Optimal use of limited cognitive resources produces bias and noise in medical decisions

Bilal A. Bari[1,4*], Shuhan He[2,4], Thomas H. McCoy[3,4†], Samuel J. Gershman[5†]

[1]McLean Hospital, Belmont, MA
[2]Department of Emergency Medicine, Massachusetts General Hospital, Boston, MA
[3]Department of Psychiatry, Massachusetts General Hospital, Boston, MA
[4]Harvard Medical School, Boston, MA
[5]Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA
[*]Corresponding author. Email: bbari@mgb.org
[†]These authors contributed equally to this work.

Keywords: medical decision making; medical errors; resource rationality; reinforcement learning

Abstract

Medical providers make more errors under cognitive load, often viewed as arising from suboptimal decision-making. This interpretation relies on frameworks that ignore cognitive costs. Here, we offer a fundamentally different perspective that reveals an underlying structure to medical errors: bias (systematic deviations) and noise (variability) are inevitable consequences of optimal decision-making under cognitive resource constraints. We analyze orders placed by emergency department providers and find that cognitive load increases bias and noise, consistent with the optimal allocation of resources. Because providers near-optimally adjust to cognitive demands, this argues that guidelines that increase cognitive resources are necessary to reduce errors. Consistent with this perspective, bias and noise are minimized when multiple providers contribute to patient care. These findings have implications for optimizing medical care.

Introduction

Medical providers make decisions under significant cognitive load. They contend with high patient volumes, vast amounts of information, and limited time. It stands to reason that these cognitive factors should influence decision-making, prompting providers to adopt strategies that balance the utility of their decisions with cognitive costs. It is well established that cognitive fatigue leads providers to adopt strategies that lead to medical errors [1, 2, 3, 4, 5, 6]. This has led to the suggestion that the strategies employed by medical providers are strictly suboptimal [7], a statement based on normative frameworks that ignore cognitive costs.

Cognitive costs cannot, however, be ignored [8]. Accounting for these costs has been essential in explaining phenomena across fields ranging from neuroscience to moral judgment. It is therefore unclear whether providers employ decision-making strategies that are optimized for efficient memory use under cognitive constraints. In other words, how *should* providers adapt to demands on cognition and do their decision-making patterns follow these principles? Establishing this understanding is critical, as it provides insight into the source of medical errors and can inform interventions to support medical decision-making.

To address this question, we formalize the notion of optimal decision-making under cognitive resource constraints. We turn to rate-distortion theory, which analyzes how constraints on information transmission across a noisy channel affect distortions of input signals. Rate-distortion theory has been used to understand how a variety of human behavioral phenomena arise from information capacity limits [9, 10, 11, 12, 13]. Here, we use rate-distortion theory to model decision-making as a capacity-limited channel [14]. To test the theory's relevance in natural environments, we focus on real-world orders placed by medical providers in the emergency department, a setting with large shift-to-shift variation in patient volume that offers a natural setting to observe changes in cognitive load.

Consistent with the optimal use of limited resources, we find that medical decisions tend towards increased bias (systematic deviations) and noise (variability) as cognitive load increases. Consistent with theory, bias is not constant but adapts according to patient characteristics. A unique prediction is that adaptation produces perseveration—a tendency to repeat orders—which represents an efficient strategy to conserve resources. Across four independent measures of cognitive load, we confirm a relationship between cognitive load and perseveration. Finally, we show that bias and noise are minimized when multiple providers contribute to patient orders. This suggests that collaborative decision-making may mitigate the effects of cognitive load.

Results

Decision-making as a capacity-limited channel

We model each provider as a distribution over orders conditional on the patient (Figure 1A), which we term the *policy*, $\pi(Order|Patient)$. Throughout this manuscript, 'policy' is a technical term in reinforcement learning, referring to a provider's *internal* strategy for selecting actions, rather than an *externally* imposed set of guidelines, as the term is typically used. The policy is constrained by an upper bound on the mutual information between patients and orders, the *policy complexity* $I^\pi(Patient; Order)$, which defines the amount of memory needed to encode the policy; the more orders depend on patient-specific characteristics, the greater the policy complexity, and the more memory required to store that policy. The optimal policy maximizes expected utility U^π subject to the capacity constraint $I^\pi(Patient; Order) \leq C$, which can be reformulated as:

$$\pi^* = \underset{\pi}{\operatorname{argmax}}\, \beta U^\pi - I^\pi(Patient; Order), \quad (1)$$

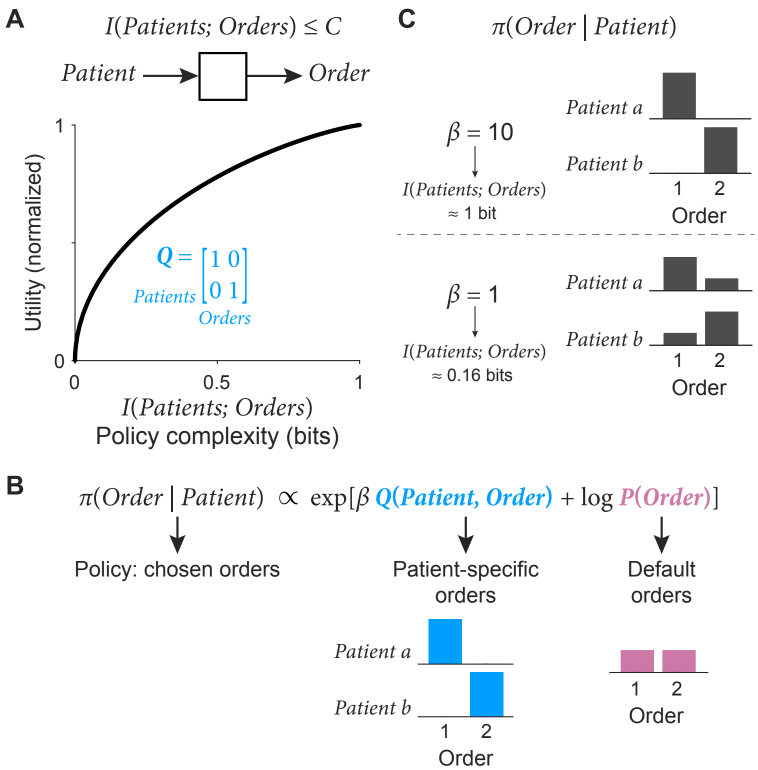

Figure 1: Decision-making modeled as a capacity-limited channel. (**A**) Patient representations are transmitted through a capacity-limited information channel to produce policy π, a conditional distribution of orders for a given patient. Given assumed utilities of patient-specific orders (blue inset), rate-distortion theory provides the optimal policy which can be used to trace out a utility/complexity frontier (black line). Each point on the frontier corresponds to a different choice of trade-off parameter β. (**B**) The optimal policy is a softmax function, seen ubiquitously in the reinforcement learning literature. It consists of two factors: patient-specific orders, $Q(Patient, Order)$, and default orders, $P(Order)$. (**C**) The parameter β controls the utility/information trade-off. When β is large (top panel), the policy results in more patient-specific orders but requires more memory (here, 1 bit). When β is small (bottom panel), the policy results in more default orders but require less memory (here, ≈ 0.16 bits). This results in noisier policies that are biased towards default orders.

