Waiting to Decide:
Revealed Uncertainty in the Emergency Department

David C. Chan
Stefano DellaVigna
David Card*

March 27, 2026

PRELIMINARY AND INCOMPLETE. PLEASE DO NOT CITE.

Abstract

This paper studies whether the time physicians take to make an admission or discharge decision in the emergency department reveals latent clinical uncertainty. We show that decision time predicts marginality rather than severity: longer decision times are associated with worse outcomes among discharged patients but less severe outcomes among admitted patients. This pattern is consistent with a sequential-information-acquisition view in which physicians devote more time and more auxiliary information-gathering activity to borderline cases near the admission threshold. We also show that decision time contains predictive information beyond rich patient observables and continues to predict clinically meaningful follow-up outcomes several visits into the future. These results suggest that behaviorally generated signals can help identify cases where additional support, observation, or decision assistance is most valuable.

*Card: UC Berkeley and NBER, card@berkeley.edu. Chan: UC Berkeley, Department of Veterans Affairs, and NBER, david.c.chan@berkeley.edu. DellaVigna: UC Berkeley and NBER, sdellavi@econ.berkeley.edu.

1 Introduction

A central problem in economics is how to infer latent information from observed behavior. In many settings, actions are informative not only because they determine outcomes, but because they reveal what decision-makers know and what the econometrician does not. A large literature on private information, screening, and adverse selection uses this logic to study how hidden types distort allocation and how equilibrium behavior reveals otherwise unobserved risk, quality, or willingness to pay (Akerlof, 1970; Spence, 1973; Rothschild and Stiglitz, 1976). In most of that literature, the revealing behavior is the *choice itself*: insurance coverage, contract selection, annuitization, educational signaling, or some related action that sorts agents monotonically by an underlying type. This paper studies a different behavioral object: the *time taken to make a decision*. We ask whether physicians' decision time in the emergency department reveals latent clinical uncertainty not fully captured by standard patient observables.

Emergency department admission decisions are a natural setting for this question. Physicians must decide, often under time pressure and with incomplete information, whether a patient should be admitted to the hospital or safely discharged home. For some patients, the appropriate disposition is clear. Patients who are obviously unstable or severely ill are admitted quickly; patients who are obviously well enough to go home are discharged quickly. The difficult cases are those near the boundary between admission and discharge, where additional observation, testing, or consultation may be valuable and where the risk of an unsafe discharge or unnecessary admission is greatest. These cases are economically and operationally important: they are where mistakes are most likely, where scarce expert attention may be most valuable, and where hospitals are most likely to benefit from better targeting of diagnostic support, second opinions, or observation-unit capacity. Yet most available data systems are much better suited to measuring *severity* than to identifying *marginality*—that is, closeness to the admit-discharge threshold.

This distinction between severity and marginality is central to the paper. Severity refers to the expected level of illness, resource use, or adverse outcomes. Patient age, comorbidity burden, triage acuity, and abnormal vital signs are all natural proxies for severity. Marginality, by contrast, refers to how close a case is to the decision boundary. A case can be severe but not marginal if admission is obvious; it can be mild but not marginal if discharge is obvious. The marginal cases are those for which the physician's posterior assessment lies close enough to the threshold that further information could change the disposition decision. If the key organizational problem is to identify where decisions are hardest and where misclassification risk is highest, then marginality may be the more relevant object. Existing decision-support systems, however, typically focus on severity: they flag high-risk patients or identify decisions that appear inconsistent with the model's predicted risk. Our premise is that this misses an important dimension of uncertainty that may instead be revealed by physician behavior.

We study whether *decision time*—the time physicians spend before reaching a disposition decision—helps identify that dimension. A natural interpretation is that physicians use time to acquire and process information. They observe the patient's trajectory, review tests, order additional diagnostics, call consults, and update their beliefs about whether inpatient care is warranted. This interpretation is consistent with standard models of sequential information acquisition and stopping, in which the value of another signal is highest near a decision threshold (Wald, 1947; Fudenberg et al., 2018). At the same time, that interpretation is not obvious *a priori* in the emergency department. Longer decision times could instead reflect workflow frictions, congestion, interruptions, or distractions. They could reflect that physicians prioritize the sickest patients, implying that shorter decision times predict severity rather than ease. Or they could simply proxy for patient complexity in a broad sense, in which case longer decision times would be associated with worse outcomes among both admitted and discharged patients. Distinguishing among these possibilities is therefore central to both the paper's interpretation and its practical relevance.

Our main finding is that decision time predicts a dimension of patient state that is distinct from standard patient characteristics. Observable patient characteristics such as age, comorbidities, and abnormal vital signs primarily predict *severity*: Among discharged patients, they are associated with worse downstream outcomes; among admitted patients, they are associated with longer lengths of stay and higher mortality. Decision time instead predicts *marginality*. Among discharged patients, longer decision times are associated with higher bounceback rates and higher mortality. Among admitted patients, longer decision times are associated with shorter hospital stays and lower mortality. As decision time rises, admitted and discharged patients increasingly resemble one another on *ex post* outcomes. This opposite-sign pattern is the paper's core empirical finding. It is difficult to reconcile with the view that decision time is simply a proxy for sickness. It is much more consistent with the idea that longer decision times identify cases closer to the admit-discharge cutoff.

We document several additional facts that reinforce this interpretation. First, longer decision time is associated with a higher probability of admission. In the emergency department, discharge is the default, so if difficult cases remain unresolved for longer, later decisions should tilt toward admission. Second, excess decision time and other information-acquisition behaviors peak where residual uncertainty appears greatest. In particular, excess decision time and consult activity are highest when the predicted probability of admission is intermediate, roughly where a physician appears most uncertain. Third, decision time continues to predict marginality even after controlling for a rich set of patient covariates, indicating that physician behavior contains information not fully recoverable from standard observables alone. Fourth, the predictive content of decision time is persistent: among patients sent home, higher decision-time cases continue to predict clinically meaningful follow-up utilization at subsequent visits, whereas more transient indicators, such as triage acuity or abnormal vital signs, appear to be much more fleeting signals of current severity. Finally, in a subsample of patients with chest X-rays, a text-based admission index constructed from radiology reports supports the same interpretation, suggesting that high-decision-time cases are difficult because they contain clinically meaningful ambiguity visible to human experts, not merely because they are administratively slow.

Our interpretation is therefore explicitly equilibrium-based. We do not treat decision time as an exogenous treatment whose causal effect on outcomes we seek to estimate. Instead, decision time is a behaviorally generated statistic that reflects how physicians respond to uncertainty and constraints. In that sense, the paper is closer in spirit to the literatures on private information, screening, and adverse selection than to reduced-form papers on delays. The key idea in those literatures is that equilibrium behavior can reveal hidden types or latent states. Here, however, the relevant behavior is not the action chosen but the *time preceding the action.* That difference matters. In much of the empirical asymmetric-information literature, the revealing behavior sorts agents monotonically by a one-dimensional notion of risk or expected loss. Decision time need not be monotone in severity: both obviously severe and obviously non-severe patients may be decided quickly, while the longest deliberation occurs near the threshold. Thus, decision time reveals *marginality rather than severity.* In this sense, the paper extends the revealed-information logic of the asymmetric-information literature from choices that sort types to deliberation times that locate cases near an action boundary. More generally, the paper can also be viewed as an outcome-test exercise, in that we ask whether a behavioral signal generated before disposition—decision time—predicts *ex post* outcomes in a way that reveals residual private information beyond observable patient severity.

The paper contributes to three literatures. First, it contributes to the economics of private information, screening, and adverse selection. Since Akerlof (1970), Spence (1973), and Rothschild and Stiglitz (1976), economists have emphasized that hidden information can distort allocation and that equilibrium behavior can reveal otherwise unobserved types. Empirically, a large literature studies these questions by examining whether observed choices are systematically related to *ex post* risk or outcomes (Chiappori and Salanié, 2000; Finkelstein and Poterba, 2004; Finkelstein and McGarry, 2006; Cohen and Siegelman, 2010; Einav et al., 2010; Einav and Finkelstein, 2011). We contribute by shifting attention from the action chosen to the time taken before the action is chosen, and by showing that this behavior reveals marginality rather than severity. This is a different revealed object and, in settings where mistakes concentrate near thresholds, a potentially more policy-relevant one.

Second, the paper contributes to the economics of information acquisition and sequential decision-making. Classical work on sequential analysis studies when decision-makers optimally continue sampling information rather than stop and act (Wald, 1947). More recent work in economics develops tractable models of costly sequential information acquisition and stopping (Fudenberg et al., 2018), and a growing empirical literature studies decision time in laboratory (Chabris et al., 2008; Clithero, 2018; Frydman and Krajbich, 2022) and field environments (Alós-Ferrer et al., 2021; Cotet et al., 2025; Card et al., 2024). Our findings are consistent with that logic in a high-stakes field setting: physicians appear to spend more time and acquire more auxiliary information precisely in cases where the posterior is closer to the admission threshold.

Third, the paper contributes to work on physician decision-making, healthcare productivity, and operational targeting. Unlike papers that study the causal effects of congestion, delay, or peer pressure on emergency department performance (Chan, 2018; Silver, 2021), we treat decision time as a realized equilibrium object rather than as a treatment. The policy implication is correspondingly different. A hospital system seeking to improve disposition decisions can identify severe patients, but severity is not the same as marginality. Our results suggest that behaviorally generated signals, such as decision time, may help identify the subset of cases where second opinions, observation-unit placement, decision support, or additional diagnostics are most valuable. The exact policy implications depend on the mechanism: If longer decision times primarily identify hard cases, they point toward decision support, consultation, or additional diagnostics; if they instead reflect operational strain or workflow bottlenecks, they point toward capacity expansion or process redesign. However, either way, decision time may help identify encounters where intervention is most valuable. More broadly, the paper highlights a general organizational problem: when scarce expert attention is endogenous, revealed behavior may help identify where uncertainty is greatest and where additional resources have the highest marginal value.

