach a sustainable steady state under the new design.

Several key performance indicators (KPIs) were defined to quantify outcomes, including the total number of consultations, the physician time required to deliver care, and the time required to eliminate the waiting list backlog. Importantly, the success of the intervention is also dependent on how Primary Care physicians use the new referral channels. The simulation is therefore also used to analyze different patterns of referral behavior and to quantify their impact on system performance.

The main contributions of this study are threefold. First, we present the construction of a simulation model based on multiple fragmented hospital electronic record systems, enabling a faithful reconstruction of patient flows, resource use, and waiting lists. The model has been designed with the structure and data integration capabilities required to evolve into a fully functional digital twin of the TS. Second, we apply this model to address a real-world strategic decision problem: the redesign of the referral interface between Primary Care and specialized hospital care, as well as the reorganization of internal workflows. The study demonstrates that simulation can serve as an effective decision-support tool for evaluating the impact of large-scale organizational changes. Third, we show that the benefits of the proposed redesign are not independent of physician behavior. The model quantifies how variations in the way Primary Care physicians use referral channels significantly influence system performance, highlighting the importance of aligning clinical practice with organizational reforms.

The remainder of this article is structured as follows. Section 2 describes the organizational setting and the proposed redesign of the TS. Section 3 details the data sources and the methodology used to estimate the parameters of the simulation model. Section 4 presents the logic of the simulation and the experimental design. Section 5 reports the simulation results, including a sensitivity analysis and recommendations for health policy-makers. Finally, Section 6 discusses the main findings, outlines the implications for strategic healthcare planning, and identifies directions for future research.

2 SIMULATION FRAMEWORK: CURRENT AND PROPOSED DESIGN FOR PRIMARY CARE INTERFACE AND PATIENT PATHWAYS IN TRAUMATOLOGY SERVICES

Specialized medical departments are complex service units due to the wide variety of patient entry points, the presence of multiple internal subspecialties, and the diverse clinical trajectories that patients may follow. In addition, these departments often interact with other areas of the hospital, such as diagnostic services, surgical units, and rehabilitation facilities. The TS, the focus of this study, is a representative example of such complexity.

Patients may access the TS through several channels: referrals from Emergency Care, other hospital specialties, Hospitalization units, or most commonly, from Primary Care. Among these, Primary Care constitutes the main gateway to the service. There are two distinct routes through which patients are referred from Primary Care: (i) an indirect route via referral requests to specialized e-consultations, where a specialist may either recommend a face-to-face appointment or resolve the case administratively; and (ii) a direct referral to general outpatient consultations. Importantly, Primary Care physicians do not have the authority to schedule appointments directly in the specialized consultation agendas (see Figure 1).

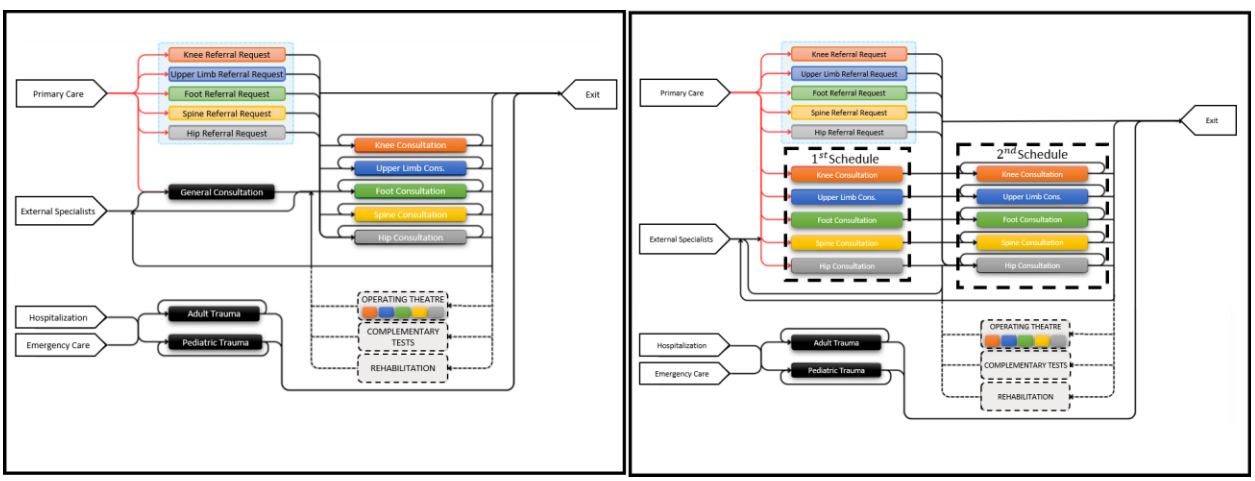

Figure 1: Patient flow in the current design of the Traumatology Service (left) and newly proposed design (right).