where $\beta \geq 0$ is a trade-off parameter implicitly reflecting the capacity constraint C. The advantage of this formulation is that the optimal capacity-limited policy can be expressed explicitly:

$$\pi(Order|Patient) \propto \exp[\beta\ Q(Patient, Order) + \log P(Order)]. \tag{2}$$

The optimal policy is a function of two terms (Figure 1B): the utility of patient-specific orders, $Q(Patient, Order)$, and the default (marginal) probability of orders,

$$P(Order) = \sum_{Patient} P(Patient)\pi(Order|Patient). \tag{3}$$

The expected utility is given by $U^{\pi} = \mathbb{E}[Q(Patient, Order)|\pi]$.

The trade-off between utility (which we seek to maximize) and policy complexity (the memory cost that we seek to minimize) is dictated by β. When β is large, the policy is strongly dictated by patient-specific orders, resulting in a policy with high utility and high policy complexity, necessitating more memory (Figure 1C). This is the solution when memory constraints are ignored. At the other extreme, when β is small, the policy is strongly dictated by the default order distribution, resulting in a policy with comparatively lower utility but with the benefit of demanding less memory—what we term *policy compression* [14]. Intermediate β values interpolate between these two regimes, mixing patient-specific orders and default orders to generate the policy. This generates the utility/complexity frontier in Figure 1A, in which utility monotonically increases with policy complexity.

We sought to understand how cognitive load affects the optimal policy. In Figure 2A, we present a toy example where we define cognitive load as the number of concurrent patients that a provider must manage. Consistent with intuition, when a resource-limited provider manages more patients, the utility of care diminishes slightly, because the same memory must now be spread across more patients. Figure 2B highlights how this arises: regardless of the provider's policy complexity, the optimal β parameter is always smaller under higher cognitive load. This has two critical consequences (Figure 2C). First, decreased β produces a smaller reliance on patient-specific orders, which makes the policy more random. We can see this in Figure 1C, where the policy under smaller β has higher entropy (i.e., increased randomness). Second, decreased β produces a greater reliance on default orders, which introduces bias. This can again be observed in Figure 1C, where the policy is biased towards the default order distribution.

Policy compression in the emergency department

To identify signatures of noise and bias in medical decisions, we analyzed orders placed by medical providers over a 5-year period in an emergency department serving a major metropolitan area. This dataset consisted of 5,934 providers placing approximately 9 million orders across 448,129 patient encounters. Owing to well-known problems with estimating information-theoretic quantities from sparse distributions [15, 16], we collapsed the 3,258 unique orders into two categories: laboratory-based orders and other orders. This yielded a near maximal entropy distribution ($P(\text{Labs}) = 0.554$, $P(\text{Other}) = 0.446$), ideal for calculating information-theoretic quantities [17]. For each patient, we calculated the policy based on the first 10 orders placed for that patient (Figure 2D). We did this to ensure information-theoretic quantities were being estimated on similar distributions for each patient. We calculated the policy exclusively for instances where only one provider placed all 10 orders because we are interested in the cognitive properties of individual providers.

Because the electronic medical record we queried does not record the number of concurrent patients under a provider's care, we operationalized cognitive load as the total number of patients in the emergency department. Intuitively, if there are more overall patients in the emergency department, then there will be an increase in the average number of patients per provider. We found that policy complexity decreases with cognitive load ($\beta_{\text{Total number of patients}} = -0.581$, $t_{168} = -9.26$, $p < 10^{-16}$; Figure 2E) consistent with the idea that cognitive resources are stripped as the emergency department becomes crowded. According to the theory, both increased cognitive load in a busy emergency department and reduced policy complexity lead to a decreased β parameter, which increases both noise and bias. To test whether decisions become noisier, we estimated the conditional entropy of the policy, which quantifies the randomness of the policy. Consistent with theory (Figure 2F), conditional entropy increases with cognitive load ($\beta_{\text{Total number of patients}} = 0.637$, $t_{168} = 10.7$, $p < 10^{19}$; Figure 2G). To test for increased bias, we calculated the Kullback–Leibler (KL) divergence between the policy, $\pi(Order|Patient)$, and the default order distribution, $P(Order)$. The KL divergence measures the statistical distance between two distributions: the more the policy resembles the default order distribution, the smaller the KL divergence. Consistent with theory (Figure 2H), the KL divergence decreases with cognitive load ($\beta_{\text{Total number of patients}} = -0.583$, $t_{168} = -9.30$, $p < 10^{-16}$; Figure 2I).

The default order distribution adapts to patient characteristics

One key feature of the theory is that the bias, which arises from the default order distribution, should capture statistical regularities in patient characteristics to allow for the efficient use of memory. This requires the default order distribution to *adapt*; it should change when patient

