Discrimination Under Time Pressure: Evidence from Emergency Triage Decisions

Ashley Litwin*

February 13, 2026

Abstract

I document large racial gaps in emergency department (ED) triage and show how acute time pressure amplifies those gaps. Using administrative data on nearly two million triage decisions by 1,800 nurses across six hospitals, I find that non-White patients are 5.4 percentage points (27 percent) less likely than observably identical White patients to receive the most urgent classification within nurse-by-shift comparisons. Exploiting quasi-random, minute-to-minute variation in available assessment time, I demonstrate that racial gaps widen monotonically when nurses are rushed, with time pressure accounting for at least 20 percent of the overall triage gap. Under tighter constraints, nurses collect less individualized clinical information, increasing the weight placed on race in triage assignments. I then trace the origins of these race-linked expectations. Using high-frequency variation in recent casemix, I show that exposure to frequent ED users—who are disproportionately non-White and typically assigned lower-acuity scores—systematically shifts subsequent racial gaps. Exposure to non-White frequent users widens disparities, whereas exposure to White frequent users attenuates or reverses them. The pattern holds within shifts: patients who follow frequent users experience different racial gaps depending on the race of the preceding patient. Finally, time-pressure-induced distortions at triage are associated with larger racial gaps in downstream wait times, diagnostic assessment, and treatment, including for time-sensitive conditions such as stroke. Together, the findings demonstrate that short-run time constraints can endogenously compress individualized assessment and amplify discrimination in high-stakes clinical decisions.

*Department of Business Economics and Public Policy, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104; alitw@wharton.upenn.edu. For their support and guidance, I thank Judd Kessler, Corinne Low, and Alex Rees-Jones. For helpful comments and suggestions, I thank Diane Alexander, Alex Chan, Zoe Cullen, Christine Exley, Marius Guenzel, Atul Gupta, Sasha Indarte, Ryan Oprea, Heather Schofield, and Kevin Volpp. Stephanie Dodd provided excellent research assistance. Funding came from the Wharton Behavioral Lab, the Wharton Coalition for Equity and Opportunity, and the Wharton Environmental, Social, and Governance Initiative.

1 Introduction

Racial health disparities remain pervasive in the United States, with non-White patients experiencing worse clinical outcomes and receiving lower-quality care relative to White patients (Currie and Schwandt, 2016; Schwandt et al., 2021; Alsan, Ianni and Peterson, 2024). While some disparities reflect differences in access, socioeconomic conditions, and environmental exposures (Alexander and Currie, 2017*b*; Grooms and Ortega, 2022; Dabiri et al., 2025; Powell et al., 2025), a growing body of evidence shows that unequal treatment also occurs within healthcare settings, where patients with similar presenting conditions receive systematically different evaluations and care (Alexander and Currie, 2017*a*; Alsan, Garrick and Graziani, 2019; Singh and Venkataramani, 2024). Despite this, we have limited evidence on how these disparities emerge within the often rapid, time-constrained decisions that structure frontline clinical care.

In this paper, I document a large racial gap in emergency department (ED) triage and identify acute time pressure as a micro-level constraint that systematically widens that gap. Using detailed, timestamped administrative data on 1.9 million triage decisions by 1,800 nurses across six hospitals between 2017 and 2025, I show that non-White patients are substantially less likely than observably identical White patients to be prioritized for care, comparing decisions by the same nurse during the same shift. I then demonstrate that this disparity increases sharply when nurses have less time available to evaluate a patient and examine how race-linked expectations shape triage decisions under time constraints.

I first document that, among patients with the same chief complaint evaluated by the same nurse during the same shift—and after controlling for a comprehensive set of clinical covariates drawn from the electronic health record[1]—non-White patients are roughly five percentage points (pp), or about 27 percent, less likely than White patients to receive the highest-priority classification. The gap is not explained by correlates of race such as primary care access or insurance coverage, nor by proxies for differences in communication or self-advocacy, suggesting that access or patient presentation alone cannot account for the

[1]Clinical covariates capture the near universe of information recorded in the electronic health record prior to and during triage, including age, sex, visit history, chronic conditions, primary care history, arrival mode, insurance type, and chief complaint, as well as detailed vital signs, pain scores, free-text triage notes, and screeners such as infectious disease, coma, and fall risk.

disparity. Similar patterns appear in nationally representative survey data, indicating that the disparity is not unique to the hospitals I study.[2]

A natural question is whether this disparity reflects racial differences in underlying illness severity or differences in how clinicians evaluate patients. Acute time pressure provides a way to distinguish between these possibilities. If racial gaps reflect unobserved differences in illness, they should not systematically vary with momentary fluctuations in available decision time. If instead disparities arise from how clinicians process information under constraint, the gaps should widen when decision time is limited. To test this prediction, I exploit quasi-random, minute-to-minute fluctuations in the time available at the start of each triage encounter. These fluctuations arise from idiosyncratic arrival patterns and variation in the duration of preceding assessments, generating sharp changes in available decision time that are orthogonal to the characteristics of the focal patient. This design allows me to compare how the same nurse evaluates clinically similar patients under tighter versus looser time constraints within the same shift. Consistent with quasi-random variation, observable pre-triage patient characteristics are balanced across high- and low-pressure moments within nurse shifts.

I show that quasi-random increases in time pressure systematically widen racial disparities in triage priority. Consistent with the differential-treatment interpretation, the racial gap increases monotonically as available decision time falls. As nurses become more rushed, triage decisions become increasingly unfavorable for non-White patients, while decisions for White patients are largely unaffected. This pattern is robust across three distinct measures of time pressure—the number of patients waiting to be triaged, the current wait time of the next patient in the queue, and the duration of the immediately preceding triage encounter—each capturing short-run variation in the time available for assessment.

I quantify the contribution of time pressure to the overall racial gap in triage priority using a counterfactual exercise based on the estimated reduced-form effects. Holding fixed patient composition, nurse assignment, and other covariates, I compare predicted triage outcomes

[2]The nationally representative data provide external validity but contain substantially less detailed clinical information and do not allow for the high-frequency, within-provider variation in time pressure exploited in the administrative data. They are therefore used to assess generalizability rather than to identify causal effects.

under the observed distribution of time pressure to a scenario in which all encounters occur under the lowest observed level of constraint. I find that variation in time pressure accounts for at least one-fifth of the overall racial gap in triage priority.

To benchmark the magnitude of this effect, I translate the implied widening of the White–non-White gap into age-equivalent units. Among otherwise similar patients, being one year older is associated with an increase of approximately 0.3 pp in the probability of receiving top triage priority. Moving from the lowest to the highest observed level of time pressure widens the racial gap by roughly 3.4 pp — equivalent to an increase of about 12 years in the implied age difference between White and non-White patients.

To understand why disparities widen under tighter time constraints, I examine how time pressure alters individual assessment behavior. Under greater time pressure, nurses spend less time evaluating patients, collect fewer diagnostic signals, and produce less detailed clinical documentation. These patterns indicate that constraint compresses individualized information acquisition, shifting judgment toward simplified assessments that rely more heavily on pre-triage information, including race.

These patterns are consistent with a simple economic framework of costly and endogenous information acquisition. When time constraints tighten, collecting individualized diagnostic signals becomes more costly, increasing reliance on baseline expectations formed prior to assessment. In this setting, any race-correlated expectations embedded in those baseline assessments exert greater influence on triage judgments when individualized information is compressed. Importantly, this mechanism does not require changes in underlying preferences or beliefs; disparities can widen solely because constraint shifts how information is acquired and weighted.

Guided by this framework, I then investigate how race-linked associations are formed and later enter triage assessments under constraint. In the ED, nurses repeatedly encounter returning patients ("frequent fliers"), who are disproportionately non-White and often perceived as lower acuity. I show that nurses' prior exposure to non-White frequent users systematically amplifies subsequent racial gaps in triage priority, while exposure to White frequent users attenuates or reverses them. Complementary within-shift evidence shows that racial gaps are larger for patients who happen to follow frequent users in the triage queue,

indicating that recent encounters shape how clinicians interpret patient race in real time. Together, these findings suggest that accumulated experience and moment-to-moment context shape what race signals about underlying severity, and that these expectations become more influential when individualized assessment is compressed.

Finally, I show that time pressure at triage is associated with larger racial gaps in downstream outcomes. Patients evaluated under greater constraint experience longer waits to see a physician, receive fewer diagnostic assessments, and are less likely to receive timely treatment. These disparities grow even for high-risk, time-sensitive conditions such as sepsis, heart attack, and stroke. Focusing on stroke presentations, the implied racial gap in permanent disability associated with delays in care increases from 0.26 pp under low time pressure to 0.64 pp under high time pressure—a nearly 150 percent widening relative to the low-pressure baseline. These results indicate that distortions introduced at the point of triage can have clinically meaningful downstream consequences.

Taken together, these results show that even temporary constraints on decision time can alter how clinicians acquire and process information, widening racial disparities in care and influencing downstream outcomes. Because time pressure shapes how individual judgments are formed—rather than which patients arrive or how sick they are—these findings suggest that disparities can emerge from the structure of the decision environment itself. Policies that reduce extreme time pressure in high-stakes clinical encounters may therefore improve not only the accuracy of triage decisions, but also their equity. More broadly, the evidence highlights how disparities can arise in standardized, protocol-driven settings when constraints shift the weight placed on individualized assessment versus baseline associations, without requiring changes in underlying preferences or beliefs.

This paper makes three main contributions to the literature. First, it contributes to the economics of discrimination by highlighting the role of decision environments in generating unequal treatment. Much of the existing literature focuses on distinguishing between belief-based and preference-based explanations for disparities across groups (Bertrand and Duflo, 2017; Lang and Kahn-Lang Spitzer, 2020; Litwin and Low, 2025).[3] The results here empha-

[3]A review by Bohren et al. (2023) finds that 62% of discrimination papers published in top economics journals from 1990–2018 discuss this distinction.

size a complementary mechanism: acute time pressure compresses individualized information acquisition and shifts decision weights toward coarse, experience-based associations. As a result, disparities can arise and intensify through changes in how decisions are made, rather than through differences in underlying beliefs or preferences alone. Importantly, this mechanism yields testable predictions about the direction and severity of unequal treatment: disparities emerge and intensify along dimensions that become salient through experience and are embedded in simplified decision rules under pressure. In this sense, the results are consistent with constructivist perspectives that view group characteristics not as a fixed input into decision-making, but as categories whose meaning and relevance are inferred from experience and context (Rose, 2023). By identifying this mechanism at the point of decision-making, the paper provides a micro-level foundation for downstream disparities in care and outcomes documented in high-stakes clinical settings (Singh and Venkataramani, 2024).[4]

Second, the paper contributes to the literature on coarsened or simplified decision-making. A growing body of theoretical and empirical research documents that decision-makers rely on coarse representations of information (Mullainathan, Schwartzstein and Shleifer, 2008; Bordalo et al., 2025) rather than fully individualized assessment in a variety of settings, including courts, labor markets, and finance (Chen, Moskowitz and Shue, 2016; Radbruch and Schiprowski, 2024; Barberis, Shleifer and Vishny, 1998; Akepanidtaworn et al., 2023). I provide large-scale field evidence on the microfoundations of such coarsening, showing when simplified rules are activated by acute time constraints, how their content is shaped by accumulated experience, and how they influence subsequent decisions.

Third, the paper contributes to the economics of clinical decision-making by isolating acute time pressure as a distinct margin shaping professional judgment. Prior work documents that physician behavior responds to fatigue (Jin et al., 2020), recent experience (Singh, 2021; Jin et al., 2023), peer behavior (Silver, 2021), and cognitive load (Handel et al., 2025). Related work shows that time pressure can reduce diagnostic documentation in primary care

[4]Singh and Venkataramani (2024) document that as hospitals approach capacity, in-hospital mortality rises for Black but not White patients, consistent with race-based rationing of inpatient beds. While these patterns are consistent with multiple underlying channels, the present paper complements this evidence by identifying a micro-level mechanism at the point of decision-making: as acute time pressure increases, decision-makers substitute toward coarse, readily observable characteristics, reweighting race and other demographic cues in ways that can plausibly generate the downstream disparities in care and outcomes documented in that setting.

(Freedman et al., 2021).[5] The present paper extends this evidence by examining how decisions are made once information use is truncated: which coarse associations substitute for missing information and how experience shapes the content of these simplified rules. Moreover, I demonstrate how rapid, high-stakes decisions made under constraint can generate systematic disparities in treatment.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and data, and Section 3 presents the conceptual framework. Section 4 outlines the empirical strategy and identifying equations. Section 5 summarizes the main results, while Section 6 and Section 7 provide mechanistic evidence on why race enters triage decisions when individualized information is compressed and quantify the contribution of time pressure to the overall racial gap. Section 8 examines downstream consequences, showing how time pressure-induced triage disparities are associated with increased racial gaps in diagnostic testing, treatment, and health outcomes. Section 9 concludes.

2 Institutional Setting and Data

I study triage decisions for the universe of patients who arrived at one of six EDs within a large U.S. health system between January 2017 and April 2025.[6] The system's hospitals are non-profit and include urban teaching hospitals as well as other acute-care hospitals serving nearby suburban communities, all of which use a unified electronic health record and standardized triage protocol. The analysis draws on detailed administrative data extracted from the electronic health record, which captures every clinical and operational data point entered by staff in real time, providing a near-census of ED activity.

My dataset comprises approximately 1.9 million triage encounters conducted by roughly 1,800 unique nurses. Each visit record includes pre-triage patient characteristics (age, sex,

[5] Lab experimental studies show that limited time affects risk-taking, strategic reasoning, cooperation, fairness, and decision quality more broadly (Kocher, Pahlke and Trautmann, 2013; Lawson, Larrick and Soll, 2020; Imas, Kuhn and Mironova, 2022; Kessler et al., 2025). A smaller number of studies document effects of deadlines and time limits in the field (De Paola and Gioia, 2016; Haji et al., 2019; Freedman et al., 2021; Carow and Witzig, 2025). This paper complements that work by providing large-scale field evidence from a high-stakes professional setting, exploiting minute-to-minute variation in decision time to show not only that these constraints matter, but how, when, and why they reshape decisions.

[6] Use of the de-identified administrative data was approved by the University of Pennsylvania Institutional Review Board.

race and ethnicity, arrival mode, and chief complaint), detailed information collected during triage (vital signs, triage notes, pain score, and the assigned priority level), and post-triage outcomes (waiting times, diagnostics, treatments, and discharge diagnosis). Crucially, the data contain precise timestamps for every key event—including patient arrivals, triage start and end times, and downstream encounters—which allow me to reconstruct the complete sequence of arrivals and triage decisions for each nurse and shift. This rich temporal structure links each decision to preceding and subsequent activity in the department and enables me to construct high-frequency measures of time pressure, information acquisition, and patient outcomes.