Upon entering the system, patients are categorized as either *ordinary* or *preferential*, depending on the clinical urgency and the need for expedited care. After their initial consultation—whether in a general consultation or a specialized consultation (either by direct access or after an e-consultation)—patients may follow various clinical pathways, which often include a sequence of specialized consultations. These pathways may also involve diagnostic procedures, surgical interventions, and rehabilitation sessions, which may take place within the TS or in coordination with other hospital departments. As a result, some patients remain continuously within the TS throughout their treatment, while others exit the system temporarily and re-enter at a later stage— commonly after undergoing diagnostic tests or surgical procedures.

The current organizational model of the TS has revealed several limitations that affect both clinical efficiency and professional satisfaction. One of the main concerns raised by hospital management is the limited clinical value of the initial general consultation, particularly when patients are seen by physicians whose expertise does not align with the patient's condition. This mismatch can delay appropriate care and reduce the overall effectiveness of the clinical pathway. Furthermore, the general consultation unit has become increasingly unattractive for senior physicians, who are reluctant to participate in this type of activity. As a result, this responsibility is often assigned to junior doctors at the beginning of their careers. This practice not only creates a bottleneck in patient access but also undermines the department's ability to attract and retain skilled professionals, especially when competing with other hospitals in the region. These combined factors have motivated the leadership of the TS to consider a structural reorganization that addresses both operational inefficiencies and human resource challenges.

In response to these challenges, hospital management has proposed a strategic redesign of the TS based on two main changes: (i) the elimination of the general consultation unit and (ii) the introduction of a new referral protocol from Primary Care. In the redesigned system, each subspecialty within the TS will maintain its own schedule for initial consultations. All patient entry points—including Primary Care—will be authorized to refer patients directly to these consultations with the relevant specialist. When additional follow-up is required, the same physician will continue treating the patient through their subspecialty consultation schedule, ensuring continuity of care. This approach removes the need for a general consultation, thereby reducing delays and increasing the clinical relevance of the first encounter. It is also expected to enhance physician engagement and improve the department's appeal to new medical professionals (see Figure 1).

The implementation of the proposed redesign involves multiple actors across different levels of the healthcare system. Primary Care physicians will need to adopt new referral practices; hospital administrative units must adapt scheduling protocols; IT services are required to develop and deploy new digital tools; and Traumatology physicians will need to reorganize their agendas to accommodate the new patient flows. Given the scale and complexity of these changes, the proposed system can only be adopted if there is clear evidence that it offers substantial improvements over the current model. However, evaluating its potential impact is not straightforward. The functioning of the TS is intricately linked to other hospital departments and external services, making it difficult to isolate and measure the effects of a structural change. Moreover, a key source of uncertainty lies in how Primary Care physicians will use the new patient entry channel. Their decisions could either relieve or increase the burden on specialist physicians. For instance, redirecting patients who would have been discharged through the e-consultation system toward in-person visits may increase demand and strain capacity. Conversely, bypassing the e-consultation step might reduce specialist workload in some cases, but in others, it may generate redundant visits if important preliminary information is missing. Thus, the overall outcome will depend not only on the design of the system itself but also on how it is used in practice by referring physicians.

The complexity of the TS and the multiple uncertainties associated with the proposed changes, makes difficult to predict the system's response to the new design. In this context, simulation emerges as a powerful tool to replicate the structure and dynamics of the service, allowing for the evaluation of interventions before they are implemented in practice. Specifically, this study employs a digital twin of the TS—a data-driven simulation model that mirrors the real system and can be calibrated with historical and real-time information. This digital twin enables hospital decision-makers to test different implementation scenarios, assess their impact on patient flows and waiting lists, and explore how the behavior of Primary Care physicians influences the system's overall performance.

3 DATA-DRIVEN ESTIMATION OF THE SIMULATION MODEL PARAMETERS

The construction of the simulation model required a detailed estimation of key parameters based on historical data. However, clinical and administrative information in the hospital system is not centrally organized around complete patient care trajectories. Instead, data are distributed across multiple databases, each associated with specific healthcare activities. As a result, reconstructing the full patient pathways through the TS required the integration of several heterogeneous data sources. These included consultation records, electronic referral requests (e-consultations), surgical interventions, and diagnostic procedures. Each dataset contains anonymized patient-level records covering the period from 2018 to 2024, and provides information such as referral origins, appointment and procedure dates, medical specialties involved, and care outcomes. Together, these databases comprise a total of 601,604 entries, offering a comprehensive yet fragmented view of patient care processes. A dedicated data integration process was therefore necessary to extract consistent, structured information (see Table 1) that was used to parameterize the simulation model accurately. Patient trajectories were reconstructed using data from 2018 to 2024, while demand and capacity were estimated using data only from 2022 to 2023 to avoid COVID-19 bias.

Table 1: Data sources and main extracted variables.