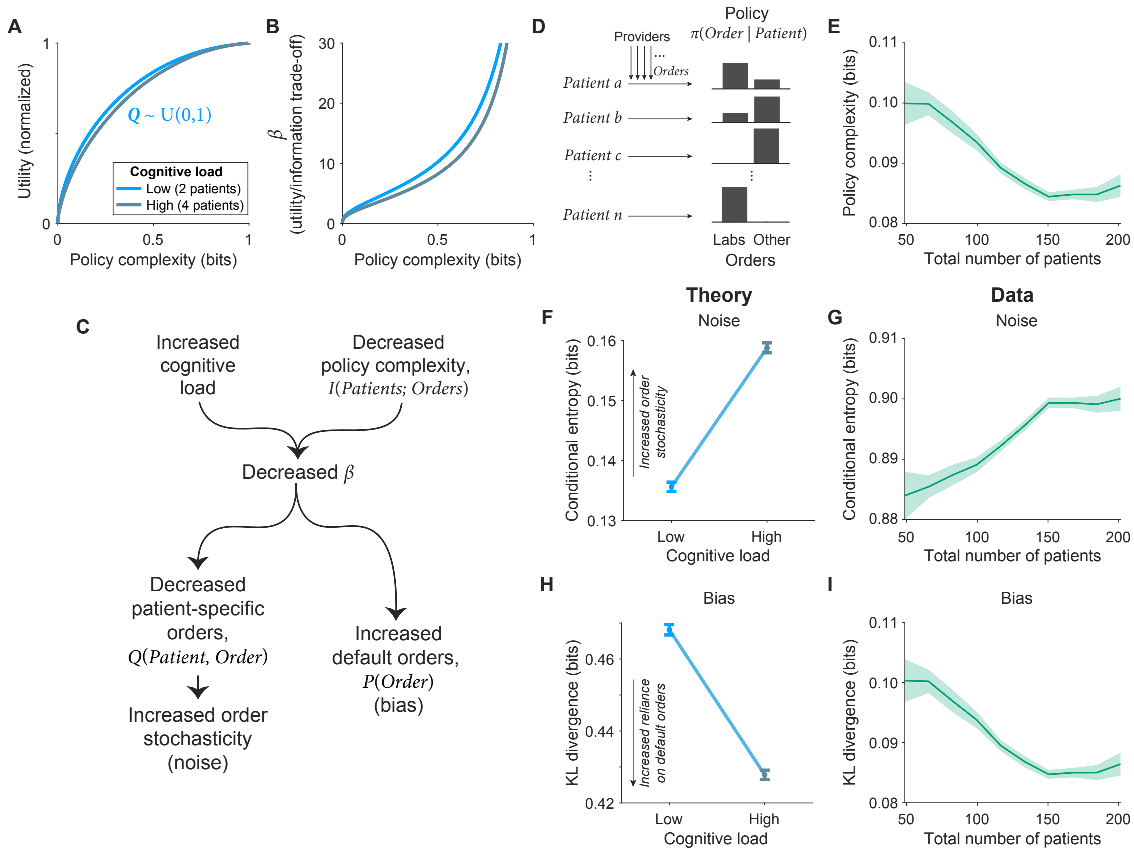

Figure 2: Cognitive load produces bias and noise in decisions. (**A**) The optimal utility/complexity frontier is lower under higher cognitive load because the same resources are spread over more patients. For simulation, the utility of patient-specific orders was drawn from a uniform distribution and averaged over 1,000 replicates. Cognitive load was operationalized as management of 2 (low load) or 4 (high load) simultaneous patients. (**B**) For the same policy complexity, the optimal β parameter is strictly smaller under higher cognitive load. (**C**) Higher cognitive load and decreased policy complexity both lead to smaller β resulting in increased noise and bias. (**D**) We categorized orders into lab-based orders and other orders and calculated the policy as the first 10 orders placed for patients. We use these empirical distributions to estimate information-theoretic quantities. (**E**) Policy complexity decreases with total number of patients in the emergency department, our proxy for cognitive load. (**F**) The theory predicts an increase in conditional entropy as a function of cognitive load, defined as in panel A. (**G**) Empirical conditional entropy increases with total number of patients. (**H**) The theory predicts a decrease in $KL(\pi(Order|Patient)||P(Order))$ with cognitive load due to increased bias towards the default order distribution. (**I**) Empirical $KL(\pi(Order|Patient)||P(Order))$ decreases with total number of patients. Error bars are SEM.

characteristics change. We have previously observed signatures of this adaptation process in well-controlled behavioral experiments [13, 18].

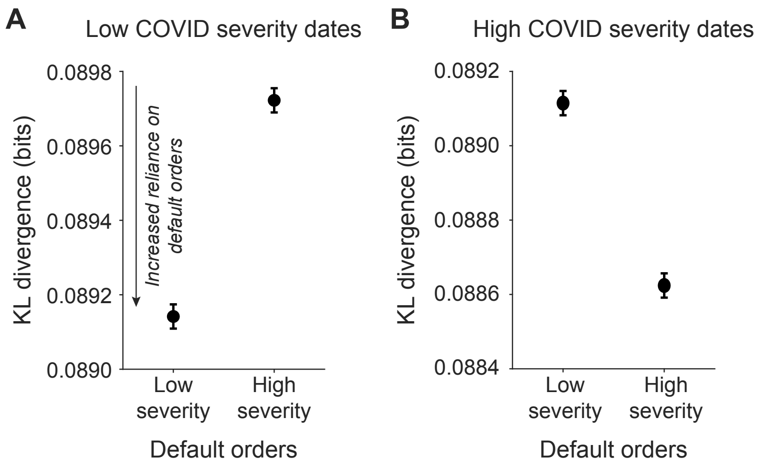

Figure 3: Default orders adapted to changing patient characteristics during the COVID pandemic. We split orders into low and high COVID severity dates, defined by the number of daily COVID cases, and calculated the default order distribution in each condition. For low and high COVID severity, we calculated the Kullback-Leibler (KL) divergence between the policy and each of these default order distributions. (**A**) For low COVID severity dates, the KL divergence between the policy and the default order distribution, $KL(\pi(Order|Patient)||P(Order)$, separately for low severity and high severity default orders. (**B**) The same analysis but for the policy during high COVID severity dates. All error bars are within-subject SEM.

To test whether providers adapt their default order distributions, we leveraged the COVID pandemic, since this resulted in large-scale changes in patient characteristics [19, 20]. Intuitively, if more patients present to the emergency department with respiratory concerns, then providers should be more likely to order respiratory-related orders across the board. We reasoned that default orders would adapt to COVID severity, capturing changes in ordering patterns among providers. We calculated terciles of daily COVID cases and split orders into low COVID severity (first tercile) and high COVID severity (third tercile) dates between January 1, 2020 and December 31, 2022, and calculated the default order distribution under each condition. We predicted that the policies implemented by providers on low COVID severity dates would resemble the default orders placed on low COVID severity dates. Further, the default orders learned on low COVID dates should provide a relatively poor description of orders placed on high COVID severity dates. For low COVID severity dates, we calculated the KL divergence twice: once between the policy and low COVID default orders, and again between the policy and high COVID default orders. We found a smaller KL divergence with low COVID default orders (KL divergence with low COVID default orders (mean $\pm$ within-subject SEM): $0.0891 \pm 3.25 \times 10^{-5}$; KL divergence with high COVID default orders: $0.0897 \pm 3.25 \times 10^{-5}$; paired t-test: $t_{42,229} = 8.92$, $p < 10^{-18}$; Figure 3A). We repeated this process for high COVID severity dates and found the inverse—the high COVID policy now resembled the high COVID default orders (KL divergence with low COVID default orders (mean $\pm$ within-subject SEM): $0.0891 \pm 3.27 \times 10^{-5}$; KL divergence with high COVID default orders: $0.0886 \pm 3.27 \times 10^{-5}$; paired t-test: $t_{41,119} = -7.51$, $p < 10^{-13}$; Figure 3B). These findings were exceedingly unlikely to be due to chance (difference in KL divergence for low COVID severity dates: 5.81×10^{-4}, shuffled data 95% CI $[-1.26 \times 10^{-6}, 3.17 \times 10^{-5}]$; difference in KL divergence for high COVID severity dates: -4.91×10^{-4}, shuffled data 95% CI $[-3.10 \times 10^{-5}, 1.00 \times 10^{-6}]$). These results strongly suggest that the default order distribution adapts to changing patient characteristics.