The remainder of the paper proceeds as follows. Section 2 presents a conceptual framework and derives the empirical predictions. Section 3 describes the institutional setting, the data, and the construction of decision time. Section 4 presents the main results on admission probability, excess decision time, consult use, and downstream outcomes. Section 5 examines the incremental information in decision time beyond patient covariates, the radiology-text validation, and the persistence of its predictive content. Section 6 assesses the persistence of decision time in predicting marginality across future visits. Section 7 concludes.

2 Conceptual Framework

Consider a doctor who has to make a binary decision $D \in \{0, 1\}$, whether to admit a patient who shows up at the emergency room. Admission ($D = 1$) is the optimal decision for patients with symptom severity q above a certain threshold, whereas for those with less severe symptoms, it is optimal to send them home ($D = 0$). The doctor observes a noisy signal s of the patient's seriousness q and can decide whether to obtain more signals or to decide based on the information at hand. Obtaining more signals is costly, in terms of time or money, so there is a trade-off between obtaining more information at a cost and acting sooner with less precise information.

There are multiple versions of this type of model with sequential sampling, all with similar solution properties. A first class of models, e.g., Ratcliff and Rouder (1998); Fudenberg et al. (2018), has decisions in continuous time, as in the Drift-Diffusion Model: individuals receive continuous signals and decide when to stop sampling. Other models, e.g., Card et al. (2024), have a discrete-time set-up that simplifies the sequential sampling into a 2-period model: individuals observe a first signal and then decide whether to draw a second signal at a cost.

The key predictions from these models are similar across the different versions. We describe them here, referring to the discrete-time 2-period model in Card et al. (2024), but we stress that the results are parallel in the continuous-time sequential sampling models.

Consider the simple case in which a doctor has received a first signal s_1 of the seriousness of the condition of the patient, for example, by doing a first set of testing, or by taking a set of vitals and doing a physical examination. The doctor has to decide whether to acquire, at a cost of c, a second signal s_2 which will allow for a more precise diagnosis. As it turns out, the optimal solution is a simple threshold rule. If the first signal is low enough, that is, $s_1 < S_L$, the doctor will send the patient home without acquiring further information. If the first signal is high enough, that is, $s_1 > S_H$, the doctor will admit the patient overnight without extra testing. It is for intermediate values of the signal, for $S_L \leq s \leq S_H$, that the doctor will pay for extra testing, since in this central region of uncertainty, the value of extra information is higher, as it is more likely to change the ultimate decision. Instead, if the initial signal is extreme, acquiring a second signal is unlikely to change the final decision, and it thus does not make sense to pay the cost to acquire it. The same logic applies to continuous-time drift-diffusion models, in which the agent stops acquiring information if the signal falls below a lower threshold or rises above an upper threshold. This decision-making rule directly implies the prediction below.

Prediction 1. *For discharged individuals, those discharged right away (without acquiring extra information) are healthier than the ones discharged later (after acquiring extra information). For individuals admitted overnight, instead, the opposite is true: the ones admitted right away (without acquiring extra information) are sicker than the ones admitted later (after acquiring extra information).*

Prediction 1 is implied by the sequential decision-making outlined above: the individuals discharged right away are the ones for whom the doctor received an initial, more positive signal about their health, and thus, on average, will tend to be healthier. Conversely, those admitted right away, without an additional signal, are those who are sick enough that the doctor does not deem it necessary to obtain additional information before making the admission decision. This channel is the *marginality* channel: only on the marginal cases does the doctor decide to acquire extra information.[1]

A second prediction concerns the information acquisition. Suppose we have a set of proxies X that predict the admission decision D, such as age, prior comorbidities, and initial vitals in the ED. We can use such proxies to construct $P(D = 1|X)$, that is, the probability of admission given the proxies. Then, as Card et al. (2024) shows, the following holds.

Prediction 2. *The probability of acquiring an extra signal (e.g., a consult, extra testing, or just longer deliberation) is inverse U-shaped in* $P(D = 1|X)$*, and it peaks at* $P(D = 1|X) = 0.5$*.*

That is, the acquisition of extra information (which in a 2-period model is the same as longer decision time) is lowest in cases in which the ex-ante probability of admission is very high or very low, and is highest when the ex-ante probability of admission is at 50-50. This is once again an application of the principle of marginality: It is the cases that are ex ante closest to a toss-up, the ex-ante marginal cases, where acquisition of information is most valuable.

A final prediction also follows from the marginality.

Prediction 3. *The probability of admission* $P(D = 1)$ *converges towards 50-50 as the decision time increases.*

In our setting, the average probability of admission is below 50-50, at about 20 percent. For cases decided quickly (short decision time), most will be discharged, and the probability of

[1]A caveat to this prediction is that, in principle, there is a counterveiling learning force. If the second signal is strong enough, it is in principle possible that the prediction above can be reversed in some cases, as Card et al. (2024) points out. In practice, though, we are not aware of any instances in which this is the case.

admission will be fairly low. As the decision time increases, the cases become marginal cases that the doctor needs to ponder, and they tend to be closer to a 50 percent probability of admission (approaching it from below).

In the next sections, we provide evidence supporting these predictions, with a particular focus on Prediction 1, which shows that decision time can predict the marginality of a medical case and, when conditioned on the decision taken, the severity of a case. We then present evidence on information acquisition consistent with Prediction 2 and also provide evidence for Prediction 3.

3 Setting and Data

Patient Care and Disposition Decisions. Emergency department patients are assigned to physicians, and each patient is typically managed by a single physician over the course of the visit. A crucial—arguably the central—decision for ED physicians is whether to admit or discharge the patient. Admitted patients receive further monitoring, diagnostic workup, or treatment in the hospital for a period of time, with sicker patients tending to stay longer. Discharged patients are sent home or to a residential facility, typically with instructions for outpatient follow-up or return precautions if their condition worsens. Unscheduled ED return visits for discharged patients (which we label as "bouncebacks") are commonly treated in policy and quality-improvement settings as relevant signals of discharge quality and transition quality, even though they are an imperfect measure of inappropriate discharge (Agency for Healthcare Research and Quality, 2014; Centers for Medicare & Medicaid Services, 2023).

This disposition decision is often made under uncertainty. In many cases, the patient does not leave the ED with a fully resolved diagnosis. Instead, the physician must decide whether the information produced during the visit is sufficient to conclude that discharge is appropriate, or whether the remaining possibility of a serious condition is substantial enough to warrant admission. Physicians make this decision using both structured patient information and information generated during the visit itself. Relevant patient characteristics include measures of severity and clinical risk such as triage acuity, comorbidities, age, and abnormal vital signs. These variables are useful inputs into the physician's assessment, but they do not by themselves resolve the admission decision in all cases.

Data and Sample. Our primary sample consists of roughly 30 million emergency department (ED) visits that occurred between January 2011 and December 2023 in the Veterans Health Administration (VHA) health care system. Our data contain rich versions of these patient characteristics from the VHA electronic medical record. In particular, we observe demographics (e.g., age, sex, race, marital status), comorbidities, prior utilization, triage information, and vital signs, along with the timing of care and disposition. These data allow us to construct flexible predictions of the probability of admission using observables alone. We later use these predictions both descriptively and as inputs into tests motivated by our conceptual framework.

We restrict our sample as follows. First, we drop visits where a patient's arrival time is later than the disposition or departure time. We further assume that extremely long ED visits of a week or longer were recorded in error and drop these observations. Overlapping ED visits by the same patient are treated as a single visit. We also drop visits where a physician is not recorded. For our main analyses, we further restrict our sample to visits in which the patient had not visited the ED in the past 30 days. Finally, we omit visits with missing demographic, mortality, or decision time data. Visits by patients younger than 18 and older than 100 are dropped, as are observations where the patient's vitals are outside of possible ranges. Appendix Table A.1 presents our sample selection process and the sample size at each step.

For analyses conducted in section 5.4, we further restrict our sample to ED visits where a single- or 2-view chest X-ray was taken during the visit.

Measuring Decision Time. We calculate our main variable of interest, decision time, using the sample of ED visits described above. The object is intended to capture the amount of physician time effectively allocated to a patient's case prior to disposition. Because we do not directly observe minute-by-minute physician effort, we infer it from the physician's patient census over time. Specifically, we assume that when a physician is responsible for multiple patients during a given hour, each patient receives an equal share of that physician's care during that hour.

Formally, for patient i, decision time is

$$T_i = \sum_{t=\underline{t}(i)}^{\overline{t}(i)} \frac{60}{N_{j(i),t}},$$

where $\underline{t}(i)$ is patient i's arrival hour, $\overline{t}(i)$ is the hour in which patient i receives a disposition decision, $j(i)$ is the assigned physician, and $N_{j,t}$ is the number of patients under physician j's care during hour t.

This construction is in the spirit of the "workload-adjusted length of stay" measure in Chan (2018). The logic is similar: Raw elapsed time is an imperfect measure of physician attention because an hour represents less case-specific time when the physician is simultaneously managing more patients. Our decision-time measure rescales calendar time by the physician's contemporaneous workload.[2] The measure should be interpreted as an approximation to cumulative physician attention rather than raw elapsed time. The equal-sharing assumption is, of course, an approximation, but the measure relies solely on systematically observed information and provides a way to convert elapsed time into an estimate of effective case-specific physician time.