Data Source	Main Variables Extracted	Estimation Purpose
Consultation records	Appointment request and execution dates, referral source, attending physician, specialty, visit type.	Arrival rates, pathway reconstruction, consultation durations, capacity by type and physician
E-consultation referrals	Referral request and resolution dates, referring unit, specialty.	Pathway reconstruction, timing of pre-specialist referrals
Clinical tests	Request and execution dates, procedure type, specialty.	Pathway reconstruction and duration of the out of the system period, re-entry into the system
Surgical interventions	Procedure request and execution dates, specialty.	Pathway reconstruction and duration of the out of the system period, re-entry into the system
Rehabilitation	Procedure request and execution dates, specialty.	Pathway reconstruction and duration of the out of the system period, re-entry into the system

Patient Pathways Analysis. First, database integration enabled the reconstruction of 203,624 individual clinical trajectories, tracing the entry point and transitions across different stages of care, including general and specialized consultations, as well as exits and re-entries to the TS for diagnostic tests, surgical procedures, and rehabilitation sessions when applicable. Process mining, carried out using the *pm4py* Python programming library (Berti et al. 2023), revealed over 7,000 distinct care pathways, reflecting the high variability in clinical needs (see Figure 2). These trajectories form the structural foundation of the simulation model, allowing it to reproduce realistic patient flow patterns and to account for the different combinations and sequences of services that patients may undergo during their treatment.

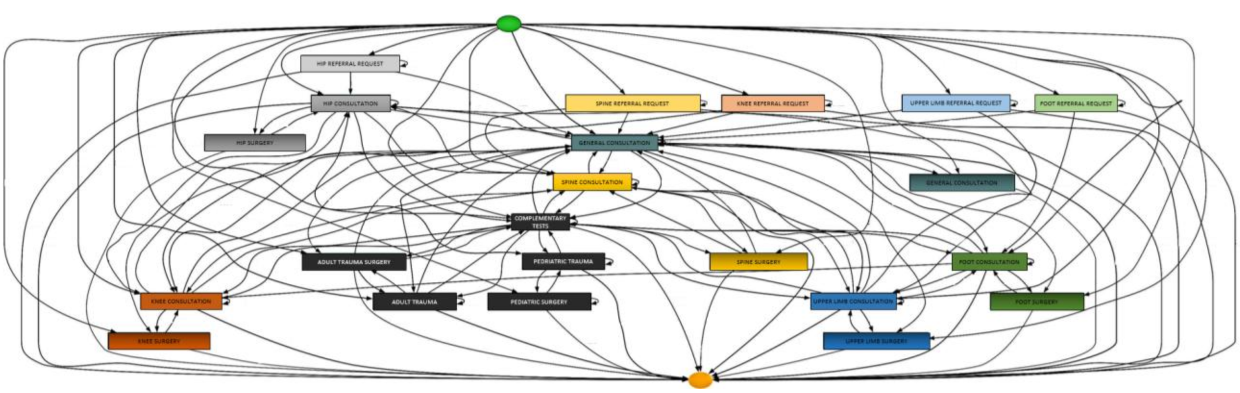

Figure 2: Patient flow chart in the Traumatology Service, extracted using process mining from real data.

Figure 1 shows the theoretical pathways that traumatology patients would follow in an ideal system, under both the current and proposed designs, whereas Figure 2 presents the actual patient trajectories in the current system, extracted from historical data and used in the simulation model.

While most patient trajectories could be directly adapted to the new organizational model, some required additional processing. This was the case for patients who had only attended a general consultation and were discharged without further specialist care. In these situations, there was no explicit information indicating which specialty would have been appropriate under the new system. To address this, a sample of such patients was reviewed and classified by trauma specialists based on their diagnosis into one of five destination specialties. The resulting distribution was then used to assign these cases probabilistically within the simulation of the proposed design.

Patient arrival rate. To estimate the rate at which new patients enter the TS, daily counts of initial care requests were extracted from hospital records over the entire study period. These included referrals to general consultations, first specialized consultations across the five trauma specialties, and e-consultations initiated from Primary Care. All arrival rates were disaggregated by referral source—Primary Care, Emergency Care, Hospitalization, and other specialties—and by patient priority level (ordinary or preferential). Average daily rates were computed by month to account for seasonal variations. The model also considers the effects of public holidays.

Service capacity. The service capacity of the TS was estimated based on consultation slot data from the last two years, which recorded the actual appointment availability for each physician. The Hospital Service of Management, Information, and Evaluation provided the assignment of each physician to their corresponding specialty. These data allowed us to estimate the average number of patients attended per day, disaggregated by patient priority level (ordinary or preferential), consultation type (general or specialized), and visit type (initial or follow-up). The analysis also differentiates between the structural capacity of the service (provided during regular working hours) and temporary capacity increases introduced through reinforcement policies aimed at reducing waiting times. As these reinforcement measures are expected to remain active until the long waiting lists are reduced and the system reaches a steady state, both capacity levels were included in the simulation analysis. Under the proposed reorganization, physicians who currently provide general consultations are reassigned to the first consultation schedules of their respective specialties.