How do providers adapt the default order distribution? We have previously proposed that the brain learns the default order distribution by incremental updating [13, 18, 21, 22]:

$$\Delta P(Order) \propto \pi(Order|Patient) - P(Order). \quad (4)$$

This update ensures that if an order was recently placed, it is more likely to be chosen again, which results in perseveration. Note that perseveration is a consequence of optimal resource-constrained decision-making. Our use of the term is technical and does not imply its colloquial, often pejorative, connotation. We have observed signatures of such perseveration in well-controlled tasks [18, 22, 23].

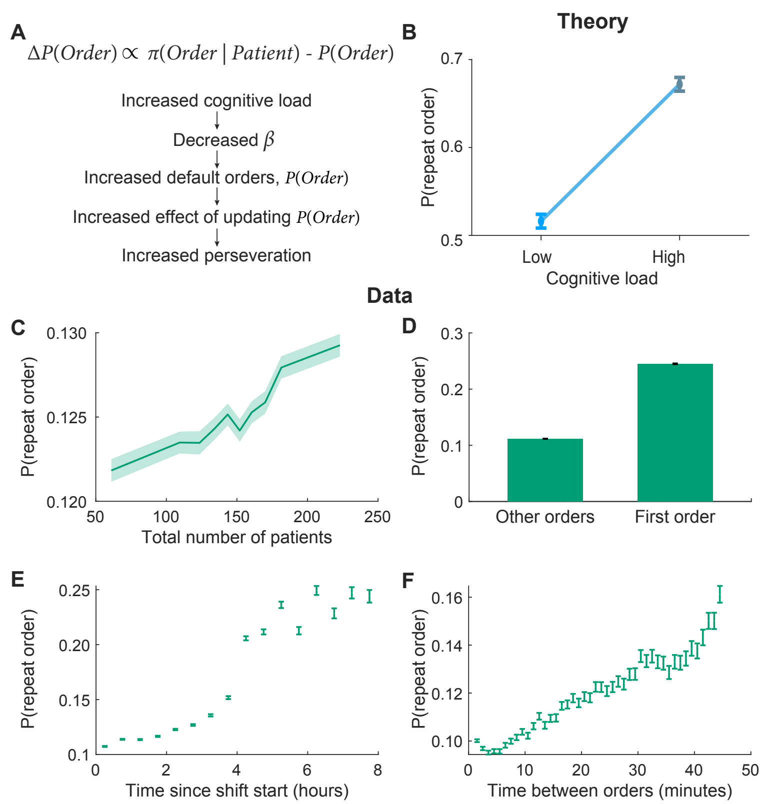

Figure 4: Adaptation of the default order distribution produces perseveration. (**A**) An iterative algorithm for updating the default order distribution produces perseveration (the tendency to repeat orders), which is magnified under increased cognitive load. (**B**) The theory predicts increased perseveration under increased cognitive load. (**C-F**) Perseveration increases with the total number of patients in the emergency department (C), for the first order placed for a patient (D), with time in shift (E), and with time between orders (F). Error bars are SEM.

We predicted increased perseveration under cognitive load, because default orders influence the policy more strongly in this regime due to reduced β (Figures 4A,B). Because we are no longer calculating information-theoretic quantities, we revert back to the raw order distribution used by providers (3,258 unique orders). We do this to ensure minimal preprocessing of the data and to assess whether providers repeat the raw orders placed (rather than order categories), which is a stronger test of the theory. Consistent with theory, providers perseverate more when there are more patients in the emergency department ($\beta_{\text{Total number of patients}} = 0.0228$, $t_{2,467,468} = 10.1$, $p < 10^{-23}$; Figure 4C), taxing memory. This effect held across a number of control analyses, including using different definitions of perseveration, only looking at perseveration to different patients, and controlling for the number of orders placed (figure 1A). This is consistent with findings from the working memory literature where subjects perseverate more when more items need to be remembered [24].

We next identified a more direct effect of a patient being incorporated into a provider's memory: when the first order for that patient is placed. Prior to the first order, providers are likely interviewing, assessing, and developing an initial plan for that patient—all events that tax memory. The first order should therefore mark when a provider has moved into a state of higher cognitive load (Figure 2A). This is indeed what we observe: perseveration increases when the order placed is the first order for a patient ($\beta_{\text{First order}} = 0.264$, $t_{2,411,398} = 143$, $p < 10^{-100}$; Figure 4D, 1B).

We next identified a signature of learning the default order distribution. Because the bias reflects continuous incremental adjustments, the tendency to perseverate should increase over time [24, 25, 26]. Consistent with this hypothesis, perseveration increases with the time that has elapsed since a provider started their shift ($\beta_{\text{Time since shift start}} = 0.213$, $t_{2,467,468} = 114$, $p < 10^{-100}$; Figure 4E, 1C).

Finally, we investigated the effect of memory retention interval, operationalized as the time between orders. We reasoned that the longer it has been since a provider placed an order, the longer they are likely to have to hold onto information related to the upcoming order. In studies of working memory, longer retention intervals tend to increase perseveration [24, 27]. Indeed, we found that perseveration increases as the time between orders increases ($\beta_{\text{Time between orders}} = 0.190$, $t_{2,467,468} = 114$, $p < 10^{-100}$; Figure 4F, 1D).

To ensure the independence of these effects, we fit a logistic mixed effects regression to predict perseveration as a function of the previous 4 factors as well as chance levels of perseveration. We found that all regression coefficients were positive and significant, confirming that each factor contributes to perseveration ($\beta_{\text{Total number of patients}} = 0.0207$, $t_{2,411,395} = 9.01$, $p < 10^{-18}$; $\beta_{\text{First order}} = 0.256$, $t_{2,411,395} = 138$, $p < 10^{-100}$; $\beta_{\text{Time since shift start}} = 0.138$, $t_{2,411,395} = 42.0$, $p < 10^{-100}$; $\beta_{\text{Time between orders}} = 0.0789$, $t_{2,411,395} = 26.3$, $p < 10^{-100}$).

Taken together, we find that medical providers systematically perseverate when placing medical orders. Far from a hindrance, this perseveration serves to economize limited cognitive resources.

Counteracting the effects of cognitive load

Although our results argue that providers optimally adapt to cognitive load, the resulting bias and noise in medical decisions is unwanted from a systems perspective. How then can bias and noise be reduced, to limit downstream medical errors? It is recognized that effective teamwork can offset the deleterious effects of cognitive fatigue [28, 29, 30, 31, 32], though this introduces miscommunication, which lead to medical errors [33, 34, 35, 36]. We therefore sought to quantify the effect of multiple providers collaborating on care, to see if and how teamwork improves patient care. Concretely, we partitioned the data into instances where one provider placed the first 10 orders for a patient, and instances where two or more providers placed these orders (our operationalization of teamwork). We found that when multiple providers contribute to patient care, decisions are improved across the board. Policy complexity increases across all levels of cognitive load (policy complexity under 1 provider (mean $\pm$ SEM): $0.0906 \pm 7.74 \times 10^{-4}$; policy complexity under 2 or more providers: $0.116 \pm 1.66 \times 10^{-3}$; two-sample t-test: $t_{338} = -13.9$, $p < 10^{-34}$; Figure 5A). Commensurate with the increase in policy complexity, noise and bias both decrease (conditional entropy under 1 provider (mean $\pm$ SEM): $0.893 \pm 7.45 \times 10^{-4}$; conditional entropy under 2 or more providers: $0.883 \pm 1.80 \times 10^{-3}$; two-sample t-test: $t_{338} = 5.46$, $p < 10^{-7}$; KL divergence under 1 provider: $0.0908 \pm 7.77 \times 10^{-4}$; KL divergence under 2 or more providers: $0.116 \pm 1.65 \times 10^{-3}$; two-sample t-test: $t_{338} = -13.8$, $p < 10^{-34}$; Figure 5B,C). Teamwork therefore effectively increases the cognitive resources available, in service of optimal decision-making.