Interpreting Decision Time. Although decision time may reflect information acquisition, that interpretation is not obvious *a priori*. In the ED, longer delays are often discussed in terms of congestion, workflow bottlenecks, interruptions, distractions, or broader operational strain rather than deliberate information production (Hoot and Aronsky, 2008; Morley et al., 2018; American College of Emergency Physicians, 2016, 2024). Related work in operations research and economics also shows that provider speed and service time respond to workload, peer pressure, and scheduling incentives (Kc and Terwiesch, 2009; Chan, 2018; Silver, 2021). A separate possibility is that physicians prioritize the sickest patients, in which case short decision times could be associated with greater severity rather than ease. More generally, decision time could simply proxy for patient complexity in a broad sense.

[2] For example, suppose patient i arrives at 10:15am and receives a disposition decision at 1:45pm. If the assigned physician has 2, 4, 3, and 1 patients during the 10am, 11am, 12pm, and 1pm hours, respectively, then $T_i = 60/2 + 60/4 + 60/3 + 60/1 = 125$ minutes. Although more than three clock hours elapse between arrival and disposition, the implied physician time devoted to the case is 125 minutes once physician workload is taken into account.

For these reasons, the central empirical question is what kind of underlying patient state decision time appears to reveal. Different interpretations imply different patterns of downstream outcomes. If longer decision times primarily reflect information acquisition or prioritization that improves sorting, then they may be associated with safer discharges and more appropriate admissions, and hence with better outcomes on both sides of the disposition decision. If they instead reflect harmful delay that worsens patient health per se, then they may be associated with worse outcomes for both discharged and admitted patients. By contrast, if longer decision times identify cases near the admit/discharge threshold, then they should predict a *marginality* pattern: worse outcomes among discharged patients, but less severe admitted patients, such as shorter inpatient stays or lower mortality.

4 Decision Time and Patient Outcomes

In this section, we present evidence regarding Prediction 1 outlined above—Does decision time predict the seriousness of a health condition as predicted by the marginality hypothesis? In particular, Prediction 1 suggests that among individuals discharged after an ED visit, the seriousness of the condition should be higher for those for whom the doctor took longer to decide, as they are more likely to be marginal cases, compared with those discharged quickly.

Bouncebacks. We plot a binned scatter plot of the probability of 7-day bouncebacks as a function of the decision time measure outlined above. Bouncebacks are defined as ED visits within the VHA network within 7 days of the initial ED visit. The bin scatter controls for provider-year-month fixed effects, so the variation is within that group. Panel A of Figure 1 shows that decision time is indeed strongly associated with a higher incidence of bouncing back. A higher decision time, from 32 to 64 minutes, predicts a higher bounceback rate, from 7 to nearly 8 percentage points. The effect is monotonic and close to linear (given that the x-axis is on a log scale).

Length of Stay. The complementary prediction for admitted patients is that those admitted after a longer decision are less seriously ill. For hospitalized patients, we use the length of stay (in natural log units) as a measure of illness severity. As Panel B of Figure 1 shows, there is a strongly monotonic and fairly linear relationship between the decision time and length of stay, in the direction predicted by the model: Longer decision times predict shorter lengths of stay. Given that both axes are in log units, we can interpret the slope as an elasticity: For each doubling of decision time, the length of stay decreases by about 8 log points.

While bouncebacks and length of stay are important health outcomes in their own right, we can also examine whether decision time predicts mortality among patients, measured within 30 days of the ED visit.

Mortality. Figure 2 shows 30-day mortality rates, with decision time on the x-axis, separately for individuals admitted and for those released from the hospital. The two binned scatter plots are run separately and then juxtaposed in a single graph. As one can imagine, there is a large difference in mortality between the two groups: while the average 30-day mortality for patients who are released is about 0.5 percent, for patients who are admitted, it is around 4 percent. Given this difference in level, we turn to the slopes of the two lines with respect to decision time. Decision time is a strong positive predictor of mortality among individuals released from the ED. There is also a relationship, of the opposite sign, for individuals admitted: longer decision time for that group is associated with lower mortality. While this second slope is less apparent, it is highly statistically significant. The evidence from mortality is thus also consistent with Prediction 1.

These figures thus support the marginality prediction we outlined. This evidence, though, does not speak to whether this pattern is distinctive of decision time rather than other patient characteristics. It is possible that this predictive pattern also holds for other patient-level variables of medical interest, such as age, whether the patient has abnormal vital signs at the start of the ED visit, and whether the patient has Elixhauser comorbidities indicating a more serious condition.

Decision Time versus Health Controls. In Panel A of Table 2, we estimate the impact of decision time on the four key outcomes above — bounceback, length of stay, mortality among discharged patients, and mortality among admitted patients. We first present the results using only provider-year-month fixed effects for the sample used for the binscatter plots above, then

Figure 1: Downstream Signals of Marginality

A. Bounceback

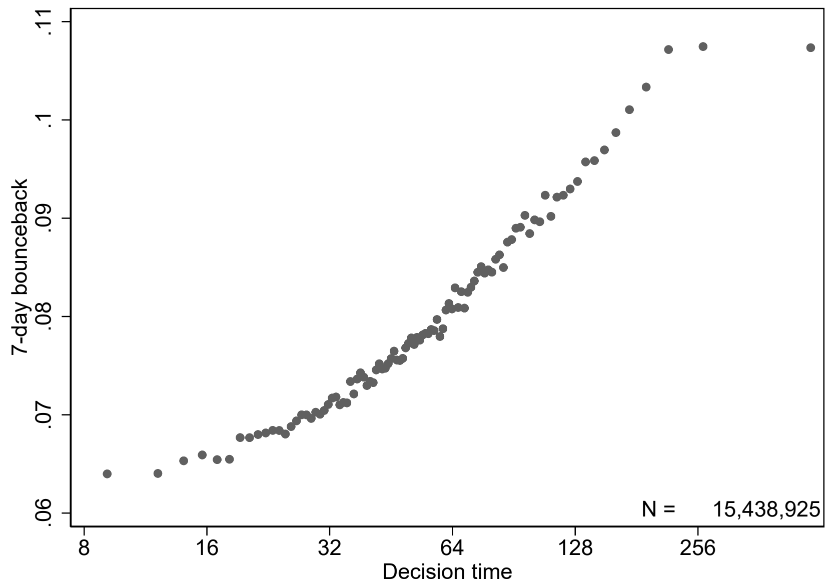

B. Length of Stay

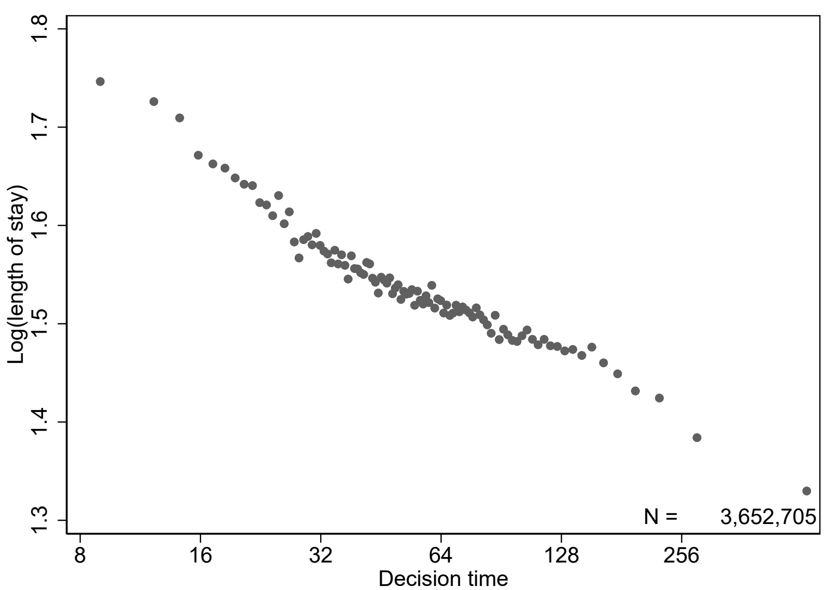

Notes: This figure shows binned scatter plots of patients' outcomes on the y-axis against decision time on the x-axis. We control for provider-year-month fixed effects. Panel A reports results for 7-day bounceback for discharged patients; Panel B reports results for log length of stay for admitted patients. Bins of decision time are formed within the station.

Figure 2: Mortality

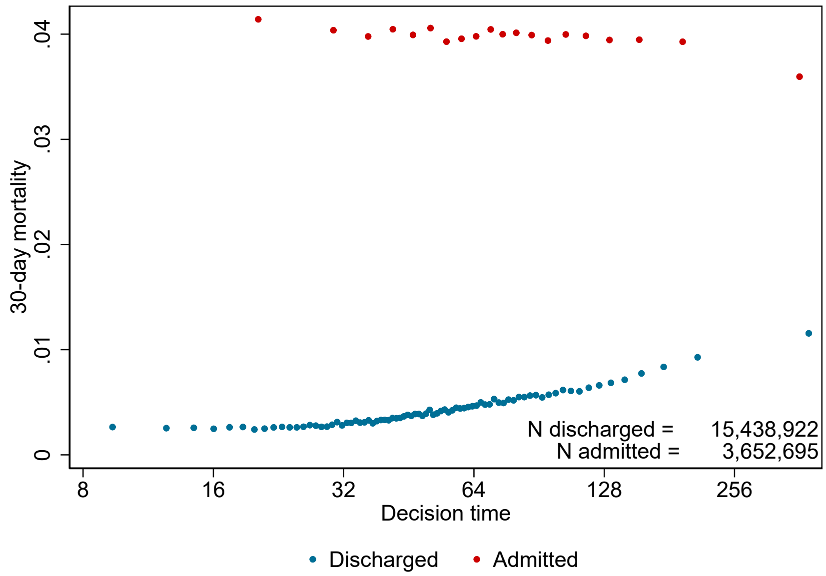

Notes: This figure shows a binned scatter plot of mortality on the y-axis against decision time on the x-axis, split by discharged and admitted patients. We control for provider-year-month fixed effects. Bins of decision time are formed within the station.

for the restricted with non-missing values of the key health controls, which we then add in a third column.