Service times. Consultation times in the TS are predefined according to hospital scheduling protocols. Physicians are allocated 20 minutes for initial patient visits and 10 minutes for follow-up appointments. These fixed durations are used to schedule patients and populate physicians' agendas, and were therefore directly incorporated into the simulation model as deterministic service times for in-person consultations. In contrast, the time required to resolve e-consultations varies depending on the case and the specialty. To capture this variability, appointment timestamp data were analyzed to estimate the probability distributions of time spent on each e-consultation, separately for each specialty.

Initial system state. The simulation model is designed to replicate the operation of the TS starting at its state at a selected point in time. To construct this initial state, hospital databases (described in the previous section) were used to identify, for a given reference date, the number of patients waiting for each type of consultation, along with information about any previous care received. Additionally, the number of pending e-consultations per specialty was determined, as well as the patients who had temporarily exited the service to undergo diagnostic tests, surgical procedures, or rehabilitation sessions but were expected to return for further care. Together, these data provide a detailed snapshot of the system's status at any chosen start date and allow the simulation to begin from a realistic and data-informed configuration. When connected to real-time data sources, the model can accurately mirror the current state of the system, functioning as a digital shadow of the TS.

4 SIMULATION MODEL AND EXPERIMENTAL DESIGN

Discrete-event simulation model. The simulation model follows a discrete-event framework with a daily time-step, advancing the simulation clock one day at a time. The state of the system is described by a set of demand- and capacity-related state variables. On the demand side, the model tracks the number of patients waiting for each type of consultation—either general or specialized—across all five specialties, as well as the number of e-consultations pending response. Each waiting patient is represented individually, with a unique identifier and a record of the trajectory completed so far within the TS. On the capacity side, the model tracks the number of physicians from each specialty assigned to attend each type of consultation and to resolve e-consultations.

At each simulated day, new patients are added to the appropriate waiting lists based on Poisson-distributed arrival processes, using the daily rates estimated in the previous section. Each new patient's attributes and clinical trajectory are sampled from the historical database of treated patients. The available service capacity for each type of consultation is also calculated daily, and patients are removed from the corresponding waiting lists according to availability. Attended patients are then routed based on the next step in their assigned trajectory: they may be placed on the waiting list for another consultation, discharged permanently from the service, or temporarily leave the TS to undergo diagnostic tests, surgery, or rehabilitation. In the latter case, each external activity has an associated duration (as defined in the patient's trajectory), after which the patient automatically re-enters the TS and is placed in the waiting list of the appropriate follow-up consultation.

Simulation objectives and key performance indicators. The primary objective of the simulation model is to assess the performance of the TS under two alternative configurations: the current organizational structure and the new design proposed by hospital management. To compare both scenarios, several key performance indicators (KPIs) were defined: (i) the average number of consultations per patient until resolution (**CN**); (ii) the total average number of consultation-related activities, including both in-person visits and e-consultations (**TN**); (iii) the average physician time per case (**ATC**); and (iv) the TS recovery time (**RT**), defined as the time required for the system to clear the existing waiting list and reach a steady state. Steady state is considered to be reached when the total waiting list falls below the average number of patients observed during the steady phase of a representative instance of the proposed alternative system, which is computed a priori using a sample parameter set.

As these indicators are intended to reflect the transient dynamics of the system during a recovery phase, the experimental design follows a *finite-horizon* approach, focusing on system behavior before equilibrium is reached. This makes an accurate reconstruction of the initial state—described in the previous section—especially important. Preliminary simulation tests indicated that a 10-year simulation horizon, using current patient arrival rates and extended service capacity, is sufficient both to reach steady state and to generate a robust sample of around 300000 patients for reliable estimation of the three main KPIs (CN, TN, ATC).

Experimental setup and exploration of referral behavior. For each of the two studied designs of the TS (current and proposed), the simulation model was run over a time horizon of ten years and replicated 30 times. Each replication used a different random seed, and common random numbers were applied to ensure that both designs were evaluated under identical stochastic conditions. For every replication, values of the four KPIs (CN, TN, ATC, RT) were recorded. Final estimates, including point values and confidence intervals, were obtained by analyzing the resulting samples of 30 observations for each scenario.

In addition to comparing the two system designs, the simulation model was used to explore how different referral behaviors by Primary Care physicians might influence system performance under the proposed scenario. This behavior was modelled using three parameters: α_1 (the tendency of Primary Care physicians to submit simple e-consultations referrals as in-person specialized consultations, due to the new availability of this option), α_2 (the tendency of Primary Care physicians to bypass e-consultations at the beginning of a patient's treatment process, as they can now refer directly to specialized consultations), and α_3 (the proportion of cases in which the e-consultation is bypassed, and the lack of information that would have been provided during the e-consultation makes it necessary to schedule an additional in-person consultation). Each parameter was varied from 0% to 100% in 10% increments, resulting in a full-factorial design of 11 × 11 × 11 = 1,331 scenarios. For each parameter combination, the simulation was executed and the KPIs were estimated, providing a comprehensive view of how the interaction between referral patterns and service configuration affects system performance. The next section presents the results of this analysis, along with recommendations for regional health policy-makers and hospital management.