Upon arrival to the ED, each patient first completes a brief registration process in which basic information — age, sex, race and ethnicity, arrival mode, and presenting complaint — is recorded by intake staff. Patients then wait in a queue to be evaluated by the triage nurse, who calls them in order of arrival as capacity allows. Patients are effectively assigned to whichever nurse becomes available next, so assignment is quasi-random conditional on arrival time and overall department busyness.[7] During the triage encounter, the nurse measures vital signs and completes standardized screening questions — for example, assessing fall risk, pain level, or recent travel—and records brief clinical notes that synthesize the nurse's assessment of the patient's condition and provide context to downstream providers. The nurse then assigns an Emergency Severity Index (ESI) level, a standardized five-point scale used nationally in U.S. EDs to prioritize patients by clinical urgency: 1 indicates life-threatening emergencies requiring immediate attention, and 5 denotes non-urgent visits that can safely wait.[8]

The triage decision is the first formal point of clinical judgment and directly affects how quickly patients are roomed, evaluated by a physician, and begin diagnostic testing and treatment. The assessment typically lasts only a few minutes and occurs continuously throughout a 12-hour shift as patients arrive. Because arrivals are inherently stochastic, the

[7]Although triage nurses have discretion to call visibly ill patients ahead of others, hospital protocol is to triage patients in order of arrival. In practice, approximately 15 percent of encounters in my data occur out of sequence. All results are robust to including fixed effects for whether a patient was skipped and to excluding out-of-sequence cases altogether.

[8]The Emergency Severity Index (ESI) is used by approximately 94 percent of U.S. hospitals (Worth et al., 2019).

conditions under which each triage occurs varies naturally throughout a shift, creating short-run fluctuations in the time pressure nurses face. I leverage this variation in my empirical analysis outlined in Section 4.

2.1 Sample Construction and Key Variables

The unit of observation is a patient visit to the ED. Encounters with missing nurse or hospital identifiers, implausible timestamps, or incomplete information on patient race or gender are excluded.[9] After applying these restrictions, the analysis sample includes approximately 1.9 million visits conducted by 1,800 triage nurses across six hospitals.[10]

Patient Race Race and ethnicity are defined using self-reported information in the electronic health record. I collapse these categories into a binary indicator for non-Hispanic White versus non-White because several racial and ethnic subgroups appear infrequently in the data, making disaggregated estimates imprecise. This approach aligns with common practice in the health disparities literature and provides sufficient power for within-nurse-by-shift comparisons.[11]

Time Pressure I define time pressure as the momentary time constraint the nurse faces at the start of each triage assessment. Because triage encounters occur sequentially and patient arrivals vary unpredictably, this urgency fluctuates from minute to minute within a shift. I capture this variation using three real-time proxies measured at the start of each triage: (i) the number of patients waiting to be triaged behind patient i, (ii) the current waiting time of the *next* patient in line $(i+1)$ when patient i begins triage, and (iii) the duration of the immediately preceding triage encounter $(i-1)$. These measures provide complementary proxies for short-run time constraints and generate high-frequency within-shift variation in the environment in which each triage decision is made.

[9] I trim implausible values of time-in-triage and waiting time at the 5^{th} and 95^{th} percentiles; results are qualitatively unchanged when using the full sample or alternative trimming rules.

[10] Less than 5 percent of observations are excluded due to missing demographic information.

[11] See, for example, Alsan, Ianni and Peterson (2024) and Singh and Venkataramani (2024).

Information Acquisition and Processing I capture how much information the nurse collects and processes during triage along three complementary dimensions. First, I use the duration of the triage encounter, measured in minutes, which reflects the time spent gathering and processing clinical information and also provides a behavioral check on the time-pressure mechanism. Second, I record the number of structured fields completed during triage — vital signs (temperature, pulse, respiratory rate, oxygen saturation, and blood pressure) and standardized screening questions — out of twelve possible items, reflecting the extent of diagnostic information acquired. Third, I assess information processing using the length of the free-text triage note, defined as the log number of words recorded, which reflects the level of narrative detail the nurse records during the assessment. Together, these measures proxy the degree of individualized assessment, consistent with the framework in Section 3.

Outcome Variables The main outcome variable is an indicator for receiving top triage priority (ESI level 1 or 2), which designates that a patient either requires immediate life-saving intervention (ESI 1) or is considered high-risk and warrants rapid evaluation (ESI 2). I also examine three sets of secondary outcomes that capture the immediate and downstream consequences of triage decisions. First, I measure wait times, defined as the number of minutes between the recorded end of triage and first physician contact. To capture clinically meaningful delays in care, I additionally construct indicators for whether patients exceed recommended benchmarks for timely evaluation, including condition-specific targets for high-risk and time-sensitive presentations such as stroke, heart attack, and sepsis. Second, I examine downstream information acquisition intensity, including laboratory tests, imaging, and reassessments of vital signs or pain. Third, I assess downstream treatment intensity, including clinical procedures, pain medication administration, and whether the patient was admitted to the hospital. Together, these outcomes provide a comprehensive view of how triage decisions influence patients' access to timely evaluation and the course of care that follows.

Table 1 presents summary statistics for the analysis sample. The average patient is 47 years old, 44 percent male, and 59 percent non-White, and 21 percent arrive by ambulance.

Table 1: Summary Statistics

	(1) Mean	(2) SD	(3) N
Panel A. Patient Characteristics			
Age	47.47	21.78	1,896,860
Male	0.44	0.50	1,896,860
Non-Hispanic White	0.41	0.49	1,896,860
Arrived by ambulance	0.21	0.40	1,896,860
Panel B. Nurse-Shift Context			
Patients per nurse-shift	24.13	20.13	1,896,860
Panel C. Triage Assessment			
Top triage priority (ESI 1–2)	0.24	0.43	1,896,860
Minutes in triage	3.49	2.37	1,716,349
Clinical signals collected (out of 12)	8.89	3.40	1,896,860
Note length (words)	37.48	35.00	1,896,860
Panel D. Time Pressure			
Patients waiting	1.19	1.69	1,810,631
Next-patient wait time (min)	6.52	12.64	1,810,631
Duration of prior triage (min)	3.55	2.29	1,415,765

Note: The unit of observation is a patient visit to the emergency department. The table reports means, standard deviations, and non-missing counts. "Arrived by ambulance" equals one for patients transported by EMS (ambulance, fire rescue, police, and air/ground medical transport). "Patients per nurse–shift" is the total number of patients triaged by a nurse during a 12-hour shift. "Minutes in triage", "clinical signals collected", and "note length" summarize information recorded during the triage encounter. Time pressure proxies are measured at the start of triage; counts for these variables are lower when queue positions or adjacent-triage durations cannot be computed (e.g., the first patient of a shift or invalid timestamps).

Triage nurses evaluate roughly 24 patients per 12-hour shift, providing ample scope for within-shift identifying variation. During the triage assessment, nurses spend an average of 3.5 minutes with each patient, complete 8.9 of the twelve structured clinical fields, and record notes averaging 37 words, indicating that the triage encounter is both brief and information-dense. The time-pressure measures—patients waiting, the next patient's waiting time, and the duration of the prior triage—exhibit substantial dispersion across encounters, reflecting the highly dynamic decision environment in which these assessments occur.

3 Conceptual Framework

Motivation. The summary statistics in Table 1 highlight two defining features of the triage decision environment. First, assessments are brief: nurses spend only a few minutes with each patient on average. Second, the time available for assessment varies substantially within shifts due to fluctuations in queue length, next-patient waiting times, and the duration of prior encounters. Individualized evaluation is not costless—taking vital signs, conducting screeners, and documenting findings all require time and attention. When available time is limited, the feasible intensity of assessment declines. As a result, nurses may rely more heavily on information immediately available at the start of the encounter, including demographic characteristics such as race, rather than investing time in collecting and processing additional clinical evidence.

This framework formalizes that mechanism by modeling triage as a decision under costly and endogenous information acquisition. It clarifies how time pressure shifts reliance from individualized evidence toward baseline expectations, and how disparities can arise when those expectations are shaped by recent, potentially non-representative exposure. The structure is related to models of statistical and inaccurate statistical discrimination (Arrow, 1971; Phelps, 1972; Aigner and Cain, 1977; Bohren et al., 2023; Bohren, Imas and Rosenberg, 2019), which study how group-based priors interact with observed signals. In those settings, signals are typically treated as exogenously observed. Here, institutional constraints such as time pressure influence whether patient-specific evidence enters the assessment at all.

Clinical severity and decision problem. Each patient i arrives with observable pre-triage characteristics s_i (e.g., age, arrival mode, chief complaint) and an unobserved true clinical severity θ_i. The nurse must assign a binary triage decision $d_i \in \{0, 1\}$, where $d_i = 1$ denotes top triage priority.[12]

Endogenous measurement intensity. Before assigning priority, the nurse chooses whether to conduct an intensive individualized assessment. Such an assessment may include reviewing

[12] In practice, nurses assign patients a five-level ESI score. For simplicity and alignment with the empirical analysis, the model treats the decision as binary.

detailed vital signs, conducting additional questioning, or completing screening assessments. For simplicity, this choice is modeled as binary. Let $m_i \in \{0, 1\}$ denote measurement intensity, where $m_i = 1$ indicates a more thorough evaluation and $m_i = 0$ indicates reliance on baseline pre-triage information alone.

If $m_i = 1$, additional patient-specific evidence becomes available and informs the assessment of severity. If $m_i = 0$, no further individualized evidence is incorporated.

Intensive assessment requires time. At any moment, the time available for evaluation is limited and varies across encounters. Let time pressure at moment t, denoted τ_t, capture the tightness of available decision time. Time pressure may rise due to queue length, the wait time of the next patient in line, or unusually long prior encounters. Let the cost of intensive measurement be $c(\tau_t)$, where

$$c'(\tau_t) > 0.$$

Let Δ_i denote the expected improvement in triage accuracy from incorporating individualized evidence, which may depend on the degree of uncertainty about patient severity. The nurse chooses intensive measurement when its expected benefit exceeds its cost:

$$m_i = 1 \quad \text{if and only if} \quad \Delta_i \geq c(\tau_t).$$

As time pressure increases, $c(\tau_t)$ rises and the likelihood of intensive measurement declines. Time pressure therefore affects not only how information is weighted, but whether individualized information enters the decision at all.

Baseline expectations and exposure-shaped beliefs. When intensive measurement is not conducted ($m_i = 0$), the nurse forms a baseline assessment of severity:

$$b_i = \phi(s_i, g_i; C_t),$$

where g_i denotes patient group (e.g., race) and C_t summarizes recent caseload composition.

Group-based expectations need not reflect stable long-run averages. In practice, nurses may form expectations about the severity of different patient groups from the cases they

encounter (Tversky and Kahneman, 1973; Malmendier and Nagel, 2011, 2016). Let recent caseload composition at time t, denoted C_t, summarize the mix of patients seen in the recent past.

Certain patient types, such as frequent emergency department users or "frequent fliers," may differ systematically in perceived acuity. If non-White patients are disproportionately represented among such frequent users, then recent exposure to frequent users can shift expectations about patient race and acuity. In this case, the mapping from group membership to expected severity is shaped by the intersection of recent exposure and group composition.

Because the distribution of frequent users differs from the broader population distribution, recent caseload composition may not reflect the overall distribution of clinical severity by group. As a result, baseline assessments can depend on recent, institutionally structured exposure rather than long-run population averages.

If nurses form expectations based on their recent caseload — for example, because recent encounters are more cognitively accessible (Tversky and Kahneman, 1973) or because individuals overweight their own experienced history when forming beliefs (Malmendier and Nagel, 2011, 2016) — baseline assessments will reflect this selective exposure. Thus, even without animus or explicit bias, institutionally structured exposure can shape how group membership maps to expected severity, and this mapping can become increasingly consequential under time pressure.

The framework does not require that such exposure-shaped expectations be incorrect relative to true severity. However, when recent caseload composition differs from the underlying population distribution, baseline expectations may diverge from long-run averages.

Overall assessment and triage assignment. If $m_i = 1$, the nurse combines individualized evidence with the baseline assessment to form an overall risk index

$$r_i = \lambda x_i + (1 - \lambda) b_i,$$

where x_i denotes individualized clinical evidence observed under intensive measurement and $\lambda \in (0, 1)$ captures the weight placed on that evidence relative to the baseline assessment. If

$m_i = 0$, no individualized evidence is incorporated and the assessment relies entirely on the baseline:

$$r_i = b_i.$$

The nurse assigns top triage priority when

$$d_i = 1 \quad \text{if} \quad r_i \geq \kappa,$$

for some threshold κ.

As time pressure increases and intensive measurement becomes less likely, decisions depend more heavily on exposure-shaped baseline expectations. Any distortions embedded in those expectations therefore exert greater influence on assignments precisely when individualized information is compressed.

Implications. The framework yields three testable predictions that guide the empirical analysis. First, time pressure reduces measurement intensity: as τ_t increases, the probability of intensive individualized assessment declines. Second, if baseline expectations differ across groups, time pressure amplifies disparities by increasing the relative influence of those expectations. Third, if baseline expectations evolve with recent caseload composition, disparities will vary systematically with exposure history. Because individualized measurement is less likely under high time pressure, the effect of recent exposure on disparities will be strongest precisely when time pressure is high. Disparities therefore need not arise from animus. Instead, they can emerge from the interaction between time-constrained information gathering and exposure-shaped baseline expectations.

4 Empirical Strategy

This section describes how I identify the effect of time pressure on triage decisions and, in turn, on racial disparities in emergency care. The empirical strategy exploits minute-to-minute fluctuations in nurses' available decision time that arise naturally from the stochastic flow of patients through the ED. By comparing patients evaluated by the same nurse during

the same shift under slightly different time constraints, I isolate the effect of acute time pressure while holding constant the broader clinical environment, case mix, and institutional workload. This design links the conceptual framework of decision-making under costly information to real-world behavior and provides field-based evidence on how short-run time constraints translate into disparities.

4.1 Identifying Variation in Time Pressure

I exploit three complementary, quasi-random sources of variation in the time available for each triage. The first is the number of patients waiting to be triaged when the nurse begins the current evaluation. A larger queue increases urgency and may induce nurses to move more quickly through each assessment. Conditional on time of day, hospital, and nurse-shift fixed effects, short-run fluctuations in queue length arise primarily from random patient arrivals and are orthogonal to the underlying severity of the focal patient. To address concerns that busier periods might coincidentally involve compositional shifts in patient race or acuity, I include hospital-by-hour fixed effects in all specifications.

The second source of variation holds queue length constant and instead exploits the waiting time of the next patient in line. Even when the number of patients waiting is the same, the pressure to move faster is greater when the next patient has already been waiting longer. For example, consider a nurse triaging two patients, Patient A and Patient B, who both have exactly two other patients behind them in the triage queue. Suppose the next patient in line behind Patient A arrived 10 minutes earlier, while the next patient in line behind Patient B arrived 3 minutes earlier. The nurse may experience more urgency to complete Patient A's assessment, precisely because the next patient in line has been waiting longer for triage. Since the next patient's wait time depends on the stochastic timing of prior arrivals, it varies independently of the current patient's characteristics and provides a within-queue measure of time pressure.