The evidence reveals several noteworthy patterns. First, the health controls are strong predictors, as one would expect, increasing the R-squared by 10 to 20 percent. Second, decision time remains strongly predictive of health outcomes in these multivariate regressions, as in the univariate scatter plots displayed above. In fact, while the coefficient on bouncebacks is attenuated by 20 percent, the coefficients on length of stay and mortality for the admitted are larger after adding the health controls. Third, none of the health coefficients exhibit behavior consistent with the pattern we find for the decision-time variables. That is, a suite of other patient characteristics—age, vitals, and Elixhauser comorbidities—predict health outcomes in the same direction for both admitted and discharged patients. Older age, abnormal vitals, and more comorbidities predict a higher likelihood of bouncebacks, longer length of stay, and higher mortality among both those discharged and those admitted. Only the decision time variable predicts outcomes in the opposite direction, as per the marginality hypothesis, for both admitted and discharged cases.

To visualize the patterns above, we plot the coefficient estimates from multi-variate specifications as scatter plots in Figure 3. For expositional purposes, we model decision time, age and the Elixhausers in quintiles, with the first quintile as omitted variable. We also display controls for the case of 1 abnormal vital or 2+ abnormal vitals. Panel A of Figure 3 shows the impact of various variables on length of stay (x axis) and bouncebacks (y axis), while Panel B shows the impact on mortality for the admitted (x axis) versus those discharged (y axis). The figure shows clearly that all the traditional health variables are in the North-East quadrant, that is, variables that worsen outcomes for the discharged also worsen outcomes for the admitted patients. It is only decision time that has the marginality property that it is associated with worse outcomes for the discharged but with better outcomes for the admitted.

Heterogeneity. Above, we showed that decision time is associated with health outcomes with a unique pattern, consistent with the marginality hypothesis. How consistent is this pattern, though, across different medical categories? To address this, we code chief complaints, and in Figure 4A. we plot, for each chief complaint, the coefficient on decision time for a specification of bouncebacks versus a specification of length of stay. Both specifications include the battery of health controls used above. The scatter plot shows that across all chief complaints, the key pattern holds: decision time is associated with higher bouncebacks but a shorter length of stay. In Figure 4B., we show that a similar pattern holds for the mortality outcomes. We thus conclude that the pattern observed above is pervasive across different health conditions.

5 Mechanisms and Other Model Predictions

In the previous section, we presented evidence supporting the model's marginality prediction (Prediction 1). Namely, we showed that decision time predicts a range of important health outcomes, even after controlling for traditional health determinants, in a direction consistent with high decision time identifying more marginal patients, as predicted by the model of sequential information acquisition. In this Section, we present evidence for additional model predictions and mechanisms.

Figure 3: Marginality vs. Severity

A. Bounceback and Length of Stay

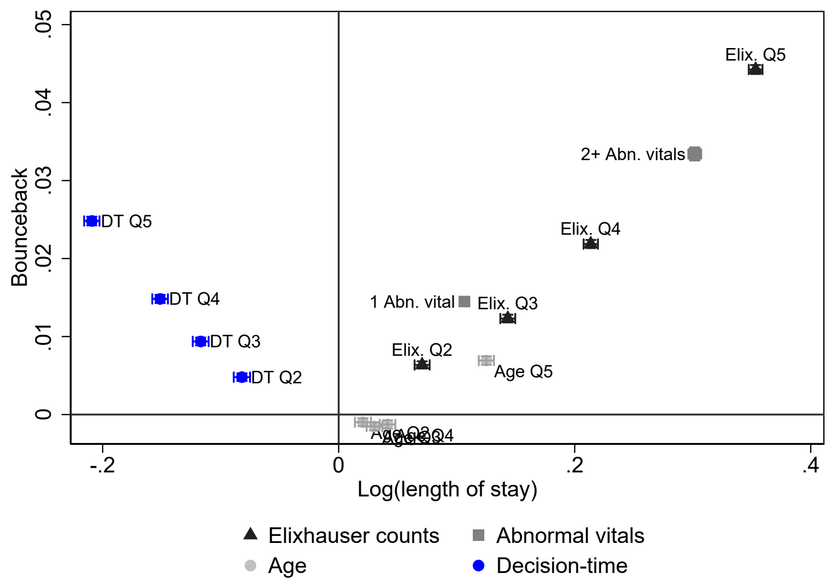

B. Mortality

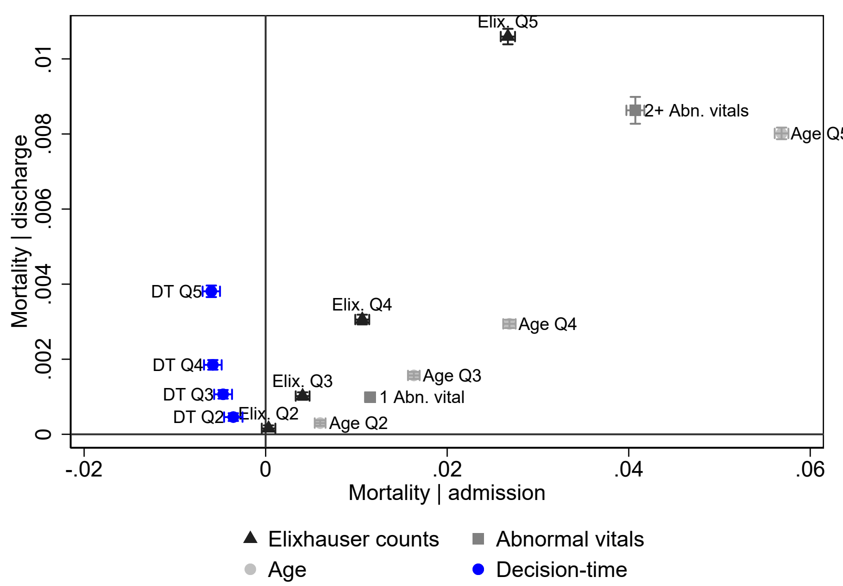

Notes: This figure shows coefficients from regressions of patient outcomes on quintiles of decision time and patient characteristics, controlling for patient-year-month fixed effects. Quintiles of decision time, patient age, and counts of Elixhauser comorbidities are formed within the station. Standard errors are clustered by provider-year-month, and whiskers display 95% confidence intervals. Panel A displays results for bounceback and length of stay outcomes; Panel B displays results for mortality.