5 RESULTS, ANALYSIS AND RECOMMENDATIONS TO HEALTH POLICY MAKERS

This section presents the results obtained from the simulation model under a wide range of scenarios. The analysis is structured into three parts: (i) a comparison between the current and proposed organizational designs under an expected behavioral scenario; (ii) a sensitivity analysis of key performance indicators (KPIs) as a function of Primary Care physicians' referral behavior; and (iii) recommendations for health service managers and policy-makers based on the simulation outcomes.

5.1 Comparison Between Current and Proposed System Designs

The first analysis compares the performance of the TS under its current structure and the proposed alternative design. The expected behavior of Primary Care physicians under the proposed design—specifically their use of referral channels ($\alpha_1 = 0.25, \alpha_2 = 0.5, \alpha_3 = 0.1$) —was estimated by hospital management based on expert judgement and historical data (the use of the current channels and how their used could be transformed). This expected scenario was used as the baseline for testing the performance of the proposed system.

Table 2 presents the estimated values of the four KPIs (CN, TN, ATC, RT) for both configurations. In the current design, the system fails to reach a steady state due to the continuous accumulation of patients in the general consultation waiting list. As a result, the RT is undefined, and the waiting list grows without limit, as illustrated in Figure 3. This behavior reflects the structural limitations of the current system and supports the need for a new organizational approach.

Table 2: Current and new design KPI comparison. **CN**: average number of consultations, **TN**: total number of consultation-related activities, **ATC**: average time per consultation, **RT**: recovery Time.

	CN Mean (IC)	TN Mean (IC)	ATC Mean (IC)	RT Mean (IC)
Current Design	1.7252 (1.72371, 1.72676)	1.9967 (1.9952, 1.9982)	29.0062 (28.9874, 29.0250)	Infinite
New Design	1.6541 (1.6526, 1.6556)	1.8501 (1.8486, 1.8516)	26.8985 (26.8831, 26.9140)	1402,90 (1400.09, 1405.71)

In contrast, the alternative design produces a significant improvement across all KPIs. The mean number of consultations per patient, the total consultation load, and the physician time per patient all show statistically significant reductions (p-value < 0.001), demonstrating that the proposed changes lead to more efficient care. In addition, the system reaches the steady state and the waiting list stabilizes (Figure 3). The periodic spikes observed are due to reduced capacity during the summer period.

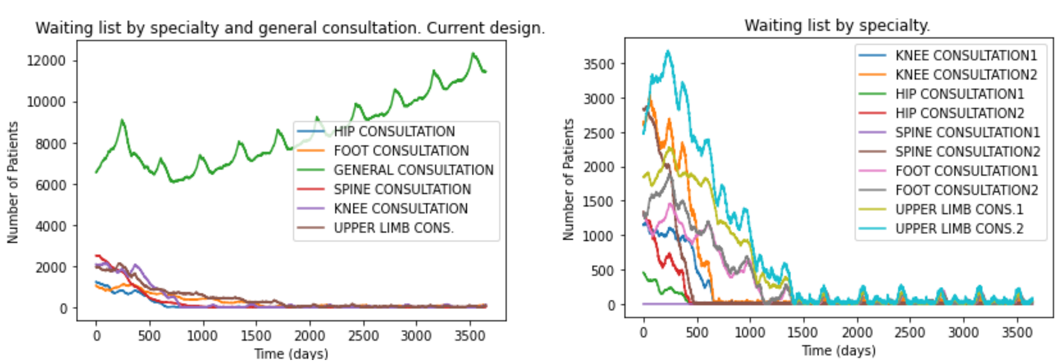

Figure 3: Patient waiting lists evolution in the current design (left) and the new design (right).

5.2 Sensitivity Analysis of Referral Behavior

Although the new design outperforms the current system under the expected scenario, its effectiveness depends on how Primary Care physicians use the available referral options. To test the robustness of the improvements, a comprehensive sensitivity analysis was conducted using the three behavioral parameters defined in Section 4:

- **α_1**: percentage of referrals submitted as in-person visits instead of e-consultations.
- **α_2**: percentage of e-consultations omitted at the beginning of the patient's trajectory.
- **α_3**: percentage of patients bypassing e-consultation and referred directly to a first specialized visit that need and additional consultation.