Discussion

Taken together, our work makes the important point that bias and noise in decisions are inevitable properties of an optimal, resource-limited system. Importantly, the theory accounts for a subset of apparent errors committed by providers. Some errors are clearly attributable to other sources. For example, a provider's distribution of patient-specific orders may be misspecified, resulting in

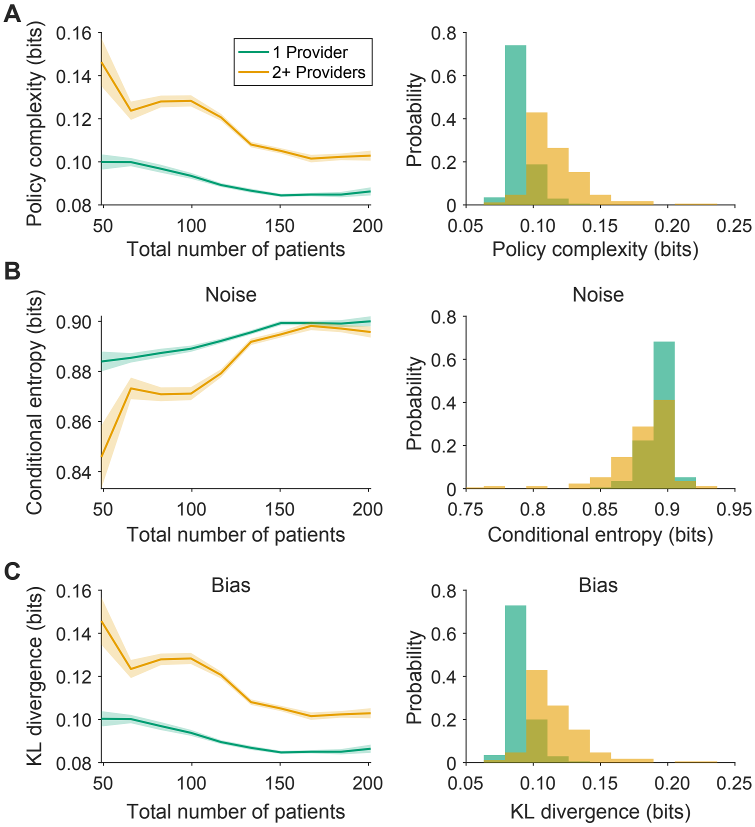

Figure 5: Effects of teamwork on bias and noise. (**A**) We calculated the policy separately when 1 provider made all orders and when 2 or more providers made all orders. When multiple providers contribute to orders, the policy complexity increases. (**B**) The conditional entropy, an estimate of the stochasticity of the policy, creases when multiple providers contribute to orders. (**C**) Bias, estimated as $KL(\pi(Order|Patient)||P(Order))$, decreases (larger KL divergence) when multiple providers contribute to orders. Error bars are SEM.

the inappropriate order being placed even under cognitively ideal conditions. Medical malpractice law recognizes these types of errors as a "failure to adhere to the standards of the profession" [37]. Quantifying the extent to which providers optimize their actions necessitates estimating the underlying reward structure that governs clinician behavior—a line of research currently under development [38, 39, 40]. Our work, however, makes an important distinction: there is a ceiling to the performance a provider can attain under a capacity limit.

We therefore suggest that guidelines focused on interventions to increase cognitive resources or minimize cognitive load are likely to have outsized impact, particularly in environments like emergency departments where cognitive demands are high. For example, interventions aimed at limiting crowding in the emergency department [41] may free up cognitive resources for providers, improving decision quality. Emergency medicine often requires rapid decision-making (e.g., head imaging without a thorough history for suspected stroke), contrasted with the more thorough history and physical examination typical in outpatient settings. These differences reflect necessary adaptations to situational demands rather than deficiencies in care, highlighting the value of considering cognitive resources explicitly. Another suggestion is that medical protocols can reify cognitive resources as if they were physical resources. Mass casualty protocols, for example, are typically framed in terms of physical resource management but implicitly also manage cognitive resources by providing simple and rapid triage guidelines to streamline decision-making [42, 43]. Explicitly accounting for cognition as a resource may allow systems to interpolate between regimes where cognitive resources are rich and where they are depleted, helping improve patient outcomes across varying clinical environments.

By integrating cognitive costs into formal models of medical decision-making, our work opens avenues for designing interventions that enhance decision-making and improve patient outcomes.

Methods

Rate-distortion theory: a capacity limit applied to decision-making

All information processing systems are subject to physical constraints that limit the ability to perfectly store and transmit information. These limits place an upper bound on achievable performance [8]. Here, we use rate-distortion theory to formalize decision-making as a constrained optimization problem that trades off the utility of decisions with the associated cognitive costs [14, 44].

We consider a decision-maker as a provider implementing a policy, $\pi(Order|Patient)$, a probability distribution that maps patients onto orders (in reinforcement learning, the standard terminology is an *agent* using a policy that maps *states* onto *actions*). From the perspective of rate-distortion theory, providers generate policies by transmitting patient information across a capacity-limited channel to generate orders. We define the cognitive cost of decision-making as the information rate across this channel: the *policy complexity*, or the mutual information between patients and orders:

$$I^{\pi}(Patient; Order) = \sum_{Patient} P(Patient) \sum_{Order} \pi(Order|Patient) \log \frac{\pi(Order|Patient)}{P(Order)} \tag{5}$$

where $P(Order) = \sum_{Patient} P(Patient)\pi(Order|Patient)$ is the marginal order distribution, which we refer to as default orders throughout.

We assume that this channel is subject to a capacity constraint, C, or an upper bound on the policy complexity. According to Shannon's noisy channel theorem, the minimum expected number of bits to errorlessly transmit a signal across a channel is equal to the mutual information. If the optimal policy requires more memory than the provider possesses, then the provider must discard some patient-specific information to reduce the policy complexity under the capacity limit. The optimal policy, π^* is defined as:

$$\pi^* = \underset{\pi}{\operatorname{argmax}}\, U^{\pi}, \text{ subject to } I^{\pi}(Patient; Order) \leq C \tag{6}$$

where U^{π} is the expected utility of the policy π:

$$U^{\pi} = \sum_{Patient} P(Patient) \sum_{Order} \pi(Order|Patient) Q(Patient, Order) \tag{7}$$

and $Q(Patient, Order)$ is the expected utility of orders for a given patient, which we refer to as patient-specific orders throughout.