Figure 4: Results by Patient Chief Complaints

A. Bouncebacks and Length of Stay

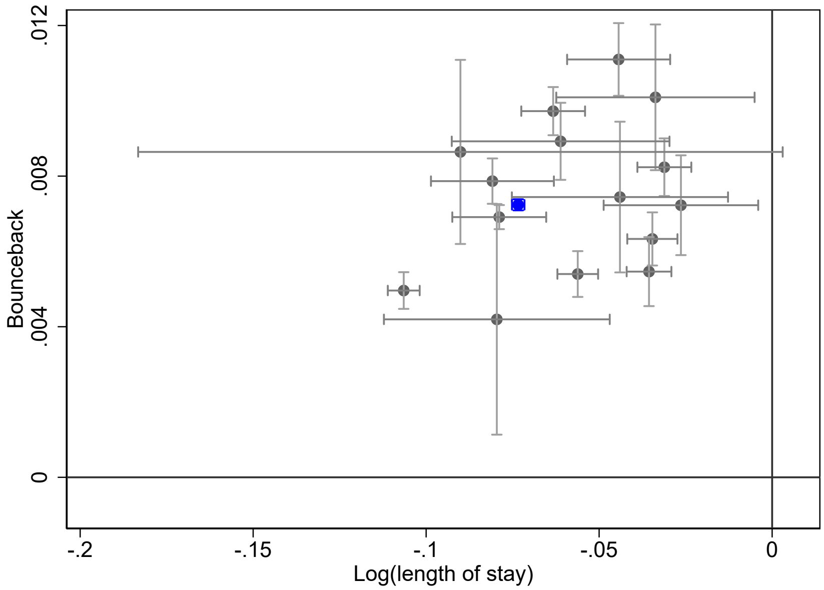

B. Mortality

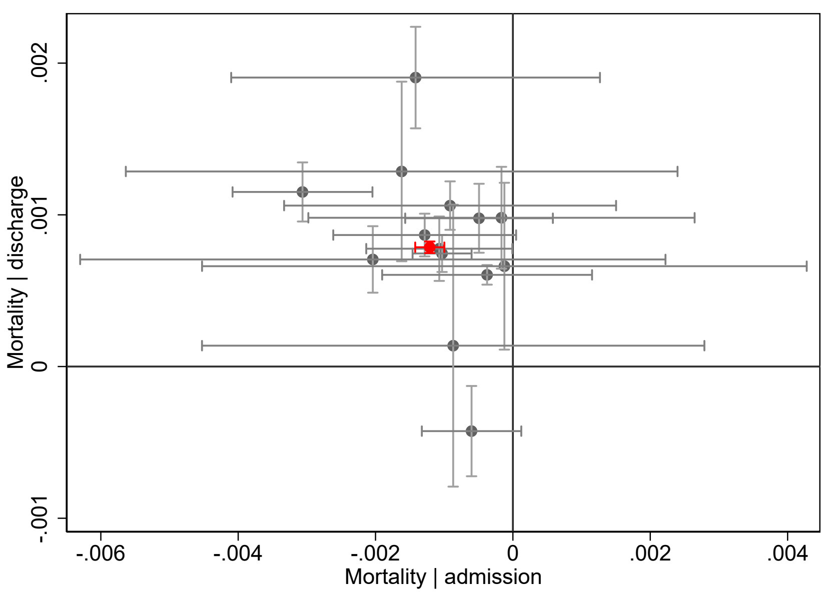

Notes: This figure shows results from regressions of patient outcomes on log decision time. Each point represents the coefficient on log decision time for patients grouped by similar chief complaints. We control for Elixhauser comorbidity indicators, demographics, vital signs, VHA visits in the past year, and patient-year-month fixed effects. Standard errors are clustered by provider-year-month and whiskers display 95% confidence intervals. Panel A displays results for bounceback and length of stay outcomes; Panel B displays results for mortality.

5.1 Information Acquisition

If doctors acquire information along lines consistent with optimal information acquisition, we should expect them to spend more time with patients in cases where the return to acquiring such information is higher. As we outlined above in Proposition 2, this is when ex ante a case (for admission) is most uncertain, that is, the probability of admission is close to 50-50. It is in such cases that the extra time and effort on a case are most likely to change the admission decision.

To test this prediction, for each ED visit, we create a prediction of the admission decision using patient characteristics, such as age and Elixhauser comorbidities, as well as vitals collected at the beginning of the ED visit. Importantly, the predictors we consider for admission probability should include only information available to the doctor at the beginning of the visit. We then use this prediction to test whether information acquisition and decision time are inverse U-shaped with respect to the predicted admission decision.

Our most direct proxy for information acquisition is consults. Specifically, the data record whether the ED physician obtained a consult on the case from another doctor or specialist before making a decision about admission. We use an indicator for one or more consults as the outcome variable. Panel A of Figure 5 displays a binned scatter plot of the probability of one or more consults as a function of the probability of admission. The figure shows that the probability of admission, as determined using the variables outlined above, ranges from near zero to 0.6 or higher. This range is important for testing the model's prediction. Turning to the key prediction, the figure shows that the probability of a consult increases sharply for predicted probabilities between 0 and 0.4, is quite flat between 0.4 and 0.6, and decreases for probabilities above 0.6. The figure is thus clearly consistent with the model's inverse-U-shape prediction: Doctors are more likely to seek additional opinions on cases in which they are (ex ante) most likely to be on the margin of the admission decision.

We can also produce parallel evidence for this prediction relating the probability of admission to our measure of decision time (as opposed to consults). To do so, we need to address the issue that, for admitted patients, the recorded "decision time" may include time spent on additional tasks required for admission (e.g., filling out paperwork, coordinating with the inpatient team). We thus adjust decision time for discharged and admitted patients as follows. For discharged patients, we subtract the mean decision time of the bottom 1% of patients by predicted admission, effectively removing a "baseline" decision time typical of discharged patients. Similarly, for admitted patients, we subtract the mean decision time of the top 1% of patients by predicted admission. We label this measure "excess" decision time. In Panel B of Figure 5, we show a binned scatter plot of this excess decision time measure as a function of the probability of admission. Once again, the shape is inverse U-shaped, peaking at an admission probability of about 0.4, close to the model's prediction of a peak at 0.5.

5.2 Decision Time and Admission

A third prediction of the model is a direct test of the "marginality" of patients with high decision times. Cases that are decided fast are likely to be easy cases, and, given that admissions are only 20 percent of cases, overwhelmingly cases in which patients are discharged. The longer it takes the doctor to decide, the more likely the case is to be marginal, and thus the closer the division between admitted and discharged cases is to 50-50.

In Figure 6, we plot a bin scatter of the share of cases admitted as a function of our decision time variable. As the plot shows, there is indeed a very strong positive slope of the admission rate with respect to decision time. Indeed, for long decision times, the probability of admission is as high as 0.45. This is economically meaningful, given that our rich set of covariates yields a wide range of predicted admission probabilities, from systematically low values at short decision times to systematically high values at long decision times.

5.3 Uncommon Diagnoses

As we discussed above, the marginality hypothesis holds that patients with long decision times are typically those with more complicated cases that require additional information from the doctor. If that is the case, we expect that the diagnoses at discharge (or at admission) for patients with long decision times tend to be less common, as the more common ones likely make for an easier assessment.

In Figure 7, we plot a binned scatter plot of a proxy of how common a diagnosis is at dis-

Figure 5: Information Acquisition for Marginal Patients

A. Consults

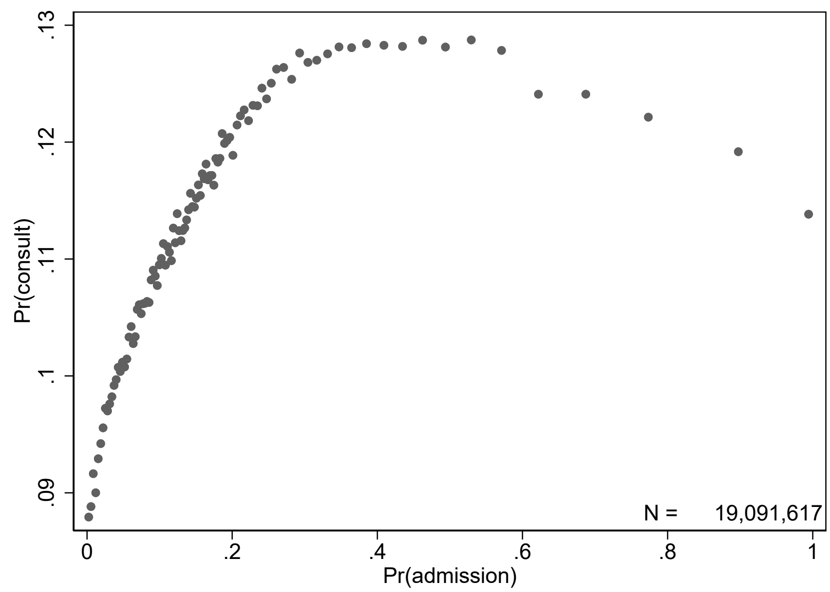

B. Excess Decision Time

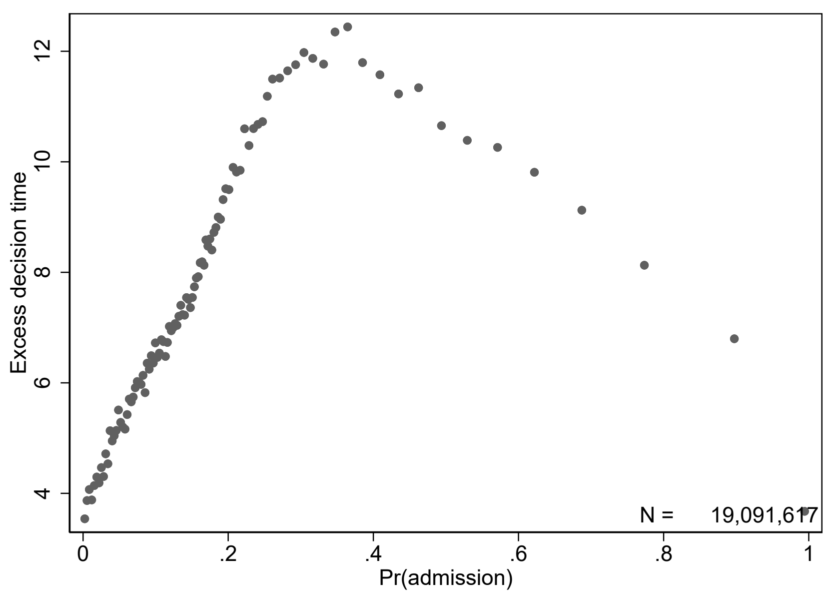

Notes: This figure shows binned scatter plots of the probability of a consult and excess decision time—in panels A and B, respectively—on the y-axis against the probability of admission on the x-axis. We control for provider-year-month fixed effects. Probability of admission is predicted using station-by-station Poisson regressions of admission on patient characteristics (Elixhauser comorbidity indicators, demographics, vital signs, and VHA visits in the past year).