The analysis of the 1,331 scenarios obtained as combination of these parameters, showed that, in a small subset of cases, the mean number of consultations per patient (CN) could increase if e-consultations were misused—particularly when they are bypassed or fail to resolve cases effectively (see Figure 4). However, in all scenarios, the remaining patient and physician related KPIs (TN, ATC) consistently improved under the alternative design.

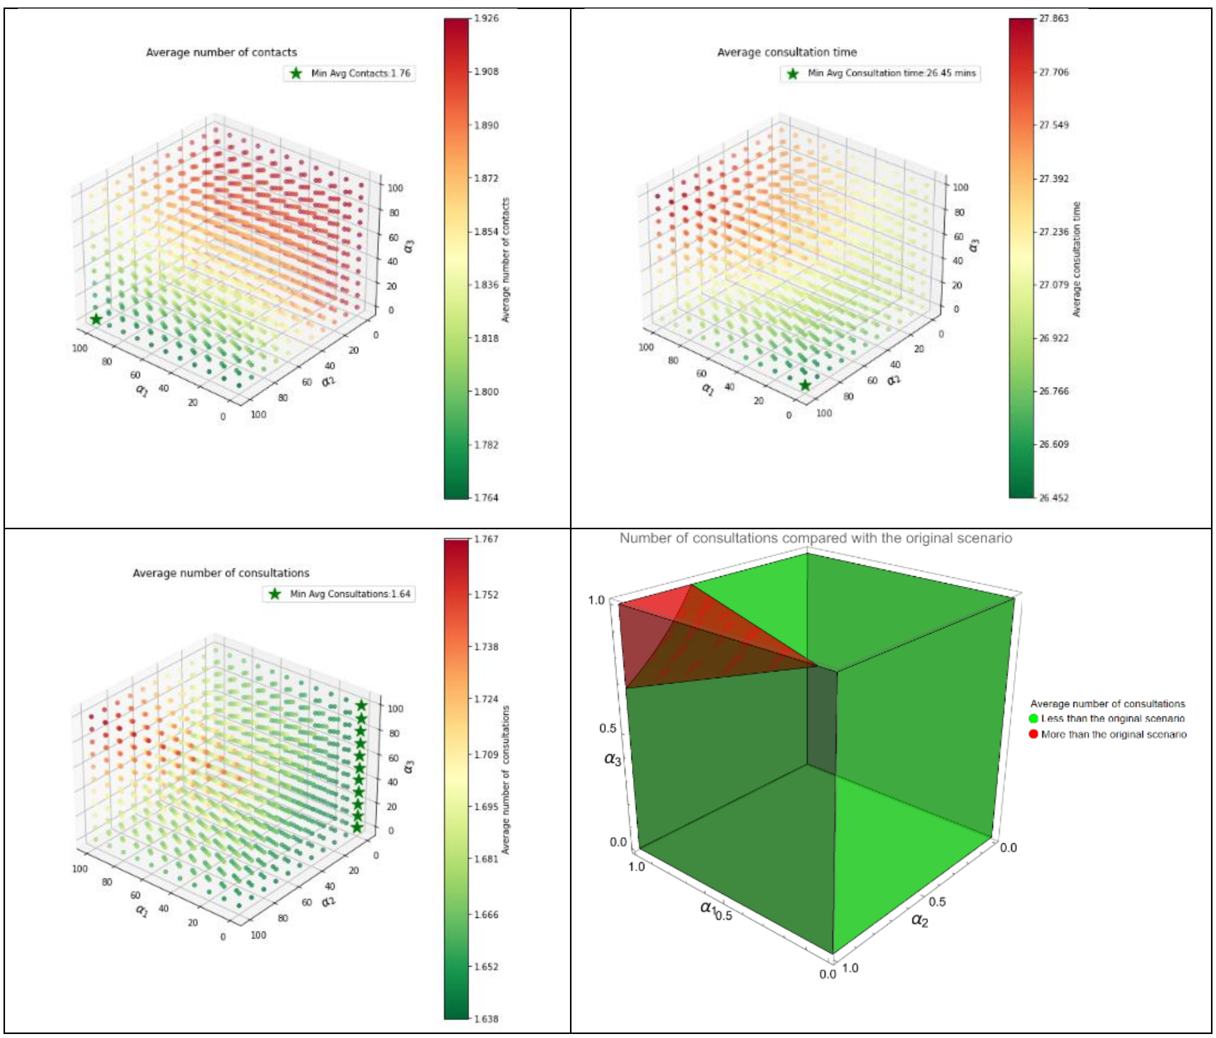

Figure 4: Top-left: TN; top-right: ATC; bottom-left: CN; bottom-right: region of the behavioral parameters (in red) that leads to an increase of CN under the new design compared to the current one.

To further quantify the influence of referral behavior on system performance, regression models were fitted to the 1,331 simulation outputs. A linear model with interaction term between two behavioral parameters explained nearly all the observed variability in the KPIs (R^2 = 99.99%). Table 3 presents the coefficients of the regression models for each KPI.