This constrained optimization problem can be be written as an unconstrained optimization problem using the following Lagrangian equation:

$$\pi^* = \underset{\pi}{\operatorname{argmax}}\, \beta U^{\pi} - I^{\pi}(Patient; Order) - \sum_{Patient} \lambda(Patient) \left(\sum_{Order} \pi(Order|Patient) - 1 \right) \tag{8}$$

where $\beta \geq 0$, $\lambda(Patient) \geq 0$ are Lagrange multipliers. The solution to this equation is:

$$\pi(Order|Patient) \propto \exp[\beta\, Q(Patient, Order) + \log P(Order)] \tag{9}$$

The optimal policy takes the form of a softmax function, ubiquitous in the reinforcement learning literature for modeling both artificial and biological agents. β plays the role of a utility/information trade-off parameter: β units of utility can be "bought" for 1 unit of information. The exact relationship between β and policy complexity is given as:

$$\beta^{-1} = \frac{dU^{\pi}}{dI^{\pi}(Patient; Order)} \tag{10}$$

which provides a geometric interpretation of β, which is the (inverse) slope of the utility/complexity curve—when the slope is shallow, β is high and when the slope is steep, β is small. In Figure 1A, at high policy complexity, where $\frac{dU^{\pi}}{dI^{\pi}(Patient;Order)}$ is shallow, the optimal β is large and the policy is largely a function of patient-specific orders, $Q(Patient, Order)$. At low policy complexity, the slope is steep and the optimal β is small, meaning default orders, $P(Order)$, dominate the policy.

Theoretical simulations

To generate the optimal policy in Figure 1, we define $Q(Order, Patient) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, where rows index patients and columns index orders, and compute the optimal policy using the Blahut-Arimoto algorithm [45, 46]. Note that this is a toy model to help build intuitions, not a literal model of a realistic provider policy.

To generate the optimal policy in Figure 2A,B, we define $Q(Order, Patient)$ as a 2x2 (Patient X Order) distribution for the low cognitive load condition and as a 4x2 distribution for the high cognitive load condition. We simulate 2 orders to mimic the 2 order categories used for empirical data analysis (lab-based orders vs other orders). For each simulation, we sample from a uniform distribution over the interval $[0, 1]$ for each entry of the Q matrix. This captures the intuition that the optimal set of orders varies across patients, such that lab-based orders are optimal for some patients, other orders are optimal for others, both may be optimal for others, and so on. We enforce the constraint that each order should be optimal for at least 1 patient. This prevents the optimal policy complexity from being 0 bits, which is highly implausible in practice and therefore a poor description of reality. We use the Blahut-Arimoto algorithm to generate the optimal policies and repeat this process 1,000 times. We average the relevant quantities (utility/complexity curves, β/complexity curves) across simulations to generate the curves in Figure 2A,B.

We estimate the conditional entropy as:

$$H(Order|Patient) = \sum_{Patient} P(Patient) \sum_{Order} \pi(Order|Patient) \log \pi(Order|Patient) \tag{11}$$

and we calculate the KL divergence as:

$$KL(\pi(Order|Patient)||P(Order)) = \sum_{Order} \pi(Order|Patient) \log \frac{\pi(Order|Patient)}{P(Order)} \tag{12}$$

In Figure 2F,G, we report the average conditional entropy and KL divergence across all policies for each cognitive load condition (i.e., by marginalizing over policy complexity).

Emergency department ordering data

We obtained approval from the Massachusetts General Brigham Institutional Review Board prior to conducting this research. We analyzed all non-medication orders placed for patients in the Massachusetts General Hospital Emergency Department from July 1, 2019 to June 30, 2024. We restricted the dataset to non-medication orders to limit the combinatorial complexity inherent with medication-based orders (e.g., different doses, formulations, route of administration, frequency, standing vs as-needed). We analyzed the following variable associated with each order: patient encounter ID (unique ID for that specific encounter), provider ID, time of order placement, order type, total number of patients in emergency department. The full dataset consisted of 448,129 unique patient encounters, 5,934 unique ordering providers, and 8,942,841 orders (3,258 unique orders). We preprocessed the data by grouping all simultaneously-released orders into 'order batches' which resulted in 2,916,690 order batches. The mean ($\pm$ SEM) order batch size was 2.67 ($\pm$ 2.18×10^{-3}) orders, with a median and mode of size 1 (68.6% of all order batches consisted of just 1 order).

Estimating information-theoretic quantities on medical orders

To estimate policy complexity, conditional entropy, and KL divergence in Figure 3E,G,I and Figure 5, we defined the policy, $\pi(Order|Patient)$, as the first 10 orders placed for a patient, where patient was defined by patient encounter ID. We collapsed the 3,258 unique orders into the two categories of lab-based orders and other orders to avoid problems with estimating entropy from sparse distributions [15, 16]. We estimated the total number of patients as the average number of patients in the emergency department over the first 10 orders and rounded down to the nearest integer. We ignored patient encounters with fewer than 10 total orders. We ignored all orders past the first 10. Because orders were occasionally batched (i.e., multiple orders released simultaneously), if a given batch resulted in more than 10 orders being placed for a patient, we randomly discarded orders until only 10 orders remained for a patient.

We computed information-theoretic quantities (Equations 5, 11, and 12) as a function of the total number of patients. Specifically, we estimated each quantity for individual values of the total number of patients and estimated the default order distribution as $P(Order) = \sum_{Patient} P(Patient)$ $\pi(Order|Patient)$ (e.g., for 50 patients, then 51 patients, and so on). Since we used the patient encounter ID to define patients, $P(Patient)$ was equiprobable for each analysis. For Figure 2, we did this for instances where only 1 provider ordered the first 10 orders, which resulted in 196,153 unique patients encounters. For Figure 5, we included instances where 2 or more providers contributed to the first 10 orders, which resulted in 89,195 unique patient encounters.

COVID severity analyses

To assess how the COVID pandemic affected the default order distribution, we analyzed COVID cases in Massachusetts from January 1, 2020 to December 31, 2022, drawn from the Oxford COVID-19 Government Response Tracker which compiled cases from open dataset (e.g., Johns Hopkins University Coronavirus Resource Center) [47]. Owing to periodic data collection (e.g., once weekly reporting of COVID cases at some timepoints), we smoothed the data with a Savitzky-Golay filter of span 50 and degree 2. We calculated terciles of COVID cases and defined the first tercile as low severity dates and the third tercile as high severity dates. We calculated the default order distributions separately for each set of dates. For the shuffled control, we randomly shuffled COVID low and COVID high dates and calculated the difference in KL divergence as the KL divergence with high severity default orders minus KL divergence with low severity default orders.