The third source of variation captures the duration of the immediately preceding triage encounter. When a nurse spends unusually long on the prior patient, she may compensate by abbreviating the subsequent evaluation. This variation reflects idiosyncratic differences in how long earlier assessments happened to take—such as the time required to document or

stabilize the prior case—rather than systematic differences in the current patient's condition.

Because these three shocks occur at a minute-to-minute frequency within a nurse's 12-hour shift, they isolate short-run fluctuations in available decision time rather than slower-moving determinants of behavior such as fatigue, staffing levels, or cumulative workload. This feature makes them well-suited for studying how acute time pressure affects decision-making within shifts. In balance tests presented in Appendix Figure A.2, I show that the three measures of time pressure are not associated with economically meaningful differences in observable, pre-triage characteristics, supporting their interpretation as quasi-random changes in the amount of time available for each decision.

4.2 Estimating Equations

I quantify how acute time pressure affects both information acquisition and triage outcomes using models that exploit within-nurse-by-shift variation in time pressure. The main specification represents time pressure using a set of discrete indicators for different pressure levels, allowing the relationship between time pressure and outcomes to be fully flexible and interpreted directly as changes relative to the lowest-pressure condition.

Baseline racial gap. I first estimate the average racial gap in top-priority assignment using a linear probability model:

$$\text{TopTriage}_{ijst} = \alpha + \beta_W \,\text{White}_i + X'_{ijst}\gamma + \mu_{js} + \kappa_{c(i)} + \tau_t + \varepsilon_{ijst}, \tag{1}$$

where TopTriage_{ijst} indicates highest triage priority for patient i evaluated by nurse j during shift s at hospital h and time t; White_i indicates non-Hispanic White; X_{ijst} includes a male indicator, a cubic in age, clinical urgency covariates (vitals with missingness indicators), arrival mode, and visit number; μ_{js} are nurse-by-shift fixed effects; $\kappa_{c(i)}$ are chief-complaint fixed effects; and τ_t are hospital-by-hour fixed effects. Standard errors are clustered at the nurse level. The coefficient β_W is the within–nurse-shift difference in the probability of top triage for White versus non-White patients, conditional on X_{ijst} and the absorbed fixed effects.

Triage priority and its racial gradient. I next examine how time pressure influences triage decisions and whether these effects differ by patient race. Let $\{1\{\text{TP}_{ijst} \in k\}\}_k$ denote mutually exclusive indicators for binned time-pressure states (e.g., number-waiting categories, next-patient-wait bins, or previous-triage-duration quintiles), with the lowest-pressure bin omitted. I estimate:

$$\text{TopTriage}_{ijst} = \alpha + \sum_k \Big[\theta_k \, 1\{\text{TP}_{ijst} \in k\} + \phi_k \, 1\{\text{TP}_{ijst} \in k\} \times \text{White}_i\Big] + X'_{ijst}\gamma + \mu_{js} + \lambda_{ht} + \varepsilon_{ijst}, \tag{2}$$

where TopTriage_{ijst} is an indicator for receiving a triage score of 1 or 2 (out of 5), and White_i indicates non-Hispanic White. The coefficients θ_k capture the effect of time pressure for non-White patients relative to the lowest-pressure state, while the interaction terms ϕ_k capture how that effect differs for White patients. Standard errors are clustered at the nurse level.[13]

Notes on fixed effects and controls. Across all specifications, X_{ijst} includes a male indicator, a cubic in age, and clinical urgency covariates (temperature, pulse, respiratory rate, oxygen saturation, and systolic and diastolic blood pressure, each with missingness indicators). The nurse-by-shift fixed effects μ_{js} absorb time-invariant heterogeneity within a nurse's shift, such as baseline stringency or overall pacing, while hospital-by-hour fixed effects λ_{ht} flexibly capture crowding, capacity, and case mix at fine temporal resolution. When focusing on a single time-pressure proxy, the non-focal proxies are included as controls to isolate the intended source of variation.

Information acquisition outcomes. To understand why racial gaps grow with increased time pressure, I also examine how time pressure affects the amount of information collected during triage, focusing on assessment-intensity outcomes such as minutes in triage, number of diagnostic signals, and triage-note length. Using the same binned representation of time

[13]Results are unchanged when clustering by nurse-shift or using two-way clustering by nurse and nurse-shift.

pressure, I estimate:

$$y_{ijst} = \alpha + \sum_{k} \delta_k \, 1\{\text{TP}_{ijst} \in k\} + X'_{ijst}\gamma + \mu_{js} + \lambda_{ht} + \varepsilon_{ijst}, \tag{3}$$

where y_{ijst} denotes an information-related outcome for patient i evaluated by nurse j during shift s at hospital h and time t; TP_{ijst} indexes the level of time pressure; X_{ijst} includes demographics, arrival mode, chief-complaint fixed effects, and clinical covariates; μ_{js} are nurse-by-shift fixed effects; and λ_{ht} are hospital-by-hour fixed effects. The coefficients δ_k measure how assessment intensity changes relative to the lowest-pressure state, isolating the effect of acute time constraints on information collection and processing.

4.3 Identifying Assumptions

Identification relies on the assumption that, conditional on rich observables and fixed effects, short-run fluctuations in time pressure are quasi-random and uncorrelated with unobserved determinants of patient severity or triage priority, including those that vary by race. Each of the three sources of time-pressure variation used in the empirical analysis rests on a distinct identifying assumption, which I describe below.

(1) Number of Patients Waiting When the waiting room is fuller, nurses experience greater pressure to move quickly through each assessment. Short-run fluctuations in the number of patients awaiting triage arise primarily from random patient arrivals. The identifying assumption is that these fluctuations are not systematically correlated with race-specific differences in the underlying severity of the current patient, conditional on patient observables such as chief complaint, vital signs, and arrival mode.

(2) Wait Time of Next Patient In Line Even holding queue length fixed, time pressure increases when the next patient in line has already been waiting longer. For instance, if two patients both have two others waiting behind them, but in one case the next person has been waiting ten minutes instead of three, the nurse perceives greater urgency to complete the current evaluation. The identifying assumption is that, conditional on queue length and

observables, variation in the next patient's wait time is unrelated to race-specific severity of the current patient.

(3) Duration of Previous Triage If the prior patient's triage took unusually long, the nurse may rush to "catch up" by shortening the next evaluation. Such variation reflects idiosyncratic differences in how long earlier assessments happened to take, rather than systematic differences in the current patient's condition. The identifying assumption is that variation in the prior triage's duration is not systematically related to race-specific severity of the current patient.

Under these assumptions, the estimates that follow can be interpreted as reflecting the effects of acute time pressure on triage decisions and racial disparities.

5 Results

I now turn to the main results. I first document a large baseline racial disparity in top triage priority using within–nurse-shift comparisons and show that the gap is robust to extensive clinical and contextual controls. I then examine how acute time pressure affects triage decisions and demonstrate that racial gaps in top triage assignment widen substantially under higher time pressure. Finally, I present evidence on the underlying mechanism, showing that increased time pressure reduces assessment time and compresses information acquisition and processing.

5.1 Racial Disparities in Top Triage Priority

Table 2 reports estimates from Equation 1, where the dependent variable is an indicator for receiving top triage priority. The table documents a striking racial disparity in emergency triage decisions. In the unconditional comparison in Column (1), non-White patients are 9.2 pp less likely than White patients to receive top triage priority. Given that only about one in five non-White patients receive top priority, a 9.2 pp difference is a substantial disparity. As progressively richer controls are added—including chief-complaint fixed effects, vital signs with missingness indicators, arrival mode, and visit history—the estimated gap narrows but

remains meaningful, with more than half of the raw disparity persisting after accounting for detailed clinical information.

In the preferred specification with nurse-by-shift fixed effects in Column (4), which compares patients evaluated by the same nurse during the same 12-hour shift, the racial gap is 5.4 pp. This is a large difference in a setting where patients look clinically identical on observables and where triage assignments are intended to follow standardized guidelines.

Column 5 adds insurance type and primary-care access—variables that are strong correlates of race and reflect longer-run differences in access to care. These factors are not included in the preferred specification, since they likely capture structural consequences of race rather than independent drivers of triage decisions. Nonetheless, adding these variables provides a sense of how much of the disparity is associated with differences in longer-run access to care. The estimated racial gap remains essentially unchanged when they are included, reinforcing that observable differences in acuity, case mix, and healthcare access do not account for the disparity in triage priority and instead reflect systematic differences in how clinically similar patients are prioritized when they arrive at the ED.

Race and Communication One hypothesis is that the racial disparity might reflect differences in communication or self-advocacy rather than differences in how clinically similar patients are assessed. Appendix Table A.1 examines whether the disparity varies with socioeconomic status (SES), using insurance type as a proxy for SES.[14] If communication or advocacy were the primary mechanism, the racial gap should be smaller among higher-SES patients, who should be better positioned to convey their symptoms regardless of race. Instead, the SES interaction is close to zero, and the racial gap is nearly identical across SES groups. Appendix Table A.2 examines reported pain scores, which provide a standardized measure of how patients describe their symptoms. Conditional on chief complaint, vital signs, and demographics, non-White patients report *more* pain than White patients on average. Taken together, these results suggest that differences in communication or willingness to express symptoms are unlikely to explain the racial disparity in triage priority.

[14]Insurance type is strongly correlated with income, employment status, and underlying access to healthcare, and is a common proxy for SES (Fischer et al., 2020; Snyder and Chang, 2020; Monuteaux, Du and Neuman, 2024).

Table 2: Racial Disparities in Top Triage Probability

	(1) Unconditional	(2) Pre-Triage	(3) Clinical Urgency	(4) Nurse-Shift FE	(5) Contextual Factors
Non-Hispanic White	0.092***	0.057***	0.052***	0.054***	0.048***
	(0.005)	(0.003)	(0.003)	(0.003)	(0.003)
Observations	1,896,898	1,896,772	1,896,376	1,800,440	1,800,440
Nurses	1,815	1,815	1,815	1,661	1,661
R^2	0.013	0.276	0.303	0.408	0.410
Chief complaint FE	No	Yes	Yes	Yes	Yes
Hospital-Year FE	No	Yes	Yes	No	No
Month FE	No	Yes	Yes	No	No
Day-of-week FE	No	Yes	Yes	No	No
Hour-of-day FE	No	Yes	Yes	Yes	Yes
Visit number FE	No	No	Yes	Yes	Yes
Nurse-Shift FE	No	No	No	Yes	Yes
Ambulance FE	No	Yes	Yes	Yes	Yes
PCP FE	No	No	No	No	Yes
Insurance Type FE	No	No	No	No	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All regressions include a male indicator so racial gaps are estimated within gender. Unconditional probability of top triage for non-white patients is 19.9%. Column (1) reports the unconditional gap. Column (2) adds pre-triage controls: a cubic in age, ambulance arrival, chief complaint fixed effects, and hospital-by-year, month, day-of-week, and hour-of-day fixed effects. Column (3) additionally controls for clinical urgency, including temperature, pulse, respiratory rate, oxygen saturation, systolic and diastolic blood pressure (with missingness indicators), and patient visit number fixed effects. Column (4) introduces nurse-by-shift fixed effects (preferred specification), comparing patients seen by the same nurse during the same 12-hour shift. This specification retains the full set of clinical urgency controls, along with hour-of-day, visit number, ambulance, and chief complaint FE. Column (5) further absorbs contextual factors: primary care provider status and health insurance type. All specifications cluster standard errors at the nurse level.

Coefficient-Stability Bounds (Oster, 2019) As an additional check for omitted-variable bias, I apply Oster's coefficient-stability method, which compares movements in the race coefficient and model fit as rich controls are added and asks how much further they could move if unobservables were as predictive as observables. In the preferred specification, the White coefficient is 0.054, and the implied bound on the race effect is $\beta^* = 0.042$, which can be interpreted as the estimated race coefficient after allowing for the strongest plausible degree of selection on unobservables. The selection ratio needed to reduce the coefficient to zero is $\delta^* \approx 4.37$, meaning that unobserved factors would need to be more than four times as strongly selected as the observed controls to eliminate the disparity.[15] Taken together, these results suggest that omitted variables — whether unmodeled, high-dimensional interactions

[15]Following Oster (2019), I set $R_{\max} = 1.3 \times R^2_{\text{full}}$ and assume equal selection on observables and unobservables ($\delta = 1$).

of observables or factors observed by the triage nurse but not by the econometrician—are unlikely to fully account for the residual White–non-White gap in triage priority.

External Validity of the Baseline Gap in Nationwide Data To assess whether the racial disparity is specific to this health system or reflects broader patterns in U.S. EDs, I replicate the specification from Table 2, Column (2) using data from the National Hospital Ambulatory Medical Care Survey (NHAMCS), a nationally representative sample of U.S. hospitals covering the years 2012–2022.[16] Appendix Figure A.1 presents estimates of the racial gap in top triage priority in the NHAMCS data, showing both the unconditional and conditional gaps. The unadjusted White–non-White gap in NHAMCS is 2.9 pp, and the adjusted gap is 1.8 pp, which are smaller than the 5.7 pp disparity observed in my data. This difference may be driven by factors such as the inclusion of hospitals with lower racial heterogeneity and potentially simpler cases in the nationwide sample, while my data reflects larger, urban hospitals with greater racial diversity and more complex cases. Additionally, the nationwide data lacks some of the more granular controls included in my analysis, such as nurse-by-shift fixed effects, which could contribute to the observed gap differences. These results indicate that the disparity observed in my setting, while larger in magnitude, is not unique to these hospitals but reflects a broader pattern seen across U.S. EDs.

5.2 The Effect of Time Pressure on Triage Decisions

Having documented a large and robust racial disparity in emergency triage decisions, I now examine how acute time pressure contributes to this gap. This section shows that quasi-random, minute-to-minute increases in time pressure systematically widen racial disparities in triage priority. I then provide evidence on the mechanism underlying this pattern, showing that time pressure compresses assessment time and reduces the amount of individualized information nurses collect during triage.

I begin by examining how time pressure affects the distribution of triage priority across racial groups. To do so, I estimate Equation 2, which interacts indicators for time pressure with a White indicator. The interaction coefficients capture whether the effect of time

[16]The available controls in the NHAMCS match those in Column (2) of Table 2 most closely.

pressure on top-priority assignment differs between White and non-White patients, holding fixed nurse-by-shift conditions, hospital-by-hour constraints, presenting complaint, arrival mode, and detailed clinical covariates.

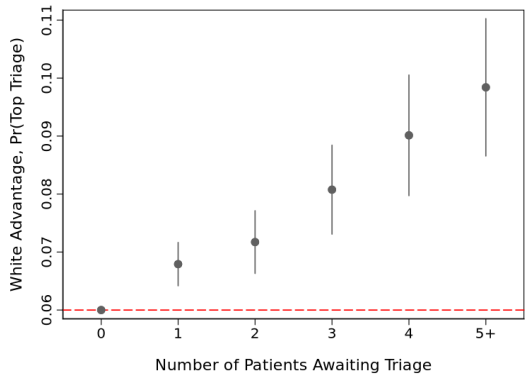
Figure 1 plots the estimated White–non-White gap in the probability of receiving top triage priority across bins of time pressure. Across all three quasi-random sources of variation, the racial gap widens monotonically as time pressure increases. In the lowest-pressure state, the estimated gap is roughly 6 pp; under the highest-pressure conditions, it is substantially larger, with point estimates increasing by roughly 2 to 4 pp depending on the measure of pressure. These effects are identified within nurse-by-shift and hospital-by-hour, ruling out differences in case mix or overall department crowding as explanations for the widening gap.