Figure 6: Decision Time and Predicted Marginality

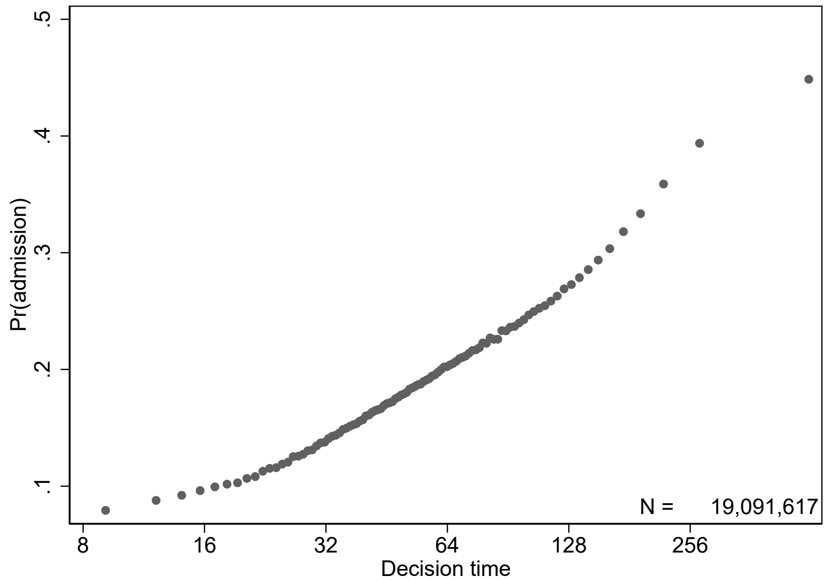

Notes: This figure shows a binned scatter plot of the probability of admission on the y-axis against decision time on the x-axis. We control for provider-year-month fixed effects.

charge, as a function of decision time. The proxy is the log of the conditional probability ratio, in which the probability of a revealed diagnosis, conditional on having a chief complaint, is divided by the probability of that same diagnosis, conditional on not having that chief complaint. The plot shows that for both admitted and discharged patients, the longer the decision time, the more likely the diagnosis is to be less common, as predicted.

5.4 Word Content from Radiology Reports

The analysis so far makes the case that longer decision time is a good proxy for instances in which doctors assess a case as marginal, and the combined evidence is consistent with this interpretation. Still, it would be useful to validate our interpretation with more direct information on the doctor's assessment of the case.

There is no general record of the primary doctors' assessments before they make an admission decision for a particular patient. However, for some patients, we can use notes from interim evaluating physicians to shed light on the patients' assessments before discharge. In particular, for patients with chest X-rays as studied by Chan et al. (2022), we observe the radiol-

Figure 7: Decision Time Predicts Uncommon Diagnoses

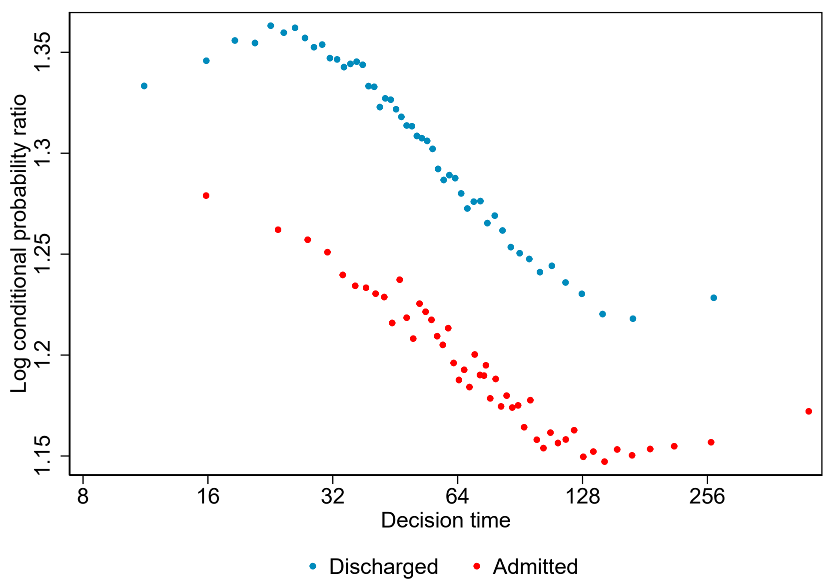

Notes: This figure shows a binned scatter plot of the log conditional probability ratio for a diagnosis given a chief complaint on the y-axis against decision time on the x-axis. The log conditional proability ratio is calculated as $\log(D|C) - log(D|C^c)$ for a diagnosis D and chief complaint C. We control for Elixhauser comorbidity indicators, demographics, vital signs, VHA visits in the past year, and patient-year-month fixed effects.

ogist's note after reading the X-rays and can thus measure the word content to assess whether it is consistent with different degrees of marginality (from the radiologists' perspective) across cases.

We proceed as follows. For all chest X-ray cases in our sample, we process the text (e.g., tokenize, remove stop words, de-stem) and identify common word combinations (i.e., n-grams). We then run a LASSO logit separately by station, predicting admission using the coded words. Panel A of Figure 8 shows an example of some of the most common diagnostic words used as coded for two different stations. The scatter plot shows that words highly predictive of admission at one station are also highly predictive at another station. Words like "fluid" and "pleural effusion" are associated with a high probability of admission, while words such as "no finding" and "mild pulmonary" are associated with lower probabilities. Based on this coding, for each patient, we compute an average text-based probability of admission from the combination of words used in the report. We then ask whether decision time predicts this text-based measure of doctor assessment of the case, and in particular, whether high-decision-time cases are assessed as more marginal.

In Panel B of Figure 8, we plot a bin scatter parallel to our main findings, except that the outcome variable is the summary measure of the words used by the radiologist. In particular, we separately examine admitted patients and those discharged home to assess how longer decision times are associated with textual assessments indicating marginality. We find a remarkably parallel picture to the benchmark one in Figures 1 and 2. That is, for patients ultimately discharged, the longer they spend in the ED, the more serious the wording in the radiologist's report. Conversely, for patients ultimately admitted, the opposite is true: Longer decision time indicates coding of less serious symptoms. Furthermore, we see strong evidence of convergence in the two text measures, as cases with high decision time have quite similar word content, regardless of whether the patient is ultimately admitted or discharged. This is exactly what marginality would imply—Patients with long decision time have radiologist assessments that are on the fence between the two decisions.

Figure 8: Marginality Predicted by Report Text

A. Predictive Words Across Stations

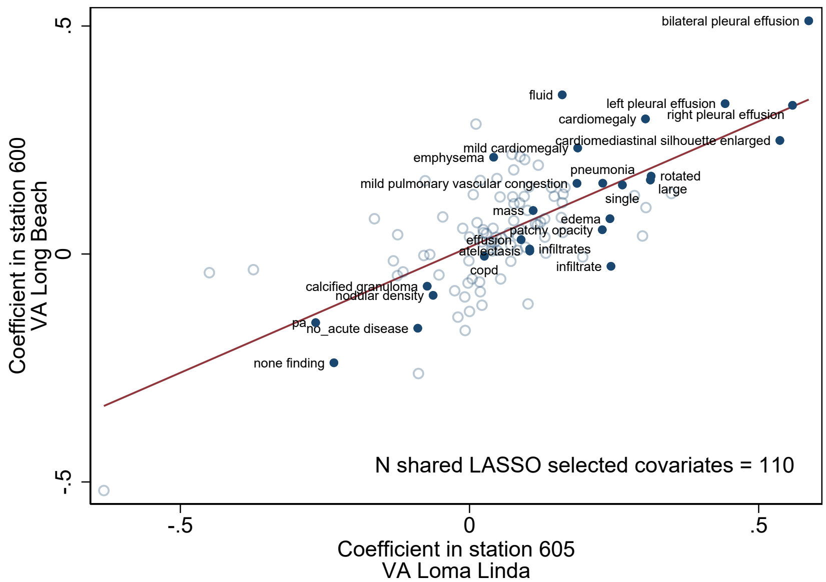

B. Predicted Admission by Report Text

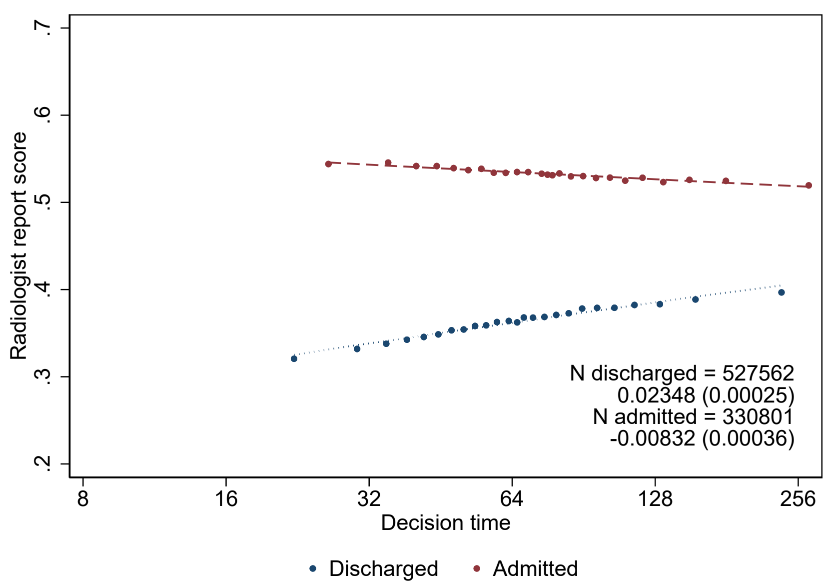

Notes: This figure shows results of analysis conducted on a subset of ED visits that included a chest X-ray. For each station, we run a LASSO logit regression of admission on indicators for coded words. Panel A displays the coefficients of words selected by LASSO in both station 600 (Long Beach) and 605 (Loma Linda). Each dot represents one of the 110 shared LASSO-selected words, and a subset of filled dots are labeled. Panel B shows a binned scatter plot of the "radiologist report score" on the y-axis against Decision time on the x-axis, split by discharged and admitted patients. The radiologist report score is defined as the probability of admission predicted by the station-specific LASSO logit models of admission on coded words from the radiology reports described above. We control for provider-year-month fixed effects.