Table 3: Coefficients of the regression models for each KPI. **CN**: average number of consultations, **TN**: total number of consultation-related activities, **ATC**: average time per consultation.

	CN	**TN**	**ATC**
Intercept	1.6378	1.9256	26.9960
$\boldsymbol{\alpha_1}$	0.0314	-0.0015	0.4943
$\boldsymbol{\alpha_2}$	0.0060	-0.1603	-0.5485
$\boldsymbol{\alpha_2 * \alpha_3}$	0.0913	0.0919	0.9160

Using these coefficients, we analyze the case in which 50% of all patients who would have previously entered the system via e-consultation are instead assumed to access it directly through in-person consultation. This behavioral shift equals setting α_1 equal to α_2 ($\alpha_1 = 0.5, \alpha_2 = 0.5, \alpha_3 = 0.1$), with the aim of illustrating the impact of behavioral changes in Primary Care physicians (Table 4).

Table 4 summarizes the results of increasing α_1 from 25% to 50% compared to the baseline scenario for the new design. The results show a 0.45% increase in the average patient consultation time, which translates into 61.79 additional hours of consultations per year—equivalent to 370 more follow-up consultations. Similarly, a 0.46% increase of TN corresponds to 232 additional consultations per year (including both first visits and follow-ups). These findings highlight that improper use of the new referral options can lead to a significant increase in the workload of the TS.

Table 4: Effects over the KPI of increasing α_1 from 25% to 50%, to match α_2. **CN**: average number of consultations, **TN**: total number of consultation-related activities, **ATC**: average time per consultation.

	CN	**TN**	**ATC**
Baseline scenario with the new design	1.8497	1.6532	26.8911
Increased α_1 scenario	1.8493	1.6610	27.0147
Variation (%)	-0.0203	0.4685	0.4595

5.3 Recommendations to Health Policy Makers and Managers

The simulation results strongly support the adoption of the proposed redesign of the TS. The new design demonstrates a generalized improvement in system efficiency, reducing waiting times, consultation load, and physician workload. The only KPI that may be negatively affected is the mean number of consultations per patient, and only in cases where the new referral channels are used inappropriately.

The degree of improvement, however, is highly dependent on how Primary Care physicians engage with the new system. Their active and informed participation is essential to ensure that the changes lead to meaningful and sustainable improvements. Therefore, alongside the implementation of structural reforms, health authorities should invest in communication and training strategies aimed at encouraging appropriate use of referral options.

By quantifying how individual professional behavior affects system-wide KPIs, the simulation provides actionable insights that make the consequences of day-to-day clinical decisions more tangible. This evidence-based approach can help foster greater alignment between operational practice and strategic system goals, ultimately contributing to more patient-centered care delivery.

6 DISCUSSION AND CONCLUSIONS

This study presents a simulation model specifically designed to evaluate the transient behavior of a complex healthcare system—the TS of a public hospital—starting from its actual state at a specific point in time. By integrating data from multiple decentralized sources, the model reconstructs patient queues, care trajectories, and clinical workloads with high fidelity. When the simulation is updated in real time and synchronized with operational data, it functions as a digital shadow of the service. Moreover, when this model is used to simulate future system evolution under different scenarios and return actionable insights to decision-makers, it begins to operate as a digital twin. The long-term goal of the collaborative development between hospital management, IT services, and the clinical leadership of the TS is to achieve this fully operational digital twin. This tool will enable efficient resource allocation across key activities such as first consultations, follow-ups, e-consultations, and surgeries.

The immediate purpose of the simulation, however, was to address an urgent question posed by health policy-makers—which could not wait for the full development of the project: whether the implementation of a redesigned access and internal workflow system in the TS would be effective in reducing waiting lists to legally mandated levels. This analysis was particularly relevant because the proposed changes involve considerable investment and organizational restructuring, including the redesign of scheduling systems, IT development, training for Primary Care and TS physicians, and communication with patients. In this context, the model provided quantitative evidence to support the viability and expected impact of the new design. It also enabled the estimation of how long enhanced medical capacity would need to be maintained until the system stabilizes and recovers from its current backlog.

Beyond the technical insights, the model has already had practical implications. An extended version of the simulation results was included in a report presented by hospital leadership to the regional health department as part of the formal request to authorize the new service configuration. The model's outputs provided the solid, data-driven justification required for such a decision. Additionally, the sensitivity analysis revealed how essential the behavior of Primary Care physicians is in determining the effectiveness of the new model. This finding was used internally to raise awareness among referring physicians of how their individual decisions can influence the functioning of the entire system—an aspect that is often difficult to convey through conventional training or guidelines. Therefore, this research highlights the critical role of mathematical modelling—and simulation in particular—in supporting strategic decision-making in healthcare.