Although the magnitudes vary across the three time-pressure measures, the qualitative pattern is remarkably consistent: acute time pressure systematically amplifies racial disparities in top triage priority. In Section 7, I quantify the contribution of these effects to the overall racial gap and show that they account for a meaningful share of the disparity.

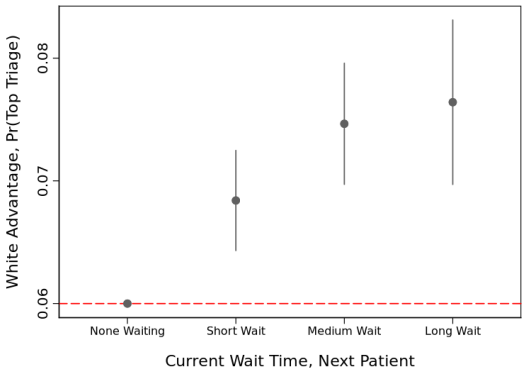
To benchmark the magnitude of these time pressure-induced disparities, I translate the resulting change in the White–non-White gap into age-equivalent units. Among otherwise similar patients, being one year older is associated with an increase of approximately 0.3 pp in the probability of receiving top triage at the sample mean age. Moving from the lowest to the highest observed level of time pressure widens the racial gap in top triage by about 3.4 pp. Expressed in age units, this corresponds to an increase of roughly 12 years in the implied age gap between White and non-White patients.[17]

I next turn to the mechanism underlying these effects. To assess how time pressure affects assessment intensity, I estimate Equation 3 with minutes spent in triage as the dependent variable. Across all three quasi-random sources of variation—the number of patients waiting, the waiting time of the next patient in line, and the duration of the immediately preceding triage—higher time pressure leads nurses to spend less time assessing the current patient.

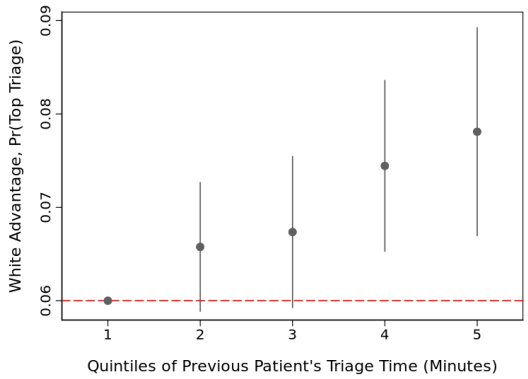
[17] I compute the change in the predicted White–non-White gap in the probability of receiving top triage when moving from the lowest to the highest observed level of time pressure, holding all other covariates fixed, and divide this gap change by the estimated marginal association between age and top triage probability at age 46, the sample mean.



(a) Patients Waiting



(b) Next Patient Wait



(c) Previous Triage Duration

Figure 1: Racial Disparities in Top Triage Probability Under Time Pressure

Note: Figure shows regression estimates of the White–non-White gap in the probability of receiving top triage across three quasi-random sources of time pressure. Panel (a) plots estimates by the number of patients awaiting triage, with 0 as the reference category (unconditional top-triage probability for non-White patients with 0 people waiting: 19.9 percent). Panel (b) plots estimates by the current wait time of the next patient in line, with "None Waiting" as the reference category. Panel (c) plots estimates by quintiles of the previous patient's triage time, with the shortest quintile as the reference category (unconditional probability for non-White patients following a prior triage in the shortest quintile: 17.0 percent). All regressions control for gender, a cubic in patient age, and an indicator for whether the patient was skipped over by someone else in the queue. The specification further absorbs nurse-by-12-hour-shift fixed effects, chief complaint fixed effects, arrival mode fixed effects, hospital-by-hour fixed effects, and the non-focal time pressure dimensions. Standard errors are clustered at the triage nurse level.

Figure 2 shows a clear and monotonic relationship in each case: relative to the lowest-pressure state, assessment duration falls sharply as time pressure increases. Moving from the reference bin to the highest-pressure bin reduces time spent in triage by approximately 4 to 6 percent, corresponding to a decline of roughly 0.15 to 0.20 minutes. While modest in absolute terms, this represents a meaningful compression of an already brief and high-stakes

clinical interaction.

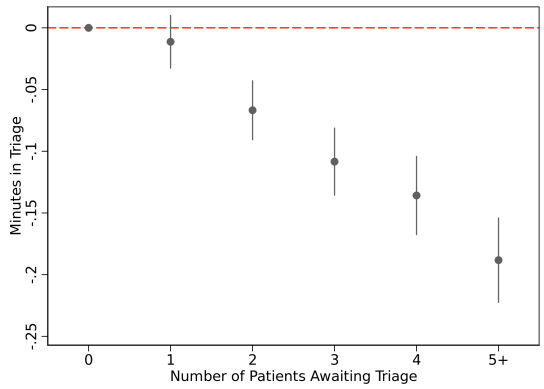
Consistent with shorter assessment times, nurses also collect and process less individualized information when they are under time pressure. Figure 3 shows that the number of structured diagnostic fields completed declines as time pressure rises, and Appendix Figure A.3 shows a corresponding reduction in triage note length.[18] Although the absolute changes in these measures are smaller than the reductions in assessment time, they represent meaningful contractions in the amount of patient-specific information incorporated into triage decisions. Together, these patterns reinforce the core mechanism described in the conceptual framework: time pressure limits information acquisition and processing, increasing reliance on coarse cues that are available prior to the assessment. As a result, when assessments are most compressed, the relative influence of race-correlated cues grows, helping explain why racial disparities in triage priority widen under time pressure, as shown in Figure 1.

Alternative Explanations and Mechanisms A remaining concern is that the widening racial disparities I attribute to time pressure may instead reflect cumulative fatigue or declining cognitive capacity as nurses progress through their shifts. Two features of the design and results make this explanation unlikely.

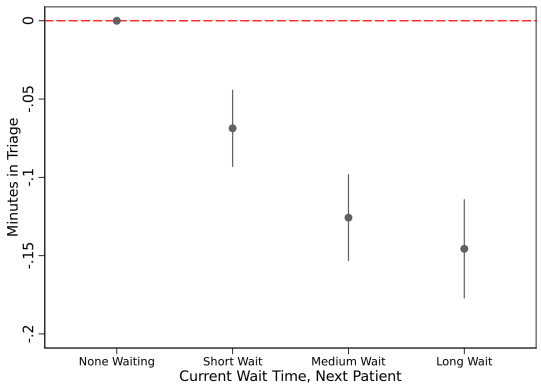
First, the sources of variation I exploit capture acute, high-frequency shocks to nurses' available decision time rather than gradual changes in energy or attention. Queue length, the waiting time of the next patient in line, and the duration of the immediately preceding triage fluctuate from minute to minute due to random patient arrivals and idiosyncratic variation in triage duration. By construction, these high-frequency changes are unlikely to be systematically related to the slow-moving accumulation of fatigue over a 12-hour shift.

Second, I present two sets of empirical results that are inconsistent with a fatigue-based interpretation. If cumulative fatigue were driving the widening disparities, the racial gap should increase systematically later in the shift. Instead, Appendix Table A.4 shows that the estimated effects of time pressure on triage decisions are broadly similar across shift

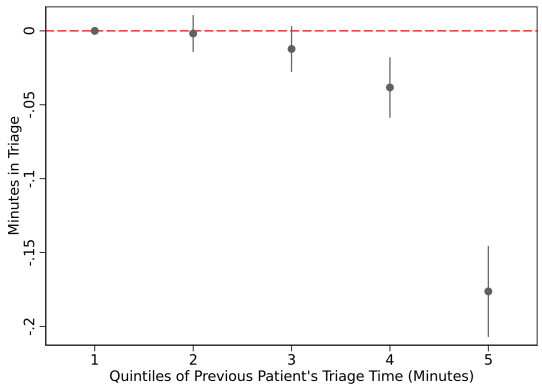
[18]Triage notes provide a useful proxy for cognitive effort and information processing: prior work suggests that when clinicians face higher workloads, they tend to write shorter, more formulaic notes, reflecting reduced cognitive engagement and processing (Hsu, Obermeyer and Tan, 2023). In this context, shorter notes likely indicate not only less documentation available to downstream providers, but also less processing of patient-specific information at the point of assessment.



(a) Patients Waiting



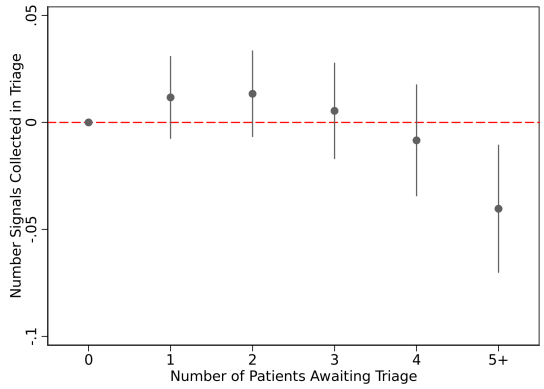
(b) Next Patient Wait



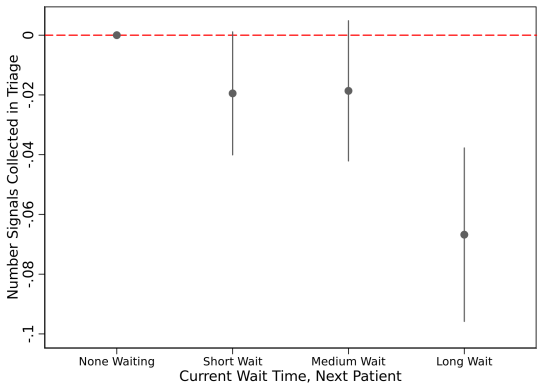
(c) Previous Triage Duration

Figure 2: Minutes in Triage Under Time Pressure

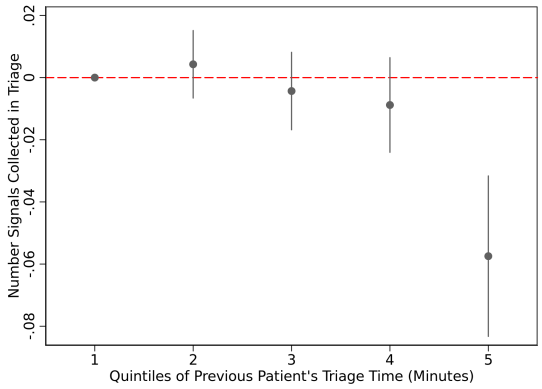
Note: Figure shows regression estimates of minutes spent in triage as a function of three quasi-random sources of time pressure. Panel (a) plots estimates by the number of patients awaiting triage, with 0 as the reference category (unconditional average triage time for patients with 0 people waiting: 4.2 minutes). Panel (b) plots estimates by the current wait time of the next patient in line, with "None Waiting" as the reference category. Panel (c) plots estimates by quintiles of the previous patient's triage time, with the shortest quintile as the reference category (unconditional average triage time for patients following a prior triage in the first quintile: 3.0 minutes). All regressions control for patient race (White vs. non-White), gender, a cubic in patient age, and an indicator for whether the patient was skipped over by someone else in the queue. The specification further absorbs nurse-by-12-hour-shift fixed effects, chief complaint fixed effects, arrival mode fixed effects, hospital-by-hour fixed effects, and the non-focal time pressure dimensions. Standard errors are clustered at the triage nurse level.



(a) Patients Waiting



(b) Next Patient Wait



(c) Previous Triage Duration

Figure 3: Number of Diagnostic Signals Collected Under Time Pressure

Note: Figure shows regression estimates of the number of diagnostic signals collected during triage (out of 12) as a function of three quasi-random sources of time pressure. Panel (a) plots estimates by the number of patients awaiting triage, with 0 as the reference category. The unconditional average number of signals collected for patients with 0 people waiting is 9.0. Panel (b) plots estimates by the current wait time of the next patient in line, with "None Waiting" as the reference category. Panel (c) plots estimates by the duration of the previous patient's triage, with the shortest quintile as the reference category. The unconditional average number of signals collected for patients following a prior triage in the shortest quintile is 9.0. All regressions control for gender, a cubic in patient age, and an indicator for whether the patient was skipped over by someone else in the queue. The specification further absorbs nurse-by-12-hour-shift fixed effects, chief complaint fixed effects, arrival mode fixed effects, hospital-by-hour fixed effects, and the non-focal time dimensions. Standard errors are clustered at the triage nurse level.

halves. High-time pressure widens the White–non-White gap for all three measures, while interactions with second-half indicators are statistically insignificant for the queue-length and previous-duration measures and small and unstable for the next-patient-wait measure. Appendix Table A.5 corroborates this conclusion: the estimated time-pressure effects remain stable when I flexibly control for within-shift progress using quintile fixed effects and their interactions with White, which absorb average changes in fatigue, pacing, and triage stringency as shifts progress. Taken together, these results indicate that the widening racial disparities are driven by acute, situational time constraints rather than by the gradual buildup of fatigue or cognitive depletion over the course of a shift.

6 Why Do Nurses Think About Race?

Having shown that time pressure increases the weight nurses place on race in triage decisions, I now examine why race in particular — rather than other patient characteristics — serves as a diagnostic cue in this setting. One plausible explanation is that nurses form associations between race and perceived illness severity through racialized patient categories. In emergency medicine, one such category is the "frequent flier" — patients with four or more ED visits in the prior 12 months who are often perceived as less sick or urgent (Grover and Close, 2009; LaCalle, Rabin and Genes, 2013; Krieg et al., 2016). Importantly, frequent fliers are a racialized category: in my data, frequent-flier encounters are disproportionately non-White. Among visits by non-White patients, 22.74 percent are frequent-flier visits, compared with 11.57 percent among White patients ($\approx$2.0$\times$). Among frequent-flier visits, 73.34 percent involve non-White patients. This racial imbalance makes the category a natural candidate for shaping racialized expectations about severity.

If nurses' decisions reflect associations between race and clinical severity that develop through experience with frequent fliers, then greater prior exposure to non-White frequent fliers should strengthen these associations, while greater exposure to White frequent fliers should attenuate them. To test this prediction, I analyze how the within–nurse-by-shift racial gap in top triage priority varies with a nurse's recent exposure to frequent fliers. I measure exposure as the share of a nurse's total patient encounters over the prior 90 days that involved

frequent fliers, computed before the start of the focal shift. Because this exposure measure is fixed within a shift and varies only across shifts of the same nurse or across different nurses, identification comes from comparing the within–nurse-by-shift racial gap across shifts with different levels of prior exposure. All specifications exclude patients who are themselves current frequent fliers, ensuring that the exposure measure is not mechanically related to the racial composition or acuity of the current patient.