6 Persistence in Predicting Marginality

A central theme of the paper is that decision time predicts *marginality* rather than simply *severity*. The evidence so far shows that, at the index visit, longer decision times are associated with a higher probability of subsequent bounceback among discharged patients and with shorter lengths of stay among admitted patients. In this section, we ask whether the predictive content of decision time persists beyond the index visit. Persistence is informative because it helps distinguish decision time from more transient markers of current severity. If decision time merely reflected short-run a patient's physiological distress, its predictive power might be expected to attenuate quickly. If instead it captures something more durable about a patient's underlying need for care or unresolved clinical ambiguity, then its predictive content may extend to future encounters.

To examine this, we compare decision time with a natural benchmark used in emergency departments: the Emergency Severity Index (ESI), a standardized five-level triage algorithm. We group ESI ratings into categories 1-2 (most severe), 3, and 4-5 (least severe, omitted category). Table 3 reports how quartiles of log decision time and ESI categories at the index visit predict three outcomes: admission from the index visit, bounceback after discharge from the index visit, and a second bounceback among patients who return and are discharged again. As expected, ESI is a strong predictor of the initial admission decision: patients with more severe ESI scores are substantially more likely to be admitted. Decision time also predicts admission, but the relative strength of ESI in this column is not surprising, since ESI is explicitly designed to summarize current urgency and resource needs at presentation.

The more interesting contrast emerges for subsequent outcomes. Decision time at the index visit strongly predicts bouncebacks among discharged patients, and its predictive content remains substantial even after a patient has already bounced back once. In particular, higher quartiles of index-visit decision time predict not only a first return to the ED, but also a second bounceback among patients who return and are discharged again. By contrast, ESI becomes much less informative once one moves beyond the initial disposition decision. Although ESI predicts admission from the index visit and, to a lesser extent, first bouncebacks, it has little predictive power for second bouncebacks. This difference is consistent with the interpretation that ESI is primarily a measure of current (and relatively transient) severity, whereas decision time reflects marginality and unresolved clinical needs that persist across encounters.

These findings suggest that decision time captures a durable dimension of patient risk that is not well summarized by traditional triage measures. Patients with high decision times appear more likely to be those whose care needs remain difficult to resolve, even after an initial discharge and even after a subsequent return visit. This does not imply that all such cases were misclassified at the index visit, nor that bouncebacks are themselves mistakes. Rather, it suggests that decision time identifies patients whose cases are persistently closer to the admit/discharge margin and therefore more likely to generate future utilization. In this sense, decision time appears to provide a signal of the need for follow-up that complements, rather than duplicates, standard clinical severity measures such as ESI.

7 Conclusion

This paper studies what physician decision time reveals in the emergency department. Our central finding is that decision time predicts *marginality* rather than *severity*. Standard patient characteristics such as age, comorbidities, triage acuity, and abnormal vital signs primarily predict severity: Among discharged patients, they are associated with worse downstream outcomes, and among admitted patients, they are associated with longer stays and higher mortality. Decision time instead predicts a different pattern. Among discharged patients, longer decision times are associated with more bouncebacks and higher mortality; among admitted patients, they are associated with shorter stays and lower mortality. As decision time rises, admitted and discharged patients increasingly resemble one another on *ex post* outcomes, consistent with the idea that longer decision times identify cases closer to the admit/discharge threshold. This interpretation is reinforced by the additional facts that longer decision times predict higher admission probabilities, excess decision time and consult activity peak at intermediate predicted admission probabilities, and the predictive content of decision time persists across future visits.

More broadly, the paper shows how a behaviorally generated signal can reveal latent information not fully contained in observed patient characteristics. In this sense, our results speak to a broader economic problem: when decision-makers allocate scarce expert attention under uncertainty, the time taken to make a choice may itself contain useful information about proximity to the decision margin. In the emergency department, this suggests that decision time may help identify encounters in which second opinions, consultations, additional diagnostics, observation-unit placement, or post-discharge follow-up are most valuable. The exact policy implications depend on the mechanism. If longer decision times primarily identify hard cases, they point toward interventions that improve information production or decision support. If they instead partly reflect operational strain or workflow bottlenecks, they point toward capacity expansion or process redesign. Either way, decision time may help identify where intervention is most valuable.

Several questions remain for future work. One is how decision time can be combined with rich electronic health record covariates to improve operational targeting in practice. Another is whether the predictive content of decision time varies systematically across physicians, emergency departments, or complaint groups. More generally, the results suggest a broader agenda of using behavioral traces generated during clinical care—not just recorded clinical covariates—to better understand where uncertainty is greatest and where additional resources have the highest marginal value.

Table 1: Summary Table

	Full sample	Admitted DT $\leq$ median	Admitted DT $>$ median	Discharged DT $\leq$ median	Discharged DT $>$ median	Chest X-ray sample
Age	59.99	65.40	66.21	57.77	59.58	65.88
Black share	0.27	0.24	0.24	0.28	0.27	0.26
# of Elixhauser comorbidities	2.94	4.26	4.27	2.45	2.84	3.76
Outpatient visits in prior year	12.46	16.05	16.19	11.06	12.24	.
Inpatient stays in prior year	0.49	1.07	1.04	0.31	0.42	0.70
Admission rate (%)	19.13	.	.	.	.	38.70
30-day mortality (%)	1.11	3.97	3.98	0.30	0.59	2.52
7-day bounceback (%)	7.88	.	.	7.11	8.81	5.08
Length of stay (days)	5.12	5.27	5.04	.	.	4.90
Observations	19,091,617	1,236,504	2,416,191	8,433,205	7,005,717	867,359

Notes: This table presents sample means of patient characteristics and outcomes for various strata of our sample. Column 1 presents means for the full sample. Columns 2 through 5 split the sample by admission/discharge status and within-station median DT (decision time). Column 6 presents means for the subset of ED visits that include a chest X-ray.

Table 2: Decision Time and Downstream Outcomes

A. Bouncebacks and Length of Stay

	Bounceback (%)			Log(LOS)		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(decision time)	0.9957	1.0032	0.8096	-0.0769	-0.0693	-0.0724
	(0.0075)	(0.0090)	(0.0090)	(0.0009)	(0.0011)	(0.0010)
Age / 100			0.6312			0.3173
			(0.0614)			(0.0075)
# of Elixhauser comorbidities			0.5976			0.0418
			(0.0040)			(0.0003)
# of abnormal vitals			1.4720			0.1305
			(0.0157)			(0.0012)
N	15,438,922	11,438,577	11,438,577	3,649,133	2,915,458	2,915,458
R^2	0.032	0.041	0.045	0.141	0.158	0.172
Sample	Full	Restricted	Restricted	Full	Restricted	Restricted

B. Mortality

	Mortality (%) | discharge			Mortality (%) | admission		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(decision time)	0.164	0.176	0.124	-0.099	-0.127	-0.184
	(0.002)	(0.002)	(0.002)	(0.011)	(0.013)	(0.013)
Age / 100			1.591			15.832
			(0.016)			(0.101)
# of Elixhauser comorbidities			0.140			0.347
			(0.001)			(0.005)
# of abnormal vitals			0.238			1.728
			(0.005)			(0.021)
N	15,438,922	11,438,577	11,438,577	3,652,695	2,918,123	2,918,123
R^2	0.030	0.039	0.044	0.096	0.114	0.134
Sample	Full	Restricted	Restricted	Full	Restricted	Restricted

Notes: The tables in panels A and B report OLS estimates of bounceback, log-transformed length of stay, and mortality outcomes on log-transformed decision time and patient characteristics. All specifications contain provider-year-month fixed effects, and standard errors are clustered at the same level. Bounceback and mortality estimates are multiplied by 100 to represent percentage-point changes. Models 1 and 4 use our main analytical sample. For models 2, 3, 5, and 6, we further reduce the sample and omit VHA patients who do not regularly use the VHA system. In particular, we omit patients who have not used primary or specialist care from the VHA in the past year. Elixhauser comorbidities are under-reported for such patients.

Table 3: Decision Time, ESI, and Persistent Predictions

	(1) Admission	(2) Bounceback	(3) Second bounceback
Log(DT) Q2	0.00215	0.00385	0.00582
	(0.00023)	(0.00022)	(0.00184)
Log(DT) Q3	0.01302	0.00880	0.00852
	(0.00028)	(0.00024)	(0.00195)
Log(DT) Q4	0.04911	0.01679	0.01399
	(0.00036)	(0.00028)	(0.00211)
ESI 3	0.1267	0.0123	0.0022
	(0.0003)	(0.0002)	(0.0016)
ESI 1,2	0.3465	0.0008	-0.0099
	(0.0005)	(0.0004)	(0.0033)
N	14,348,482	11,430,948	745,863
Mean dep. var.	0.203	0.083	0.175

Notes: This table reports OLS estimates of admission, bounceback, and second bounceback on within-station quartiles of decision time and Emergency Severity Index (ESI) group indicators. ESI categories 4 and 5 (least severe) as well as 1 and 2 (most severe) are grouped together. A second bounceback is defined as a return to the ED within 7 days of a bounceback visit for patients discharged on that visit. We control for Elixhauser comorbidity indicators, demographics, vital signs, VHA visits in the past year, and patient-year-month fixed effects.

References

Agency for Healthcare Research and Quality, "Improving the Emergency Department Discharge Process: Environmental Scan Report," Technical Report AHRQ Publication No. 14-0041-EF, Agency for Healthcare Research and Quality, Rockville, MD 2014. Accessed March 26, 2026.

Akerlof, George A., "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970, *84* (3), 488–500.