Perseveration analyses

We use the following iterative update equation to update the marginal order distribution [13, 18, 21, 22]:

$$\Delta P(Order) = \alpha(\pi(Order|Patient) - P(Order)), \tag{13}$$

where $\alpha \in [0, 1]$ is a learning rate parameter. To generate the simulated results in Figure 4B, we defined Q(Order,Patient) as a 2x2 (Patient x Order) distribution, sampled from a uniform distribution over the interval [0, 1] using $\beta = 100$ for the low cognitive load condition and $\beta = 1$ for the high cognitive load condition. $P(Order)$ was initialized as $[0.5, 0.5]$ and we randomly sampled from patients (either patient 1 or patient 2) for 1,000 trials. We calculated the policy according to Equation 2 and used it to update $P(Order)$ according to Equation 13 using $\alpha = 0.03$. We calculated the probability of repeating an order across all trials to estimate perseveration. We repeated this simulation 100 times.

In Figure 4, we define perseveration as a binary variable equal to 1 if any order was repeated and 0 otherwise.

To account for chance levels of perseveration, which increases when the order batch size is larger, or when the order placed was higher probability, we used the following process. First, we defined I_O as the set of all orders placed at a particular timepoint. We calculated the probability that no order would be repeated as $P(\text{no repeat}) = (1 - \sum_{I_O} P(I_O))^N$ where $P(I_O)$ is the probability of making a particular order, drawn from $P(Order)$ and N is the number of orders in the current batch. The chance probability of any one order is repeated is therefore $P(\text{repeat order}) = 1 - P(\text{no repeat})$. Intuitively, if the number of prior orders goes up or the probability of those orders is high (either of which increases $\sum_{I_O} P(I_O)$), or the number of orders in the current batch is large (increasing N), then the probability of repeating an order by chance increases. We include this as a regressor in all perseveration regression analyses (see Statistical analyses below).

We calculated shift duration by calculating the cumulative time elapsed from the first order to the current order. When the cumulative time elapsed exceeded 16 hours, we defined the end of the shift and considered the current order as the start of a new shift. This is a longer duration than typical shifts, which we included to allow ample time for post-shift orders to be placed, in the case of lengthy passoffs (e.g., multiple codes running at the end of shift requiring a provider's attention, limiting passoff at the assigned time). For all analyses, we only include perseveration within a shift (i.e., we ignore the first order of a shift in our analyses).

For control analyses (Figure 1), we used a different definition of perseveration (fractional), instances where perseveration was to different patients, and for instances where the order batch size for the prior and current orders was equal to 1. We used a fractional definition of perseveration to account for instances where where multiple orders were batched and released simultaneously (in Figure 1, perseveration was defined as the fraction of orders that were repeated). We used instances where perseveration was to different patients because orders may be repeated for the same patient for reasons unrelated to the theory (e.g., because an inadequate blood sample was drawn, requiring a duplicate order to be placed for another attempt). We used instances with a batch order size of 1 for prior and current orders to control for effects related to batch size.

Statistical analyses

For all regressions, we standardized all variables by z-scoring, with the exception of the binary perseveration variable used as the outcome variable with logistic regressions. For all perseveration analyses, we used random effects models which included the fixed effects listed in the text and a random intercept per provider. All error bars are SEM unless otherwise reported. We used the method of [48] to calculate within-subject SEM for Figure 3. We report a minimum p value of 10^{-100}.

Funding sources: This work was supported by National Institute of Mental Health grants R25MH135837 and T32MH125786 (B.A.B.) and a Schmidt Sciences Polymath Award (S.J.G.).

Author contributions: B.A.B. contributed to formal analysis, investigation, methodology, visualization, and writing the original draft. S.H. contributed to methodology and project administration. T.H.M. contributed to data curation, methodology, project administration. S.J.G. contributed to funding acquisition, methodology, project administration. All authors contributed to conceptualization and reviewing & editing the manuscript.

Competing interests: S.H.: Employment: Mass General Physician Organization. Mass General Institute of Health Professions. Consulting Fees: Maze Eng Inc, ConductScience Inc, VIO Elevate Beauty Partners. No other disclosures are reported. S.H. is a volunteer at Health Tech Without Borders. T.H.M.: Dr. McCoy is funded to grants to his institution from InterSystems, Philips Research, Koa Health, as well as the National Institute of Mental Health. He is a paid Associate Editor at npj Digital Medicine and paid speaker within the MGH Psychiatry Academy.

Data and materials availability: The Massachusetts General Brigham Institutional Review Board does not allow open dissemination or the deposition of these electronic medical record data in a repository because participants (patients and providers) did not provide formal consent. This is to minimize risk of individual patients and providers being identified. To obtain access to the data, investigators should collaborator with a Massachusetts General Brigham-affiliated faculty member to submit an application to the Institutional Review Board requesting access to these data.

Acknowledgments: We thank Sara Golas for extracting the data from the electronic medical record. We thank Gideon Loevinsohn for comments on an earlier draft of the manuscript.

References

[1] Donaldson, M. S., Corrigan, J. M. & Kohn, L. T. *To err is human: building a safer health system* (National Academies Press, 2000).

[2] Shojania, K. G., Duncan, B. W., McDonald, K. M., Wachter, R. M. & Markowitz, A. J. *Making health care safer: a critical analysis of patient safety practices.* 43 (Agency for Healthcare Research and Quality, U.S. Department of Health and Human Services, 2001).

[3] Landrigan, C. P. *et al.* Effect of reducing interns' work hours on serious medical errors in intensive care units. *New England Journal of Medicine* **351**, 1838–1848 (2004).

[4] Barger, L. K. *et al.* Impact of extended-duration shifts on medical errors, adverse events, and attentional failures. *PLoS medicine* **3**, e487 (2006).

[5] Lockley, S. W. *et al.* Effects of health care provider work hours and sleep deprivation on safety and performance. *The Joint Commission Journal on Quality and Patient Safety* **33**, 7–18 (2007).

[6] West, C. P., Tan, A. D., Habermann, T. M., Sloan, J. A. & Shanafelt, T. D. Association of resident fatigue and distress with perceived medical errors. *Jama* **302**, 1294–1300 (2009).

[7] Saposnik, G., Redelmeier, D., Ruff, C. C. & Tobler, P. N. Cognitive biases associated with medical decisions: a systematic review. *BMC medical informatics and decision making* **16**, 1–14 (2016).

[8] Lieder, F. & Griffiths, T. L. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and brain sciences* **43**, e1 (2020).

[9] Sims, C. R. Rate–distortion theory and human perception. *Cognition* **152**, 181–198 (2016).

[10] Sims, C. R. Efficient coding explains the universal law of generalization in human perception. *Science* **360**, 652–656 (2018).