Figure 4 plots the estimated White–non-White gap in top-priority assignment across quintiles of prior exposure to non-White and White frequent fliers. Two patterns emerge. First, the racial gap increases sharply with exposure to non-White frequent fliers: relative to the lowest-exposure quintile, the racial gap in the highest quintile is larger by approximately 8.4 pp. Second, exposure to White frequent fliers attenuates the racial gap: nurse-shifts in the highest exposure quintile exhibit a gap that is roughly 2.2 percentage points smaller than those in the lowest quintile. These relationships are estimated within nurse-by-shift and condition on chief complaint, arrival mode, detailed clinical covariates, and hospital-by-hour crowding conditions, helping to distinguish longer-run exposure effects from differences in current shift composition.

Together, these results are consistent with the interpretation that repeated interactions with specific patient types shape how race is associated with underlying illness severity. In this setting, frequent fliers are disproportionately non-White and commonly viewed as lower acuity, so greater exposure to non-White frequent fliers is associated with a widening of racial gaps under time pressure, while exposure to White frequent fliers is associated with a narrowing of these gaps. This pattern helps explain why race becomes a particularly salient cue in this environment and why it receives greater weight when nurses are constrained in the time available for individualized information gathering.

If nurses' expectations are shaped by racially patterned experience with frequent fliers, then these associations should be activated not only through accumulated exposure across shifts, but also by immediate context within the triage sequence. In particular, encounters with frequent fliers may cue race-related assessments in subsequent decisions. Table 3 reports estimates from regressions with an indicator for receiving top triage priority as the dependent variable and examines three related channels: whether frequent fliers are triaged differently

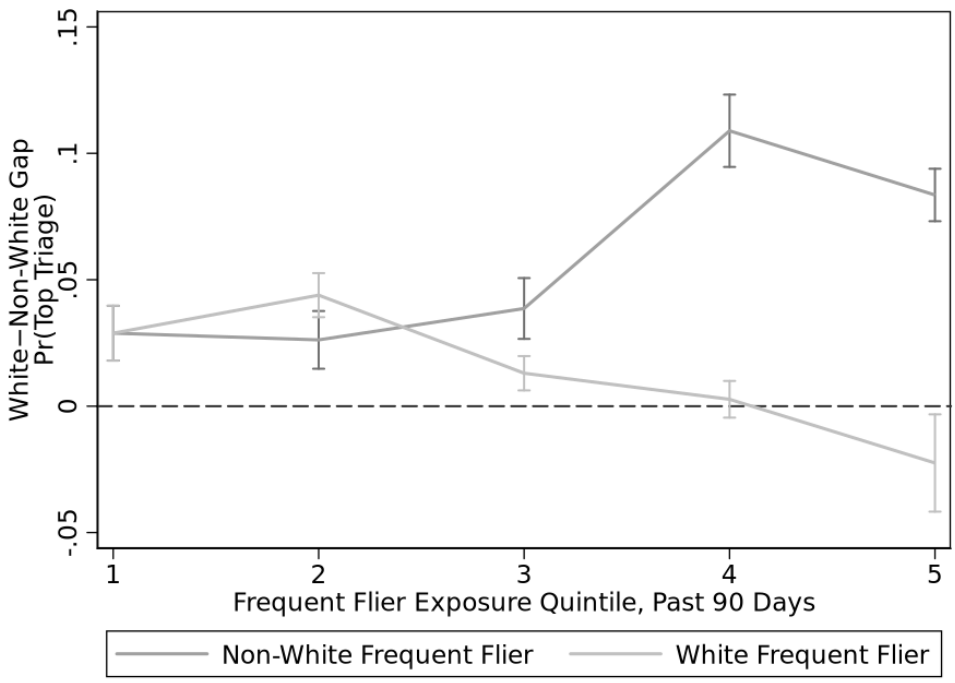


Figure 4: Racial Disparities in Top Triage Probability by Frequent Flier Exposure

Note: Figure shows regression estimates of the White–non-White differences in the probability of top triage by pooled quintiles of a nurse's recent exposure (prior 90 days, frozen at shift start) to non-White frequent fliers and White frequent fliers. Q1 (lowest exposure) is the reference category. The Q1 predicted gap is 2.89 pp. Frequent fliers are defined as patients with four or more visits in the prior 12 months. The specification includes both sets of exposure-by-White interactions simultaneously, controls for gender and a cubic in age, and absorbs nurse-by-12-hour-shift, chief complaint, and arrival mode fixed effects. The analysis sample excludes current frequent fliers and requires ≥ 30 prior visits to compute exposure. Quintile mean exposure shares: non-White FF — Q1 1.3, Q2 3.1, Q3 5.0, Q4 11.8, Q5 20.9; White FF — Q1 1.2, Q2 3.2, Q3 5.3, Q4 8.7, Q5 18.0. 95 percent confidence intervals shown; standard errors clustered at the triage nurse level.

from other patients; whether triage decisions shift after a nurse has just evaluated a frequent flier; and whether these sequential spillovers depend on the race of that preceding frequent flier. All specifications exclude current frequent fliers when appropriate and include nurse-by-shift fixed effects, so each coefficient reflects within–nurse-by-shift racial disparities under different encounter conditions, holding constant presenting complaint, arrival mode, clinical covariates, and other standard controls.

Column (1) compares frequent fliers to non–frequent fliers in the likelihood of receiving top triage priority. Among non–frequent fliers, the within–nurse-shift White–non-White gap is 6.0 pp. Among frequent fliers, this gap is 0.6 pp larger, representing a 10 percent increase in the size of the disparity. This pattern is consistent with the idea that the frequent-flier

Table 3: Frequent Flier Status and Racial Disparities in Top Triage

	(1) FF v. Non-FF	(2) Follows FF	(3) Follows NW FF
Non-Hispanic White	0.060***	0.060***	0.061***
	(0.003)	(0.003)	(0.005)
Frequent Flier (FF)	0.018***		
	(0.002)		
Non-Hispanic White × FF	0.006***		
	(0.002)		
Follows FF		-0.004***	
		(0.001)	
Non-Hispanic White × Follows FF		0.008***	
		(0.002)	
Follows Non-White FF			-0.008***
			(0.003)
Non-Hispanic White × Follows Non-White FF			0.015***
			(0.005)
Observations	1,800,441	1,800,436	278,086
Nurses	1,661	1,661	1,262
R^2	0.386	0.386	0.460
Nurse-Shift FE	Yes	Yes	Yes
Chief complaint FE	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes
Visit number FE	Yes	Yes	Yes
Frequent Flier FE	No	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Outcome is an indicator for receiving top triage priority. All regressions include a male indicator and a cubic in age, and absorb the fixed effects listed in the table. Column (1) includes an indicator for whether the focal patient is a frequent flier (FF), defined as having at least four ED visits in the prior 12 months (excluding the current visit). Column (2) includes an indicator for whether the immediately prior patient in the triage queue was an FF. Column (3) restricts the sample to patients who follow an FF and includes an indicator for whether that FF was non-White. Standard errors are clustered at the triage nurse level.

category carries additional contextual meaning that interacts with patient race in triage assessments. Although frequent fliers as a group receive slightly higher priority overall — consistent with some having chronic conditions that require recurrent care — the key pattern for understanding disparities is that the racial gap is larger within the frequent-flier category. Taken together, these results provide corroborating evidence that frequent-flier status conveys information that is correlated with race and clinical severity, shaping how race is

interpreted in triage decisions. [19]

Column (2) addresses a key limitation of Column (1): frequent-flier status may be correlated with underlying differences in patient severity, and these differences could themselves vary by race. To isolate a more plausibly exogenous source of variation, Column (2) examines whether racial disparities depend on whether the *preceding* patient in the triage queue is a frequent flier. When the preceding patient is not a frequent flier, the within–shift White–non-White gap in top triage priority is again 6.0 pp. When the current patient instead follows a frequent flier, the gap increases by 0.8 pp, representing a 13.3 percent increase. Because the identity of the immediately preceding patient is driven by the stochastic order of arrivals rather than by characteristics of the current patient, this pattern provides evidence of short-run sequential spillovers. In particular, recent exposure to a patient category that is disproportionately non-White and often viewed as lower acuity appears to cue race-linked assessments in subsequent triage decisions, exacerbating the White–non-White gap.

Column (3) sharpens the sequential-spillover pattern by conditioning on patients who follow a frequent flier and examining whether the effect depends on the race of that preceding frequent flier. Among patients who follow a White frequent flier, the within–nurse-shift White–non-White gap in top-triage assignment is 6.1 pp. When the preceding frequent flier is non-White, this gap increases by an additional 1.5 pp, representing a 24.6 percent widening of the racial disparity. The negative main effect (–0.8 pp) indicates that non-White patients who follow a non-White frequent flier are triaged more stringently, consistent with a within-group spillover. These patterns suggest that encounters with non-White frequent fliers differentially cue race-linked assessments in subsequent decisions.

Taken together, the evidence in this section shows that race becomes a salient cue in triage assessments through both accumulated experience and immediate context. Longer-run exposure to racially patterned patient categories is associated with systematic differences in how race maps to triage priority across shifts, while encounters with frequent fliers within a shift cue these associations in real time. These patterns indicate that race enters triage decisions not because it is inherently informative, but because institutional experience shape

[19]The standard clinical definition of a frequent flier—four or more visits in the prior 12 months—is conservative and groups together patients with heterogenous utilization patterns.

what race signals about underlying illness severity. This helps explain why constraints on individualized information gathering—such as acute time pressure—magnify racial disparities by increasing reliance on these associations.

7 Quantifying the Contribution of Time Pressure

I next assess how much of the overall racial disparity in triage priority can be attributed to time pressure. To do so, I estimate a linear probability model for top triage priority that includes the three time-pressure measures each interacted with race, alongside the full set of clinical controls and fixed effects. I then compute two model-based quantities using predictive margins: (i) the *as-observed* racial gap, defined as the difference in average predicted top-priority assignment between White and non-White patients when all covariates are held at their observed values, and (ii) a *counterfactual* gap in which all visits are evaluated in the lowest-pressure state (no one waiting and a short prior triage).

Table 4: Quantifying The Effect of Time Pressure

	Estimate (pp)	Std. Err. (pp)	95% CI Low (pp)	95% CI High (pp)
Observed gap	5.79	0.35	5.11	6.47
Counterfactual gap	4.62	0.52	3.60	5.64
Difference (CF–Obs)	-1.17	0.39	-1.94	-0.40
Share explained by TP (%)	20.2	6.9	6.8	33.6
Observations:	1,487,511			
Clusters (Nurses):	1,519			

Note: Entries are percentage points. "Observed gap" is the difference in average predicted top-priority assignment between White and non-White patients with all other covariates at their observed values. "Counterfactual gap" evaluates all visits in the lowest-pressure state (no one waiting and a short prior triage), averaging other covariates as observed. Standard errors are delta-method using the regression's nurse-clustered covariance matrix. Model includes ambulance, chief-complaint, 12-hour nurse-shift, visit-number, and arrival-hour fixed effects and full clinical controls.

Table 4 shows that the as-observed White–non-White gap is 5.79 pp (SE 0.35). Under the lowest-pressure counterfactual, the gap falls to 4.62 pp (SE 0.52), a reduction of 1.17 pp (SE 0.39), i.e., the gap is 20.2 percent smaller (95 percent CI [6.8 percent, 33.6 percent]). For non-White patients, top-priority assignment rises by 0.79 pp (a 4.1 percent relative increase), which closes about 14 percent of the racial gap. For White patients, the predicted probability of top triage falls by 0.38 pp (a 1.5 percent relative decrease) under low time pressure.

To check sensitivity to functional form, I collapse the three time pressure proxies into a binary *Low Time Pressure (Low TP)* indicator (no patients waiting; shortest next-patient wait bin; shortest prior-triage-duration bin). Appendix Table A.3 shows that the racial gap at typical pressure is 5.9 pp; under Low TP, it falls to 4.7 pp, a roughly 20 percent reduction. This closely matches the binned counterfactual.

Importantly, even this "lowest-pressure" state is not pressure-free. Triage is designed to be a rapid initial assessment, typically intended to last only a few minutes, so that patients can be moved quickly to a physician. Thus, even when no one is waiting in line and the prior triage was brief, nurses still operate under an institutional imperative to work quickly. As a result, my counterfactual should be interpreted as a lower bound on the contribution of time pressure, capturing the share of disparities due to additional variation in pressure beyond this baseline.

8 Downstream Implications of Triage Decisions

Triage represents the first moment of structured decision-making in the ED. In principle, it is designed to identify clinical urgency and initiate appropriate care. In practice, it serves as an early signal that may influence how downstream providers allocate attention and resources. When nurses are under time pressure, they have less capacity to collect and process information, leading to noisier assessments that may systematically understate the needs of some patients. Importantly, even small distortions in this initial signal can alter the sequence of subsequent decisions—affecting when patients are evaluated, how much additional information is gathered, how intensively they are treated, and ultimately what health outcomes they experience.

This section examines how nurse time pressure at triage is associated with racial disparities in downstream outcomes. I begin by analyzing wait times, both overall and for time-sensitive conditions, as short-run outcomes that capture how differences in triage priority coincide with delays in evaluation. I then turn to medium-run outcomes that reflect the intensity of downstream diagnostic effort and treatment provision, including measures of information acquisition (vital and pain reassessments, laboratory tests, and imaging)

and clinical interventions (procedures and pain medications). Finally, I examine longer-run outcomes including hospital admission and indicators linked to permanent disability risk. Together, these results assess whether time pressure at triage is linked primarily to short-lived delays in evaluation or to more persistent differences in how attention, information, and treatment are allocated across patients.

Table 5: Time Pressure and Racial Disparities in Wait Times

	(1) ESI 1–3 (Urgent)	(2) ESI 4–5 (Fast Track Eligible)
High Time Pressure (TP)	3.149***	0.566*
	(0.276)	(0.291)
Non-Hispanic White (Low TP)	-2.661***	-1.331***
	(0.292)	(0.333)
High TP × White	-3.315***	0.045
	(0.340)	(0.445)
Observations	932,381	299,367
Nurses	1,495	1,149
R-sq	0.506	0.628
Nurse-Shift FE	Yes	Yes
Hosp-Hour FE	Yes	Yes
Ambulance FE	Yes	Yes
Chief Complaint FE	Yes	Yes
Visit Number FE	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is minutes from triage to being seen by a physician. Average wait time for non-White patients under low time pressure is 56.04 minutes in Column (1) and 51.24 minutes in Column (2). "High Time Pressure (High TP)" equals 1 if the visit falls in the top two bins of *any* of the three time-pressure proxies (number of patients waiting, next-patient wait time, and previous triage duration; see text). Controls include a male indicator and a cubic in age. Fixed effects: nurse×shift, hospital-hour, arrival mode, chief complaint, and visit number. Standard errors (in parentheses) are clustered by triage nurse. Column (1) restricts to ESI 1–3 (urgent); Column (2) restricts to ESI 4–5 (low acuity/Fast Track).