Alós-Ferrer, Carlos, Ernst Fehr, and Nick Netzer, "Time will tell: Recovering preferences when choices are noisy," *Journal of Political Economy*, 2021, *129* (6), 1828–1877.

American College of Emergency Physicians, "Emergency Department Crowding: High-Impact Solutions," Technical Report, American College of Emergency Physicians 2016. Accessed March 26, 2026.

_ , "ACEP Provides Feedback on CMS Boarding Measure," Technical Report, American College of Emergency Physicians 2024. Accessed March 26, 2026.

Card, David, Stefano DellaVigna, Chenxi Jiang, and Dmitry Taubinsky, "Understanding Expert Choices Using Decision Time," NBER Working Paper 32515, National Bureau of Economic Research 2024.

Centers for Medicare & Medicaid Services, *State Operations Manual, Appendix V: Interpretive Guidelines — Responsibilities of Medicare Participating Hospitals in Emergency Cases (EMTALA)* Centers for Medicare & Medicaid Services 2023. Accessed March 26, 2026.

Chabris, Christopher F, David Laibson, Carrie L Morris, Jonathon P Schuldt, and Dmitry Taubinsky, "Individual laboratory-measured discount rates predict field behavior," *Journal of risk and uncertainty*, 2008, *37* (2), 237–269.

Chan, David C., "The Efficiency of Slacking Off: Evidence from the Emergency Department," *Econometrica*, 2018, *86* (3), 997–1030.

Chan, David C, Matthew Gentzkow, and Chuan Yu, "Selection with Variation in Diagnostic Skill: Evidence from Radiologists," *The Quarterly Journal of Economics*, May 2022, *137* (2), 729–783.

Chiappori, Pierre-André and Bernard Salanié, "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 2000, *108* (1), 56–78.

Clithero, John A, "Improving out-of-sample predictions using response times and a model of the decision process," *Journal of Economic Behavior & Organization*, 2018, *148*, 344–375.

Cohen, Alma and Peter Siegelman, "Testing for Adverse Selection in Insurance Markets," *Journal of Risk and Insurance*, 2010, *77* (1), 39–84.

Cotet, Miruna, Wenjia Joyce Zhao, and Ian Krajbich, "Deliberation during online bargaining reveals strategic information," *Proceedings of the National Academy of Sciences*, 2025, *122* (7), e2410956122.

Einav, Liran, Amy Finkelstein, and Jonathan Levin, "Beyond Testing: Empirical Models of Insurance Markets," *Annual Review of Economics*, 2010, *2* (1), 311–336.

_ and _ , "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 2011, *25* (1), 115–138.

Finkelstein, Amy and James Poterba, "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market," *Journal of Political Economy*, 2004, *112* (1), 183–208.

_ and Kathleen McGarry, "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review*, 2006, *96* (4), 938–958.

Frydman, Cary and Ian Krajbich, "Using response times to infer others' private information: An application to information cascades," *Management Science*, 2022, *68* (4), 2970–2986.

Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki, "Speed, Accuracy, and the Optimal Timing of Choices," *American Economic Review*, 2018, *108* (12), 3651–3684.

Hoot, Nathan R. and Dominik Aronsky, "Systematic Review of Emergency Department Crowding: Causes, Effects, and Solutions," *Annals of Emergency Medicine*, 2008, *52* (2), 126–136.e1.

Kc, Diwas S. and Christian Terwiesch, "Impact of Workload on Service Time and Patient Safety: An Econometric Analysis of Hospital Operations," *Management Science*, 2009, *55* (9), 1486–1498.

Morley, Claire, Mark Unwin, Genevieve M. Peterson, Julie Stankovich, and Leigh Kinsman, "Emergency Department Crowding: A Systematic Review of Causes, Consequences and Solutions," *PLOS ONE*, 2018, *13* (8), e0203316.

Ratcliff, Roger and Jeffrey N Rouder, "Modeling response times for two-choice decisions," *Psychological science*, 1998, *9* (5), 347–356.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 1976, *90* (4), 629–649.

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

Spence, Michael, "Job Market Signaling," *Quarterly Journal of Economics*, 1973, *87* (3), 355–374.

Wald, Abraham, *Sequential Analysis*, New York: Wiley, 1947.

Appendix: "Waiting to Decide: Revealed Uncertainty in the Emergency Department"

David Chan

Stefano DellaVigna

David Card

March 27, 2026

Figure A.1: Decision Time Distribution

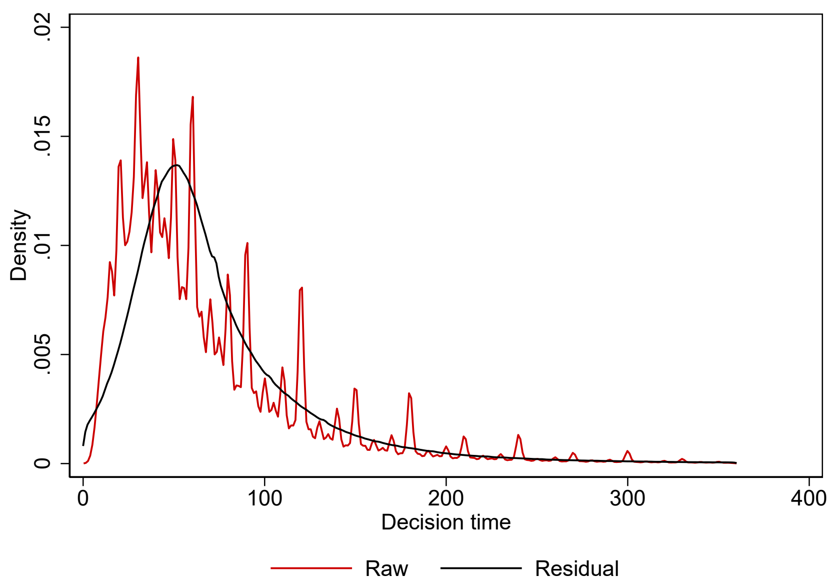

Notes: This figure shows the distribution of decision time in our main sample. The red line plots the kernel density of raw decision time, while the grey line plots the kernel density of decision time residualized by provider-year-month categories.

Figure A.2: Decision Time by Predicted Admission

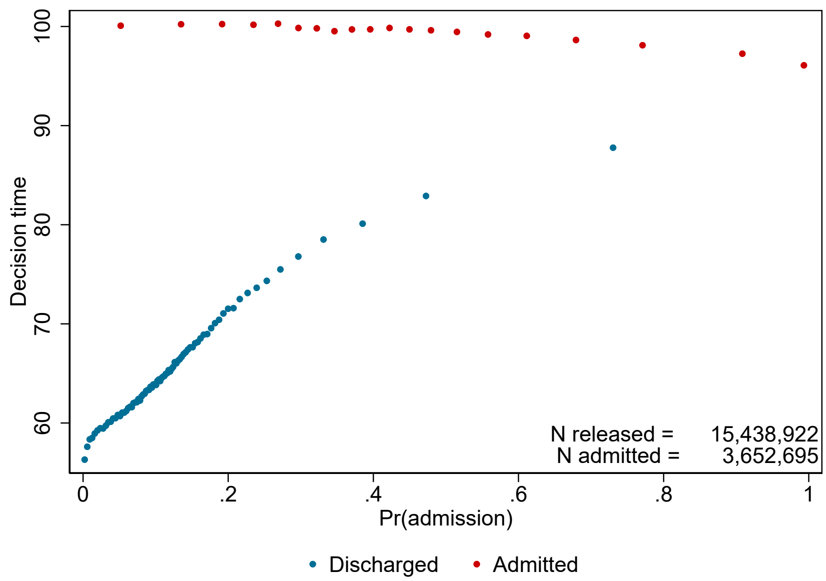

Notes: This figure shows a binned scatter plot of decision time on the y-axis against the probability of admission on the x-axis, separately for discharged and admitted patients. The probability of admission is predicted using station-specific Poisson regressions of admission on patient characteristics (Elixhauser comorbidity indicators, demographics, vital signs, and VHA visits in the past year). We control for provider-year-month fixed effects.

Table A.1: Sample Selection Steps

Sample step	Description	ED Visits	Patients	Providers	EDs/Stations
1. Build sample of ED visits from January 1, 2011 to Dec 31, 2023.	We restrict the sample to ED visits between January 1, 2011 and Dec 31, 2023.	30,540,616	5,877,728	16,893	121
2. Clean sample and collapse overlapping visits. Calculate decision time.	We drop visits missing departure time and where departure or disposition time is before the arrival time. Visits where the arrival time is equal to the departure time are dropped as well. We assume extremely long ED visits of a week or longer were recorded in error and drop such observations. Visits where a patient is recorded as having overlapping visits are interpreted as one long visit. Visits without a provider identifier are dropped as decision time per patient cannot be calculated in these cases. Decision time is calculated using the full sample of ED visits in order to encapsulate all patients under a provider's care during a given hour.	30,411,903	5,869,840	16,886	121
3. Prior visit restriction.	We restrict the analytic sample to ED visits where the patient had not been to the ED in the 30 days prior.	19,836,279	5,143,632	14,236	108
4. Merge and clean covariates.	The sample is restricted to visits for which demographic, mortality, and decision time data as calculated in Step 2 is non-missing. Visits by patients younger than 18 and older than 100 are dropped, as are observations where the patient's vitals are outside the realm of possibility.	19,091,617	5,056,384	14,234	108

Notes: This table reports changes in the sample size when applying each of the sample restrictions described in columns 1 and 2. Columns 3 through 6 report the number of visits, patients, providers, and EDs remaining after each step.