Finally, as with any model-based approach, this study has limitations. It does not yet incorporate economic evaluations, patient-reported outcomes, or staff satisfaction, which could enrich the decision-making framework. Moreover, the assumptions regarding referral behavior—while carefully constructed and validated through expert input—remain subject to real-world variability. Future developments will extend the model to dynamically allocate medical capacity across activities, enabling proactive rather than reactive management of waiting lists. The current model offers a solid methodological foundation for its extension to other departments of specialized medicine and to other geographical areas, thus contributing to broader systemic improvements in public healthcare planning.

ACKNOWLEDGMENTS

This work was supported by Proyecto PID2023-149957OB-I00 de investigación financiado por MICIU/AEI /10.13039/501100011033/ y por FEDER and Ayuda Formación de Profesorado Universitario (FPU22/00461). We acknowledge the Traumatology Service of the University Hospital of Navarre for their collaboration.

REFERENCES

Berti, A., S. van Zelst, and D. Schuster. 2023. "PM4Py: A Process Mining Library for Python". *Software Impacts* 17:100556.

Brailsford, S. C., P. R. Harper, B. Patel, and M. Pitt. 2009. "An Analysis of the Academic Literature on Simulation and Modelling in Health Care". *Journal of Simulation* 3(3):130–140.
Elkefi, S. and O. Asan. 2022. "Digital Twins for Managing Health Care Systems: Rapid Literature Review." *Journal of Medical Internet Research* 24(8):e37641.
Fone, D., S. Hollinghurst, M. Temple, A. Round, N. Lester, A. Weightman, K. Roberts, E. Coyle, G. Bevan, and S. Palmer. 2003. "Systematic Review of the Use and Value of Computer Simulation Modelling in Population Health and Health Care Delivery". *Journal of Public Health Medicine* 25(4):325–335.
Garcia-Vicuña, D., L. Esparza, and F. Mallor. 2022. "Hospital Preparedness during Epidemics Using Simulation: The Case of COVID-19". *Central European Journal of Operations Research* 30(1):213–249.
Greenwood-Lee, J., L. Jewett, L. Woodhouse, and D. A. Marshall. 2018. "A Categorisation of Problems and Solutions to Improve Patient Referrals from Primary to Specialty Care". *BMC Health Services Research* 18(1):986.
Jun, J. B., S. H. Jacobson, and J. R. Swisher. 1999. "Application of Discrete-Event Simulation in Health Care Clinics: A Survey". *Journal of the Operational Research Society* 50(2):109–123.
Kianfar, K. and A. Atighehchian. 2023. "A Hybrid Heuristic Approach to Master Surgery Scheduling with Downstream Resource Constraints and Dividable Operating Room Blocks". *Annals of Operations Research* 328(1):727–754.
OECD. 2020. *Waiting Times for Health Services: Next in Line*. OECD Health Policy Studies. Paris: OECD Publishing. https://doi.org/10.1787/242e3c8c-en, accessed 1st April 2025.
Rachuba, S., M. Reuter-Oppermann, and C. Thielen. 2024. "Integrated Planning in Hospitals: A Review." *OR Spectrum*. https://doi.org/10.1007/s00291-024-00797-5.
Rahiminia, M., S. Shahrabifarahani, M. Alipour-Vaezi, A. Aghsami, and F. Jolai. 2025. "A Novel Data-Driven Patient and Medical Waste Queueing-Inventory System under Pandemic: A Real-Life Case Study". *International Journal of Production Research* 63(2):418–434.
Wang, L. and E. Demeulemeester. 2023. "Simulation Optimization in Healthcare Resource Planning: A Literature Review". *IISE Transactions* 55(10):985–1007.

AUTHOR BIOGRAPHIES

MARTA CILDOZ is an Assistant Professor in the Department of Statistics, Informatics and Mathematics at Public University of Navarre. Her research interests include the optimization of the hospital patient flow by using simulation models, the development of digital twins and the scheduling of human resources. Her email address is marta.cildoz@unavarra.es.

MIGUEL BAIGORRI is a PhD Student in the Department of Statistics, Informatics and Mathematics at Public University of Navarre. His research interests include the modeling of human behavior by agents to develop hybrid simulation models of health care systems. His email address is miguel.baigorri@unavarra.es.

ISABEL RODRIGO-RINCÓN is Deputy Medical Director at the University Hospital of Navarra and has led and contributed to several research projects in patient safety, health information systems, and pandemic management. Her current work focuses on improving hospital planning and management systems through simulation- and optimization-based methods, with a particular emphasis on the development of digital twins. Her email address is mi.rodrigo.rincon@navarra.es.

FERMIN MALLOR is a Professor of Statistics and Operations Research at the Public Univeristy of Navarre. His research interests include stochastic modeling applications in healthcare and energy, hybrid simulation and simulation based optimization. He serves as an associate editor for Flexible Services and Manufacturing journal. His email address is mallor@unavarra.es and his website is https://www.unavarra.es/quphs.