Wait Times and Delays in Time-Sensitive Care I begin by examining short-run differences in patient wait times, which provide the most immediate indication of how nurse time pressure is associated with variation in the allocation of attention across patients. Table 5 reports how racial disparities in wait times vary with nurse time pressure. The dependent variable is the number of minutes between the end of triage and the patient's first physician contact. The key regressor, High Time Pressure (High TP), is an indicator equal to one when the nurse was in the top two bins of any of three quasi-random time pressure measures.

Column (1) restricts the sample to ESI 1-3 visits, which are higher-acuity cases treated in the main ED, while Column (2) reports the same specification for ESI 4–5 visits, which are lower-acuity cases often routed to Fast Track areas where staffing, physician availability, and queuing operate differently.[20]

Among urgent (ESI 1–3) visits, high nurse time pressure increases average wait times for non-White patients by about three minutes ($p < 0.01$; 5.6 percent). Under low time pressure, White patients are seen roughly three minutes faster than non-White patients. This disparity roughly doubles under high time pressure, reflecting substantially longer delays for non-White patients. Relative to the typical 30-minute operational target for door-to-provider times, this corresponds to a 10-20 percent widening in access disparities (Reiter, 2025). In contrast, for low-acuity (ESI 4–5) visits, many of which occur in Fast Track areas where patients are seen immediately, time pressure has little effect on either overall wait times (approximately 0.6 minutes, $p \approx 0.05$) or racial disparities.

Delays in initial evaluation are not only inequitable but can be clinically consequential for conditions in which timely diagnosis and treatment are critical. Table 6 examines how time pressure at triage is associated with racial disparities in delays beyond established clinical targets for time-sensitive conditions. The dependent variable measures the number of minutes by which a patient's diagnostic evaluation exceeded the recommended benchmark — 20 minutes for stroke, 30 minutes for sepsis, and 10 minutes for heart attack (Leisman et al., 2016; Barbour and Thakore, 2017; Stocker et al., 2025).[21] The coefficient on *High Time Pressure (High TP)* reflects the association between being triaged under tighter time constraints and delays for non-White patients, while the interaction *High TP* × *White* captures how that relationship differs for White patients. Across all specifications, high time pressure is associated with substantially longer delays for non-White patients: evaluations occur roughly four minutes later in the pooled sample (Column 1), five minutes later for stroke (Column 2), six minutes later for sepsis (Column 3), and two minutes later for heart-attack presentations (Column 4). Even under low-pressure conditions, White patients are evalu-

[20]Fast Track areas are designed to reduce congestion in the main ED by diverting and treating low acuity patients quickly. Because wait times are determined by different staffing models, patient flows, and queueing rules in the main ED and Fast Track, pooling the two settings would mix fundamentally different care processes.

[21]Patients seen more quickly than the clinical target are coded as zero delay.

Table 6: Time Pressure and Racial Disparities in Delays Beyond Clinical Targets

	(1) All Time-Sensitive (ESI 1–3)	(2) Stroke	(3) Sepsis	(4) Heart Attack
High Time Pressure (TP)	3.941***	4.666***	5.686***	2.301**
	(0.586)	(0.944)	(1.513)	(1.072)
Non-Hispanic White (Low TP)	-3.209***	-2.977***	-3.735***	-3.182***
	(0.562)	(0.823)	(1.327)	(1.123)
High TP × White	-4.570***	-4.433***	-5.117**	-1.530
	(0.828)	(1.192)	(2.115)	(1.707)
Observations	168,393	66,163	20,302	50,251
Nurses	1,195	1,053	789	913
R-sq	0.600	0.623	0.649	0.677
Nurse–Shift FE	Yes	Yes	Yes	Yes
Hospital–Hour FE	Yes	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes	Yes
Condition FE	Yes	No	No	No
Visit Number FE	Yes	Yes	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is the number of minutes by which the evaluation exceeded the clinical target for time-sensitive conditions (stroke, sepsis, and heart attack). Clinical targets are defined as 20 minutes for stroke (door-to-CT), 30 minutes for sepsis (door-to-physician evaluation), and 10 minutes for heart attack (door-to-ECG). Average delay for non-White patients under low time pressure is 43.72 minutes in the pooled sample, 36.87 minutes for stroke, 27.03 minutes for sepsis, and 56.50 minutes for heart attack. "High Time Pressure (High TP)" equals one if the triage nurse was in the top two bins of any of the three time-pressure proxies (number of patients waiting, next-patient wait time, or previous triage duration; see text). Controls include a male indicator and a cubic in age. Fixed effects: nurse×shift, hospital-hour, arrival mode, and visit number. Column (1) additionally includes indicators for the specific time-sensitive condition. Standard errors (in parentheses) clustered by triage nurse. Column (1) pools all time-sensitive visits (ESI 1–3); Columns (2)–(4) restrict the sample to stroke, sepsis, and heart attack presentations, respectively.

ated faster than non-White patients — by roughly three to four minutes across conditions —indicating a baseline racial gap in timeliness. Under high pressure, this gap widens further: for stroke, the *High TP* × *White* of -4.43 ($p < 0.01$) implies that the White advantage grows from about 3 to 8 minutes, and for sepsis, the corresponding coefficient (–5.12, $p < 0.05$) indicates a comparably large widening. The heart-attack estimate (–1.53) is directionally consistent with these patterns but not statistically significant. Overall, the results indicate that time pressure is associated with an amplification of pre-existing racial disparities in the timeliness of evaluation for the most time-sensitive emergency conditions.

To gauge the clinical relevance of the estimated delays, Table 7 benchmarks the racial gaps against medical evidence linking time to patient harm. I focus on stroke because it has well-defined, minute-level benchmarks relating delays to both biological damage and functional outcomes. Saver (2006) estimates that each minute of untreated stroke destroys

roughly 1.9 million neurons — about 22 days of normal brain aging per minute. Applying this benchmark, the additional 4.4 minutes of delay experienced by non-White patients under high time pressure corresponds to roughly 8 million extra neurons lost, or about three months of accelerated brain aging relative to White patients seen under the same conditions. Although this may appear modest, concentrated and rapid injury within a single brain region — as occurs in stroke — can eliminate functions such as speech or movement (Ding et al., 2020; Salvalaggio et al., 2020). Additionally, Saver et al. (2016) shows that each hour of treatment delay for stroke is associated with an approximately 5.2 pp increase in the probability of disability at 90 days. The estimated 4.4-minute racial delay under high time pressure therefore corresponds to an increase of roughly 0.4 pp in disability risk for affected stroke patients. While modest at the individual level, this magnitude illustrates how minute-level disparities in evaluation speed can translate into meaningful differences in health outcomes.

Table 7: Clinical Consequences of Evaluation Delays for Stroke Patients

	Racial Gap in Delay (Minutes)	Neurons Lost (Millions)	Increased Pr(Disabled) (pp)
Low time pressure	2.98	5.66	0.26
High time pressure	7.41	14.08	0.64
Gap widening (High – Low)	4.43	8.42	0.38

Note: Neuron-loss equivalents for stroke are computed using Saver (2006), which estimates that each minute of untreated ischemic stroke destroys approximately 1.9 million neurons. Values are calculated as *Delay (minutes)* $\times$ 1.9 million. Increases in the probability of long-term disability are calculated using the Saver et al. (2016) meta-analysis, which estimates that each hour of treatment delay increases the probability of 90-day disability by 5.2 pp ($\approx$ 0.0867 pp per minute). Values are calculated as *Delay (minutes)* $\times$ 0.0867 pp/minute.

Downstream Information Acquisition and Treatment Beyond delays in initial evaluation, time pressure at triage is also associated with differences in the intensity of downstream diagnostic assessment. Table 8 shows that, among urgent (ESI 1-3) encounters, non-White patients receive consistently less downstream assessment than White patients, and this racial gap widens substantially under high time pressure. Under low time pressure, non-White patients already receive fewer imaging, laboratory tests, vitals reassessments, and pain reassessments, than clinically similar White patients. Under high time pressure,

these gaps widen further: non-White patients become 0.2-0.4 pp less likely to receive any reassessment or diagnostic test and receive substantially fewer total diagnostic assessments—for example, 0.7 fewer lab tests and 0.1 fewer vitals reassessments on average.[22] For White patients, diagnostic intensity under high time pressure declines by significantly less or even increases modestly. For example, the racial gap in the number of laboratory tests ordered grows by roughly 1.7 tests under high time pressure (Column 4), a large increase relative to the baseline White–non-White difference. Because these downstream assessments play a central role in confirming or correcting the initial triage signal, these patterns indicate that time pressure is associated with increasingly divergent information environments for White and non-White patients as care unfolds.

Table 8: Time Pressure and Racial Disparities in Downstream Information Acquisition

	Imaging		Labs		Vitals		Pain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Any	Number	Any	Number	Any	Number	Any	Number
High Time Pressure (TP)	-0.001	-0.013**	-0.002**	-0.747***	-0.002**	-0.106***	-0.004**	-0.053*
	(0.001)	(0.006)	(0.001)	(0.169)	(0.001)	(0.026)	(0.002)	(0.031)
Non-Hispanic White (Low TP)	0.021***	0.089***	0.004***	2.775***	0.017***	0.646***	0.022***	0.644***
	(0.002)	(0.009)	(0.001)	(0.337)	(0.001)	(0.051)	(0.002)	(0.049)
High TP × White	0.004**	0.027***	0.003***	1.724***	0.002	0.237***	0.013***	0.114**
	(0.002)	(0.008)	(0.001)	(0.255)	(0.002)	(0.039)	(0.002)	(0.045)
Observations	1,337,499	951,099	1,334,655	1,164,606	1,262,436	1,019,105	1,260,625	620,912
Nurses	1,629	1,593	1,626	1,614	1,613	1,595	1,617	1,535
R-sq	0.331	0.319	0.302	0.334	0.223	0.312	0.245	0.323
Nurse–Shift FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hospital–Hour FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Chief Complaint FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Visit Number FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Outcomes capture diagnostic information collection during the emergency department visit after triage: any imaging ordered and number of distinct imaging tests (Cols. 1–2), any laboratory tests and number of distinct lab tests (Cols. 3–4), any vitals reassessed and number of vitals reassessments (Cols. 5–6), and any pain reassessment and number of pain reassessments (Cols. 7–8). The omitted group (non-White patients under low time pressure) has an average outcome of 69.40% in Col. 1, 2.20 in Col. 2, 87.45% in Col. 3, 43.16 in Col. 4, 80.44% in Col. 5, 6.72 in Col. 6, 49.53% in Col. 7, and 6.78 in Col. 8. "High Time Pressure (High TP)" equals one if the triage nurse was in the top two bins of any of the three time pressure proxies (number of patients waiting, next-patient wait time, or previous triage duration; see text). Controls include a male indicator and a cubic in age. All models include nurse×shift, hospital-hour, arrival mode, chief complaint, and visit number fixed effects. Standard errors (in parentheses) are clustered by triage nurse. Sample includes all ESI 1–3 patients.

A similar pattern emerges for downstream treatment decisions, summarized in Table 9. Among ESI 1–3 visits, high time pressure is associated with a lower likelihood that non-White

[22]For outcomes measured as the number of labs, imaging tests, vitals checks, or pain reassessments, the coefficients capture intensive-margin effects, i.e., conditional on receiving any testing.

patients receive any procedure or pain medication. White patients again experience smaller declines — or, in some cases, modest increases — in treatment intensity under pressure. The interaction terms indicate that under high time pressure, the White–non-White gap widens along both the extensive and intensive margins of treatment. Because downstream clinicians observe triage scores and nurse documentation at the start of evaluation, these patterns are consistent with time pressure at triage being associated with differences in how diagnostic and treatment decisions unfold downstream.

Table 9: Time Pressure and Racial Disparities in Downstream Treatment

	Procedures		Pain Meds		Hosp. Admission
	(1)	(2)	(3)	(4)	(5)
	Any	Number	Any	Number	Any
High Time Pressure (TP)	-0.005***	-0.007	-0.005***	-0.022	-0.007***
	(0.001)	(0.005)	(0.001)	(0.017)	(0.001)
Non-Hispanic White (Low TP)	0.036***	0.094***	0.005***	0.410***	0.046***
	(0.002)	(0.007)	(0.002)	(0.024)	(0.003)
High TP × White	0.006***	0.018**	0.005***	0.051**	0.016***
	(0.002)	(0.007)	(0.002)	(0.026)	(0.002)
Observations	1,278,090	638,330	1,265,674	406,582	1,093,712
Nurses	1,613	1,549	1,617	1,431	1,610
R-sq	0.354	0.249	0.261	0.316	0.359
Nurse–Shift FE	Yes	Yes	Yes	Yes	Yes
Hospital–Hour FE	Yes	Yes	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes	Yes	Yes
Chief Complaint FE	Yes	Yes	Yes	Yes	Yes
Visit Number FE	Yes	Yes	Yes	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Outcomes capture treatment decisions during the emergency department visit after triage: any procedures and number of distinct procedures (Cols. 1–2), any pain medication and number of pain medications (Cols. 3–4), and hospital admission (Col.5). The omitted group (non-White patients under low time pressure) has an average outcome of 55.33% in Col. 1, 1.74 in Col. 2, 36.85% in Col. 3, 3.22 in Col. 4, and 23.32% in Col. 5. "High Time Pressure (High TP)" equals one if the triage nurse was in the top two bins of any of the three time pressure proxies (number of patients waiting, next-patient wait time, or previous triage duration; see text). Controls include a male indicator and a cubic in age. All models include nurse×shift, hospital-hour, arrival mode, chief complaint, and visit number fixed effects. Standard errors (in parentheses) are clustered by triage nurse. Sample includes all ESI 1–3 patients.

In addition to diagnostic and treatment effort, high time pressure is also associated with wider racial disparities in hospital admission, a clinically central and economically meaningful outcome. As shown in Column (5), Table 9, the White–non-White gap in admission probability increases under high time pressure, with non-White patients less likely to be admitted when nurses are rushed. Because admission decisions synthesize information from triage, diagnostic testing, and physician evaluation, this pattern is consistent with differences

in information accumulation and assessment that emerge earlier in the visit. Admission is among the most consequential decisions made during an ED encounter, governing access to inpatient monitoring, diagnostics, and treatment. The presence of widening racial disparities at this stage highlights the potential clinical and policy relevance of time pressure as a determinant of inequity in emergency care.

9 Conclusion

This paper shows that acute time pressure—a pervasive feature of frontline clinical care—can activate and amplify discriminatory treatment. Using rich, real-time administrative data from emergency department triage, I document a substantial racial gap in triage priority and show that this gap widens systematically with minute-to-minute fluctuations in nurses' available assessment time. When nurses are rushed, they spend less time evaluating patients, collect fewer individualized diagnostic signals, and place greater weight on coarse cues, including race. I further show that these race-linked expectations are shaped by recent exposure to racially patterned patient types, providing evidence on how discriminatory treatment emerges and intensifies under constraint.

These distortions do not remain localized to triage. Because triage scores shape downstream clinical decisions, time-pressure-induced differences at the point of initial evaluation propagate throughout the episode of care. Under higher time pressure, non-White patients wait longer to see a physician, experience larger delays relative to clinical targets for time-sensitive conditions, receive fewer diagnostic assessments and treatments, and are admitted at lower rates. These findings link short-run constraints on decision time to systematic differences in subsequent care and outcomes, consistent with the idea that seemingly small frictions in information acquisition can compound into larger disparities as decisions accumulate.

More broadly, the results highlight how the structure of the decision environment can generate discrimination, even in standardized, protocol-driven settings. In emergency triage, discriminatory treatment does not require changes in preferences or animus; it can arise when time constraints compress individualized assessment and increase reliance on experience-

based expectations. By leveraging quasi-random variation in available decision time, this paper isolates a micro-level mechanism through which disparities are triggered and amplified in high-stakes institutional contexts.

Taken together, the evidence suggests that even momentary fluctuations in available decision time can meaningfully widen racial gaps in care. Policies that reduce extreme time pressure in high-stakes clinical encounters may therefore improve not only the accuracy of medical decision-making, but also its equity. More generally, these findings underscore the importance of incorporating situational constraints into economic analyses of discrimination, particularly in environments where rapid judgments carry significant welfare consequences.

References

Aigner, Dennis J., and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *ILR Review*, 30(2): 175–187.

Akepanidtaworn, Klakow, Rick di Mascio, Alex Imas, and Lawrence DW Schmidt. 2023. "Selling Fast and Buying Slow: Heuristics and Trading Performance of Institutional Investors." *The Journal of Finance*, 78(6): 3055–3098.

Alexander, Diane, and Janet Currie. 2017*a*. "Are Publicly Insured Children Less Likely To Be Admitted To Hospital Than The Privately Insured (and Does It Matter)?" *Economics & Human Biology*, 25: 33–51.

Alexander, Diane, and Janet Currie. 2017*b*. "Is It Who You Are or Where You Live? Residential Segregation and Racial Gaps in Childhood Asthma." *Journal of Health Economics*, 55: 186–200.

Alsan, Marcella, Katherine Ianni, and Graeme Peterson. 2024. "Racial Health Disparities in The United States." *Oxford Review of Economic Policy*, 40(3): 498–517.

Alsan, Marcella, Owen Garrick, and Grant Graziani. 2019. "Does Diversity Matter for Health? Experimental Evidence from Oakland." *American Economic Review*, 109(12): 4071–4111.

Arrow, Kenneth. 1971. "The Theory of Discrimination." Princeton University, Department of Economics, Industrial Relations Section.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny. 1998. "A Model of Investor Sentiment." *Journal of Financial Economics*, 49(3): 307–343.

Barbour, Victoria, and Shobhan Thakore. 2017. "Improving Door To CT Scanner Times For Potential Stroke Thrombolysis Candidates–The Emergency Department's Role." *BMJ Quality Improvement Reports*, 6(1).

Bertrand, Marianne, and Esther Duflo. 2017. "Chapter 8 – Field Experiments on Discrimination." *Handbook of Economic Field Experiments*, 1: 309–393. eds. Abhijit Vinayak Banerjee and Esther Duflo, North-Holland.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "The Dynamics of Discrimination: Theory and Evidence." *American Economic Review*, 109(10): 3395–3436.

Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope. 2023. "Inaccurate Statistical Discrimination: An Identification Problem." *The Review of Economics and Statistics*.

Bordalo, Pedro, Nicola Gennaioli, Giacomo Lanzani, and Andrei Shleifer. 2025. "A Cognitive Theory of Reasoning and Choice." NBER Working Paper 33466.

Carow, Johannes, and Niklas M. Witzig. 2025. "Time Pressure and Strategic Risk-Taking In Professional Chess." *Journal of Economic Behavior and Organization*, 238.

Chen, Daniel L., Tobias J. Moskowitz, and Kelly Shue. 2016. "Decision Making Under the Gambler's Fallacy: Evidence from Asylum Judges, Loan Officers, and Baseball Umpires." *The Quarterly Journal of Economics*, 131(3): 1181–1242.

Currie, Janet, and Hannes Schwandt. 2016. "Mortality Inequality: The Good News From A County-Level Approach." *Journal of Economic Perspectives*, 30(2): 29–52.

Dabiri, Sanaz, Rema Raman, Linda B Cottler, Catherine W Striley, and Jevay Grooms. 2025. "Racial Climate on Willingness to Participate in Clinical Research." *Journal of Racial and Ethnic Health Disparities*, 1–14.

De Paola, Maria, and Francesca Gioia. 2016. "Who Performs Better Under Time Pressure? Results From A Field Experiment." *Journal of Economic Psychology*, 53: 37–53.

Ding, Junhua, Randi C. Martin, A. Cris Hamilton, and Tatiana T. Schnur. 2020. "Dissociation Between Frontal and Temporal-Parietal Contributions To Connected Speech In Acute Stroke." *Brain*, 143(3): 862–876.

Fischer, Igor, Hendrik-Jan Mijderwijk, Ulf D Kahlert, Marion Rapp, Michael Sabel, Daniel Hänggi, Hans-Jakob Steiger, Marie-Therese Forster, and Marcel A Kamp. 2020. "Association Between Health Insurance Status and Malignant Glioma." *Neuro-Oncology Practice*, 7(5): 531–540.

Freedman, Seth, Ezra Golberstein, Tsan-Yao Huang, David J Satin, and Laura Barrie Smith. 2021. "Docs With Their Eyes on The Clock? The Effect of Time Pressures on Primary Care Productivity." *Journal of Health Economics*, 77.

Grooms, Jevay, and Alberto Ortega. 2022. "Substance Use Disorders Among Older Populations: What Role Do Race and Ethnicity Play In Treatment and Completion?" *Journal of Substance Abuse Treatment*, 132.

Grover, Casey A., and Reb JH Close. 2009. "Frequent Users of The Emergency Department: Risky Business." *Western Journal of Emergency Medicine*, 10(3).

Haji, Anouar El, Michał Krawczyk, Marta Sylwestrzak, and Ewa Zawojska. 2019. "Time Pressure and Risk Taking In Auctions: A Field Experiment." *Journal of Behavioral and Experimental Economics*, 78: 68–79.

Handel, Ben, Louis Heizlsperger, Jonas Knecht, Jonathan Kolstad, Ulrike Malmendier, and Filip Matejka. 2025. "Thinking Versus Doing: Cognitive Capacity, Decision Making and Medical Diagnosis." *Working Paper*.

Hsu, Chao-Chun, Ziad Obermeyer, and Chenhao Tan. 2023. "Clinical Notes Reveal Physician Fatigue." *Working Paper*.

Imas, Alex, Michael A. Kuhn, and Vera Mironova. 2022. "Waiting to Choose: The Role of Deliberation in Intertemporal Choice." *American Economic Journal: Microeconomics*, 14(3): 414–440.

Jin, Lawrence, Rui Tang, Han Ye, Junjian Yi, and Songfa Zhong. 2020. "Time Dependency in Physician Decision-Making." *AEA Papers and Proceedings*, 110: 284–88.

Jin, Lawrence, Rui Tang, Han Ye, Junjian Yi, and Songfa Zhong. 2023. "Path Dependency in Physician Decisions." *The Review of Economic Studies*, 91(5): 2916–2953.

Kessler, Judd B., Hannu Kivimaki, Ashley Litwin, and Muriel Niederle. 2025. "Please Take A Minute: How Prosocial Choices Change With Deliberation." *Working Paper*.

Kocher, Martin G, Julius Pahlke, and Stefan T Trautmann. 2013. "Tempus Fugit: Time Pressure In Risky Decisions." *Management Science*, 59(10): 2380–2391.

Krieg, Cynthia, Catherine Hudon, Maud-Christine Chouinard, and Isabelle Dufour. 2016. "Individual Predictors of Frequent Emergency Department Use: A Scoping Review." *BMC Health Services Research*, 16(1).

LaCalle, Eduardo J., Elaine J. Rabin, and Nicholas G. Genes. 2013. "High-Frequency Users of Emergency Department Care." *The Journal of Emergency Medicine*, 44(6): 1167–1173.

Lang, Kevin, and Ariella Kahn-Lang Spitzer. 2020. "Race Discrimination: An Economic Perspective." *Journal of Economic Perspectives*, 34(2): 68–89.

Lawson, M. Asher, Richard P. Larrick, and Jack B. Soll. 2020. "Comparing Fast Thinking and Slow Thinking: The Relative Benefits of Interventions, Individual Differences, and Inferential Rules." *Judgment and Decision Making*, 15(5): 660–684.

Leisman, Daniel, Benjamin Wie, Martin Doerfler, Andrea Bianculli, Mary Frances Ward, Meredith Akerman, John K D'Angelo, and Jason A Zemmel D'Amore. 2016. "Association of Fluid Resuscitation Initiation Within 30 Minutes of Severe Sepsis and Septic Shock Recognition With Reduced Mortality and Length of Stay." *Annals of emergency medicine*, 68(3): 298–311.

Litwin, Ashley, and Corinne Low. 2025. "Measuring Discrimination with Experiments." *Handbook of Experimental Methods in the Social Sciences*. ed. Alex Rees-Jones, Edward Elgar Publishing Ltd.

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *The Quarterly Journal of Economics*, 126(1): 373–416.

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *The Quarterly Journal of Economics*, 131(1): 53–87.

Monuteaux, Michael C, Michelle Du, and Mark I Neuman. 2024. "Evaluation of Insurance Type As A Proxy For Socioeconomic Status In The Pediatric Emergency Department: A Pilot Study." *Annals of Emergency Medicine*, 83(6): 562–567.

Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer. 2008. "Coarse Thinking and Persuasion." *The Quarterly Journal of Economics*, 123(2): 577–619.

Phelps, Edmund. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review*, 62(4): 659–661.

Powell, David, Dan Han, Jevay Grooms, and Rosalie Liccardo Pacula. 2025. "Geography And Fentanyl: Explaining The Disproportionate Rise In Opioid Overdose Deaths Among Black Americans." *Health Affairs*, 44(9): 1050–1059.

Radbruch, Jonas, and Amelie Schiprowski. 2024. "Interview Sequences and the Formation of Subjective Assessments." *The Review of Economic Studies*, 92(2): 1226–1256.

Reiter, Mark. 2025. "11 Benchmarks That Should Matter to EPs." Accessed: 2026-01-15.

Rose, Evan K. 2023. "A Constructivist Perspective on Empirical Discrimination Research." *Journal of Economic Literature*, 61(3): 906–923.

Salvalaggio, Alessandro, Michele De Filippo De Grazia, Marco Zorzi, Michel Thiebaut de Schotten, and Maurizio Corbetta. 2020. "Post-Stroke Deficit Prediction From Lesion and Indirect Structural and Functional Disconnection." *Brain*, 143(7): 2173–2188.

Saver, Jeffrey L. 2006. "Time is Brain — Quantified." *Stroke*, 37(1): 263–266.

Saver, Jeffrey L, Mayank Goyal, AAD Van der Lugt, Bijoy K Menon, Charles BLM Majoie, Diederik W Dippel, Bruce C Campbell, Raul G Nogueira, Andrew M Demchuk, and Alejandro Tomasello. 2016. "Time To Treatment With Endovascular Thrombectomy and Outcomes From Ischemic Stroke: A Meta-Analysis." *Journal of the American Medical Association*, 316(12): 1279–1289.

Schwandt, Hannes, Janet Currie, Marlies Bär, James Banks, Paola Bertoli, Aline Bütikofer, Sarah Cattan, Beatrice Zong-Ying Chao, Claudia Costa, Libertad González, et al. 2021. "Inequality In Mortality Between Black and White Americans by Age, Place, and Cause and In Comparison to Europe, 1990 to 2018." *Proceedings of the National Academy of Sciences*, 118(40).

Silver, David. 2021. "Haste or Waste? Peer Pressure and Productivity in the Emergency Department." *The Review of Economic Studies*, 88(3 (320)): 1385–1417.

Singh, Manasvini. 2021. "Heuristics In The Delivery Room." *Science*, 374(6565): 324–329.

Singh, Manasvini, and Atheendar Venkataramani. 2024. "Rationing by Race." *NBER Working Paper 30380*.

Snyder, Rebecca A, and George J Chang. 2020. "Insurance Status As A Surrogate for Social Determinants of Health In Cancer Clinical Trials." *JAMA Network Open*, 3(4).

Stocker, Michael D., Chrissie Schaeffer, Randy Cox, Emily Tew, Kaitlyn Jensen, Kimberley Smith, Mitchell Sexton, Brian Bales, Amina Belghit, and Jonathan W Andereck. 2025. "Improving Door to ECG Time at a Quaternary Care Emergency Department." *The Joint Commission Journal on Quality and Patient Safety*.

Tversky, Amos, and Daniel Kahneman. 1973. "Availability: A Heuristic For Judging Frequency and Probability." *Cognitive Psychology*, 5(2): 207–232.

Worth, Murtis, Leslie L. Davis, Debra C. Wallace, Robin Bartlett, and Debbie Travers. 2019. "Are Emergency Departments In The United States Following Recommendations by The Emergency Severity Index To Promote Quality Triage and Reliability?" *Journal of Emergency Nursing*, 45(6): 677–684.

A Appendix

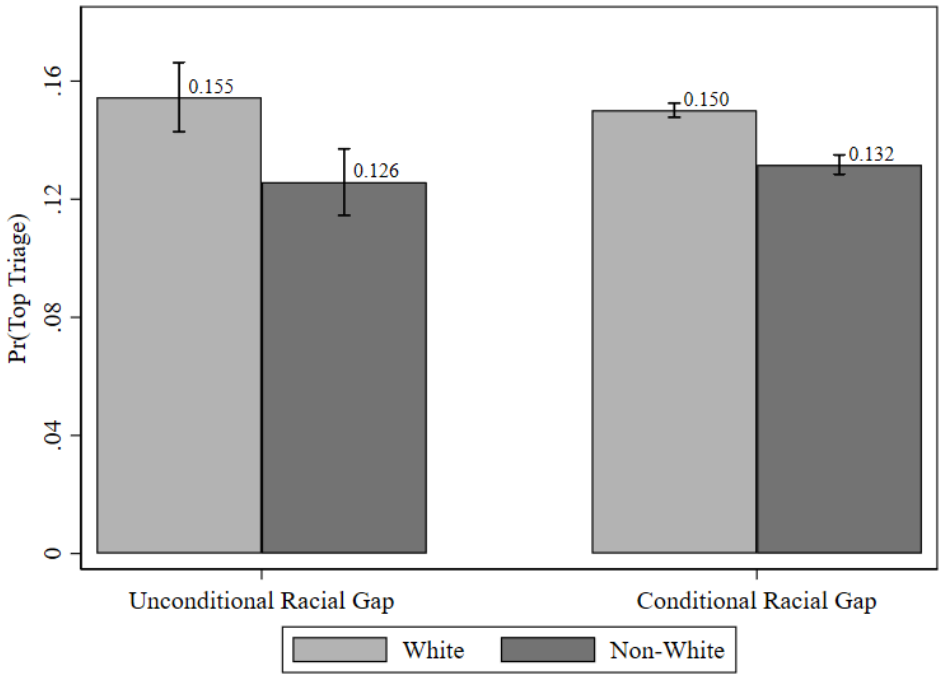


Figure A.1: Racial Disparities in Top Triage Probability in the U.S.

Note: Figure shows estimates from the National Hospital Ambulatory Medical Care Survey (NHAMCS), 2012-2022. All regressions include a male indicator so racial gaps are estimated within gender. The unadjusted White-non-White gap in the probability of receiving top triage is 2.9 pp; the adjusted gap is 1.8 pp. The adjusted specification controls for a cubic in patient age, an indicator for ambulance arrival, chief complaint fixed effects, hospital-by-year fixed effects, day-of-week fixed effects, hour-of-day fixed effects, and month fixed effects. Estimates use NHAMCS patient weights, with standard errors clustered at the primary sampling unit level to account for survey design.

Table A.1: Racial Disparities in Top Triage Probability by SES

	(1)
Non-Hispanic White	0.053***
	(0.004)
High-SES	0.015***
	(0.001)
Non-Hispanic White × High-SES	0.001
	(0.002)
Observations	1,188,279
Nurses	1,562
R^2	0.388
Nurse-Shift FE	Yes
Hour-of-day FE	Yes
Chief complaint FE	Yes
Ambulance FE	Yes
Visit number FE	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Regression includes a male indicator so racial gaps are estimated within gender. The dependent variable is an indicator for receiving top triage priority. Unadjusted probability of top triage for low-SES, non-white patients is 16.5%. "High-SES" is defined as having commercial or employer-sponsored insurance (Blue Cross, Commercial, or Managed Care). Patients with Medicaid or self-pay are classified as low-SES, and Medicare patients are excluded since eligibility is age-driven. The coefficient on "Non-Hispanic White" captures the racial gap in top triage probability among low-SES patients. The coefficient on "High-SES" captures the difference between high- and low-SES patients among non-Whites. The interaction term reports the difference in the racial gap between high-SES and low-SES patients. Specification controls for a cubic in age and clinical urgency covariates (temperature, pulse, respiratory rate, oxygen saturation, systolic and diastolic blood pressure, with missingness indicators), and absorbs fixed effects for nurse-by-shift, chief complaint, ambulance arrival, visit number, and hour of arrival. Standard errors are clustered at the nurse level.

Table A.2: Pain Score Documentation and Reported Pain by Race

	Missing Pain Score		Pain Score (0-10)	
	(1)	(2)	(3)	(4)
	Unadjusted	Adjusted	Unadjusted	Adjusted
Non-Hispanic White	0.103***	-0.004***	-1.056***	-0.699***
	(0.007)	(0.001)	(0.027)	(0.017)
Observations	1,896,898	1,800,440	1,020,298	940,458
Nurses	1,815	1,661	1,570	1,394
R^2	0.014	0.495	0.044	0.297
Nurse-Shift FE	No	Yes	No	Yes
Hour-of-day FE	No	Yes	No	Yes
Chief complaint FE	No	Yes	No	Yes
Ambulance FE	No	Yes	No	Yes
Visit number FE	No	Yes	No	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All regressions include a male indicator so racial gaps are estimated within gender. Unadjusted missingness of the pain score among non-White patients is 41.8 percent, and their unadjusted mean reported pain score (0–10 scale) is 4.34. Column (1) reports the unadjusted racial gap in the probability that a pain score is missing. Column (2) adds controls: a cubic in age, clinical urgency covariates (temperature, pulse, respiratory rate, oxygen saturation, systolic and diastolic blood pressure, with missingness indicators), and absorbs fixed effects for ambulance arrival, chief complaint, nurse-by-shift, visit number, and hour of arrival. Column (3) reports the unadjusted racial gap in reported pain score, conditional on the score being non-missing. Column (4) adds the same controls and fixed effects as Column (2). All specifications cluster standard errors at the nurse level.

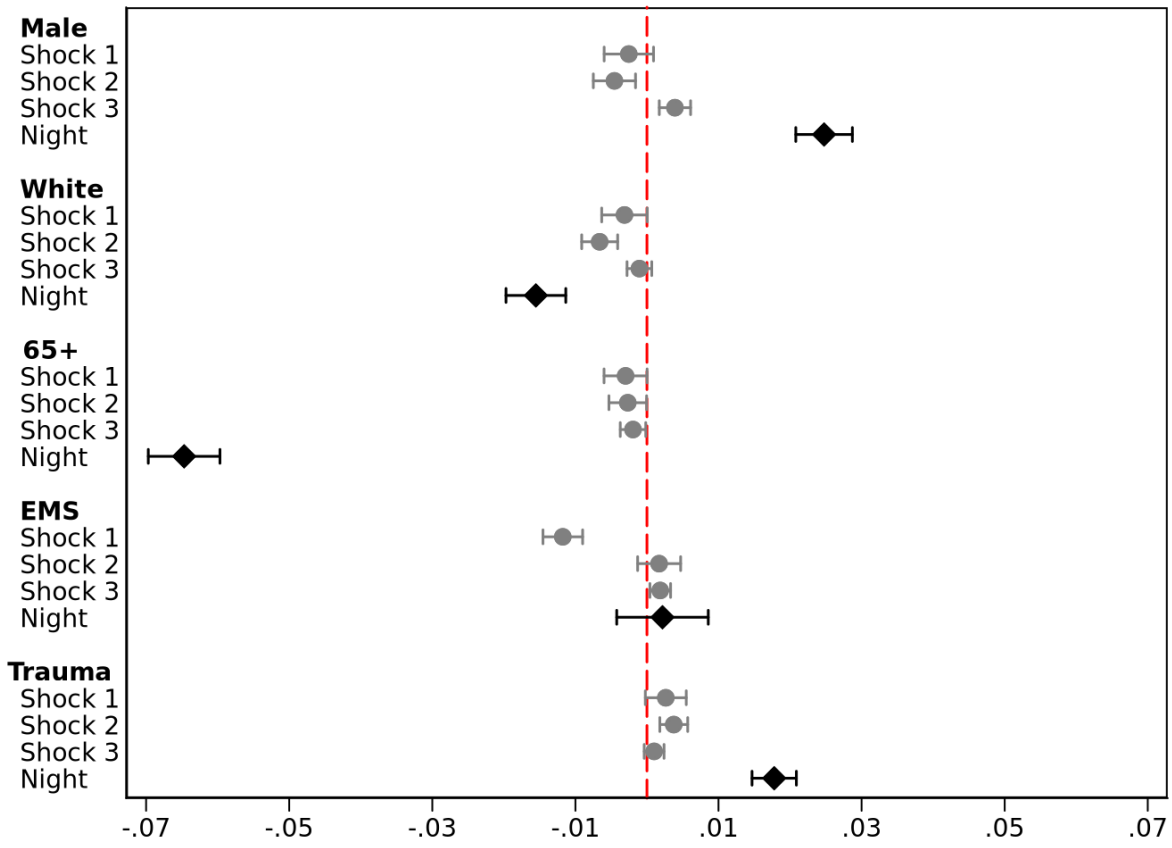
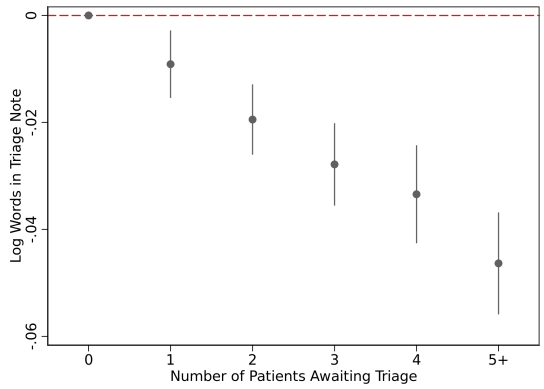
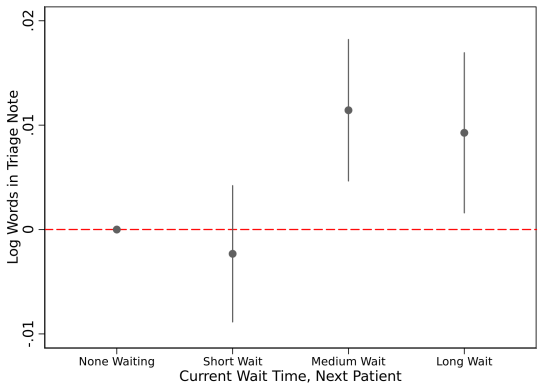


Figure A.2: Balance on Pre-Triage Characteristics

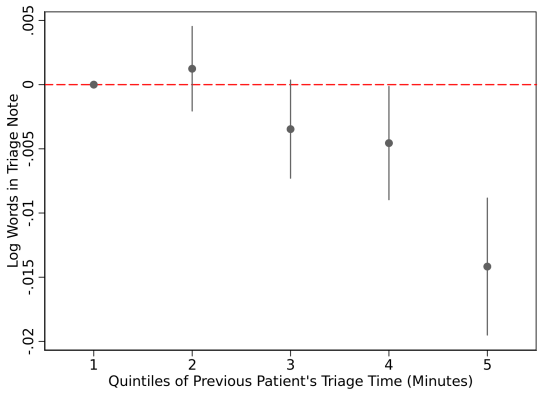
Note: Figure shows estimates from separate regressions of pre-determined patient characteristics on each indicator shown. Points denote point estimates and horizontal bars represent 95% confidence intervals. "Shock 1" is an indicator equal to one for patient visits in the top two bins of the number of patients waiting; "Shock 2" is an indicator for visits in the top two bins of the next patient's wait time; and "Shock 3" is an indicator for visits in the top two bins of the duration of the previous patient's triage. Nighttime arrival (7pm–7am) is included as an endogenous benchmark. Regressions with time-pressure indicators include fixed effects for nurse-shift, visit number, hospital-by-hour, and binned non-focal time-pressure shocks. Regressions for nighttime arrival include fixed effects for nurse-shift and visit number. Standard errors are clustered at the triage nurse level.



(a) Patients Waiting



(b) Next Patient Wait



(c) Previous Triage Duration

Figure A.3: Triage Note Length Under Time Pressure

Note: Figure shows regression estimates of the log number of words in triage notes as a function of time pressure, using three quasi-random sources of time pressure. Panel (a) plots estimates by the number of patients awaiting triage, with 0 as the reference category. The unconditional average note length for patients with 0 people waiting is 3.3 log words. Panel (b) plots estimates by the current wait time of the next patient in line, with "None Waiting" as the reference category. Panel (c) plots estimates by the duration of the previous patient's triage, with the shortest quintile as the reference category. The unconditional average note length for patients following a prior triage in the shortest quintile is 3.2 log words. All regressions control for race (White vs. non-White), gender, a cubic in patient age, and an indicator for whether the patient was skipped over by someone else in the queue. The specification further absorbs nurse-by-12-hour-shift fixed effects, chief complaint fixed effects, arrival mode fixed effects, hospital-by-hour fixed effects, visit number fixed effects, and the non-focal time dimensions. Standard errors are clustered at the triage nurse level.

Table A.3: Quantifying The Effect of Time Pressure (Robustness)

	(1)
Non-Hispanic White	0.059***
	(0.004)
Low Time Pressure	0.006***
	(0.002)
Non-Hispanic White × Low TP	-0.012***
	(0.004)
Observations	1,487,511
Nurses	1,519
R-sq	0.384
Nurse-Shift FE	Yes
Hour-of-day FE	Yes
Chief complaint FE	Yes
Ambulance FE	Yes
Visit number FE	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is an indicator for receiving top triage priority. "Low Time Pressure (Low TP)" equals one in the lowest-pressure state defined by the baseline levels of the three TP proxies (no patients waiting, lowest next-patient-wait bin, and lowest prior-triage-duration bin; see text). Controls include a male indicator, a cubic in age, and clinical urgency covariates (temperature, pulse, respiratory rate, oxygen saturation, systolic and diastolic blood pressure) with missingness indicators. Fixed effects for ambulance arrival, chief complaint, nurse×shift (12h), visit number, and hour of arrival are absorbed. Standard errors (in parentheses) are clustered by triage nurse.

Table A.4: Time Pressure vs. Cumulative Fatigue (First vs. Last Half of Shift)

	(1) Num Waiting	(2) Next-Patient Wait	(3) Prev Duration
High TP × White	0.021***		
	(0.005)		
High TP × White × First Half	-0.002		
	(0.005)		
High TP × White		0.014***	
		(0.003)	
High TP × White × First Half		-0.005*	
		(0.003)	
High TP × White			0.006*
			(0.003)
High TP × White × First Half			0.005
			(0.003)
Observations	1,487,511	1,487,511	1,337,858
Nurses	1,519	1,519	1,484
R-sq	0.363	0.363	0.374
Nurse-Shift FE	Yes	Yes	Yes
Hosp-Hour FE	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes
Chief complaint FE	Yes	Yes	Yes
Visit number FE	Yes	Yes	Yes
Num Waiting FE	Yes	Yes	Yes
Next-Patient Wait FE	Yes	Yes	Yes
Prev Duration FE	Yes	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is an indicator for receiving top triage priority. "High Time Pressure (High TP)" equals one for observations in the top two bins of each time pressure proxy. Each column reports the estimated change in the White–non-White gap under High versus Low time pressure ("High TP × White"), and the difference in that change between the first and second halves of the shift ("High TP × White × First Half"). All specifications include a male indicator, a cubic in age, and absorb fixed effects for nurse×shift, chief complaint, arrival mode, hosp-hour of arrival, visit number, and the binned versions of all time pressure proxies. Standard errors are clustered by triage nurse.

Table A.5: Time Pressure vs. Cumulative Fatigue (Shift Progress Quintiles)

	(1) Num Waiting	(2) Next-Patient Wait	(3) Prev Duration
White × High TP	0.021***		
	(0.004)		
White × High TP		0.012***	
		(0.002)	
White × High TP			0.008***
			(0.003)
Observations	1,487,511	1,487,511	1,337,858
Nurses	1,519	1,519	1,484
R-sq	0.363	0.363	0.374
Nurse-Shift FE	Yes	Yes	Yes
Hosp-Hour FE	Yes	Yes	Yes
Ambulance FE	Yes	Yes	Yes
Chief complaint FE	Yes	Yes	Yes
Visit number FE	Yes	Yes	Yes
Num Waiting FE	Yes	Yes	Yes
Next-Patient Wait FE	Yes	Yes	Yes
Prev Duration FE	Yes	Yes	Yes

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable is an indicator for receiving top triage priority. "High Time Pressure (High TP)" equals one for observations in the top two bins of each time pressure proxy. Each column reports the estimated change in the White–non-White gap under High versus Low time pressure ("High TP × White"). All specifications include a male indicator, a cubic in age, and absorb fixed effects for nurse×shift, chief complaint, arrival mode, hosp-hour of arrival, visit number, and the binned versions of all time pressure proxies. To flexibly account for within-shift drift in triage behavior—such as gradual changes in fatigue, pacing, or baseline decision stringency—each regression also includes quintile fixed effects of shift progress and their interactions with White. Standard errors are clustered by triage nurse.