[11] Gershman, S. J. Origin of perseveration in the trade-off between reward and complexity. *Cognition* **204**, 104394 (2020).

[12] Bari, B. A. & Gershman, S. J. Resource-rational psychopathology. *Behavioral Neuroscience* (2024).

[13] Lai, L. & Gershman, S. J. Human decision making balances reward maximization and policy compression. *PLOS Computational Biology* **20**, e1012057 (2024).

[14] Lai, L. & Gershman, S. J. Policy compression: An information bottleneck in action selection. In *Psychology of Learning and Motivation*, vol. 74, 195–232 (Elsevier, 2021).

[15] Nemenman, I., Shafee, F. & Bialek, W. Entropy and inference, revisited. *Advances in neural information processing systems* **14** (2001).

[16] Paninski, L. Estimation of entropy and mutual information. *Neural computation* **15**, 1191–1253 (2003).

[17] Jaynes, E. T. Information theory and statistical mechanics. *Physical review* **106**, 620 (1957).

[18] Liu, S., Lai, L., Gershman, S. J. & Bari, B. A. Time and memory costs jointly determine a speed-accuracy trade-off and set-size effects. *PsyArXiv* (2024).

[19] Clifford, C. T. *et al.* Association between covid-19 diagnosis and presenting chief complaint from new york city triage data. *The American Journal of Emergency Medicine* **46**, 520–524 (2021).

[20] Rodgers, L. *et al.* Changes in seasonal respiratory illnesses in the united states during the coronavirus disease 2019 (covid-19) pandemic. *Clinical Infectious Diseases* **73**, S110–S117 (2021).

[21] Gershman, S. J. & Lai, L. The reward-complexity trade-off in schizophrenia. *Computational Psychiatry* **5** (2021).

[22] Bari, B. A., Krystal, A. D., Pizzagalli, D. A. & Gershman, S. J. Computationally-informed insights into anhedonia and treatment by κ-opioid receptor antagonism. *medRxiv* 2024–04 (2024).

[23] Bari, B. A. & Gershman, S. J. Undermatching is a consequence of policy compression. *Journal of Neuroscience* **43**, 447–457 (2023).

[24] Jakob, A. M. & Gershman, S. J. Rate-distortion theory of neural coding and its implications for working memory. *Elife* **12**, e79450 (2023).

[25] Foster, J. J., Bsales, E. M., Jaffe, R. J. & Awh, E. Alpha-band activity reveals spontaneous representations of spatial position in visual working memory. *Current Biology* **27**, 3216–3223 (2017).

[26] Barbosa, J. & Compte, A. Build-up of serial dependence in color working memory. *Scientific reports* **10**, 10959 (2020).

[27] Bliss, D. P., Sun, J. J. & D'Esposito, M. Serial dependence is absent at the time of perception but increases in visual working memory. *Scientific reports* **7**, 14739 (2017).

[28] Manser, T. Teamwork and patient safety in dynamic domains of healthcare: a review of the literature. *Acta Anaesthesiologica Scandinavica* **53**, 143–151 (2009).

[29] Schmutz, J. & Manser, T. Do team processes really have an effect on clinical performance? a systematic literature review. *British Journal of Anaesthesia* **110**, 529–544 (2013).

[30] Welp, A. & Manser, T. Integrating teamwork, clinician occupational well-being and patient safety–development of a conceptual framework based on a systematic review. *BMC health services research* **16**, 1–44 (2016).

[31] Ma, C., Park, S. H. & Shang, J. Inter-and intra-disciplinary collaboration and patient safety outcomes in us acute care hospital units: A cross-sectional study. *International Journal of Nursing Studies* **85**, 1–6 (2018).

[32] Cho, H., Sagherian, K., Scott, L. D. & Steege, L. M. Occupational fatigue, workload and nursing teamwork in hospital nurses. *Journal of advanced nursing* **78**, 2313–2326 (2022).

[33] Aronson, J. K. Medication errors: what they are, how they happen, and how to avoid them. *QJM: An International Journal of Medicine* **102**, 513–521 (2009).

[34] Zavala, A. M., Day, G. E., Plummer, D. & Bamford-Wade, A. Decision-making under pressure: medical errors in uncertain and dynamic environments. *Australian Health Review* **42**, 395–402 (2017).

[35] Shitu, Z., Hassan, I., Aung, M. M. T., Kamaruzaman, T. H. T. & Musa, R. M. Avoiding medication errors through effective communication in a healthcare environment. *Malaysian Journal of Movement, Health & Exercise* **7**, 115–128 (2018).

[36] The Joint Commission. Sentinel event statistics (1995–2019) (2019). URL https://psnet.ahrq.gov/issue/sentinel-event-statistics-1995-2019. Accessed: 2024-12-01.

[37] Bal, B. S. An introduction to medical malpractice in the united states. *Clinical orthopaedics and related research* **467**, 339–347 (2009).

[38] Yu, C., Liu, J. & Zhao, H. Inverse reinforcement learning for intelligent mechanical ventilation and sedative dosing in intensive care units. *BMC medical informatics and decision making* **19**, 111–120 (2019).

[39] Bovenzi, I. *et al.* Pruning the path to optimal care: Identifying systematically suboptimal medical decision-making with inverse reinforcement learning. *arXiv preprint arXiv:2411.05237* (2024).

[40] Chong, P., Coleska, A., Kijpaisalratana, N., el Ariss, A. & He, S. 68 leveraging probability theory and machine learning to reduce diagnostic uncertainty. *Annals of Emergency Medicine* **84**, S32–S33 (2024).

[41] Kelen, G. D. *et al.* Emergency department crowding: the canary in the health care system. *NEJM Catalyst Innovations in Care Delivery* **2** (2021).

[42] Timbie, J. W. *et al.* Systematic review of strategies to manage and allocate scarce resources during mass casualty events. *Annals of emergency medicine* **61**, 677–689 (2013).

[43] Lin, Y.-K., Chen, K.-C., Wang, J.-H. & Lai, P.-F. Simple triage and rapid treatment protocol for emergency department mass casualty incident victim triage. *The American journal of emergency medicine* **53**, 99–103 (2022).

[44] Bhui, R., Lai, L. & Gershman, S. J. Resource-rational decision making. *Current Opinion in Behavioral Sciences* **41**, 15–21 (2021).

[45] Arimoto, S. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory* **18**, 14–20 (1972).

[46] Blahut, R. Computation of channel capacity and rate-distortion functions. *IEEE transactions on Information Theory* **18**, 460–473 (1972).

[47] Hale, T. *et al.* A global panel database of pandemic policies (oxford covid-19 government response tracker). *Nature human behaviour* **5**, 529–538 (2021).

[48] Cousineau, D. *et al.* Confidence intervals in within-subject designs: A simpler solution to loftus and masson's method. *Tutorials in quantitative methods for psychology* **1**, 42–45 (2005).

Extended Data

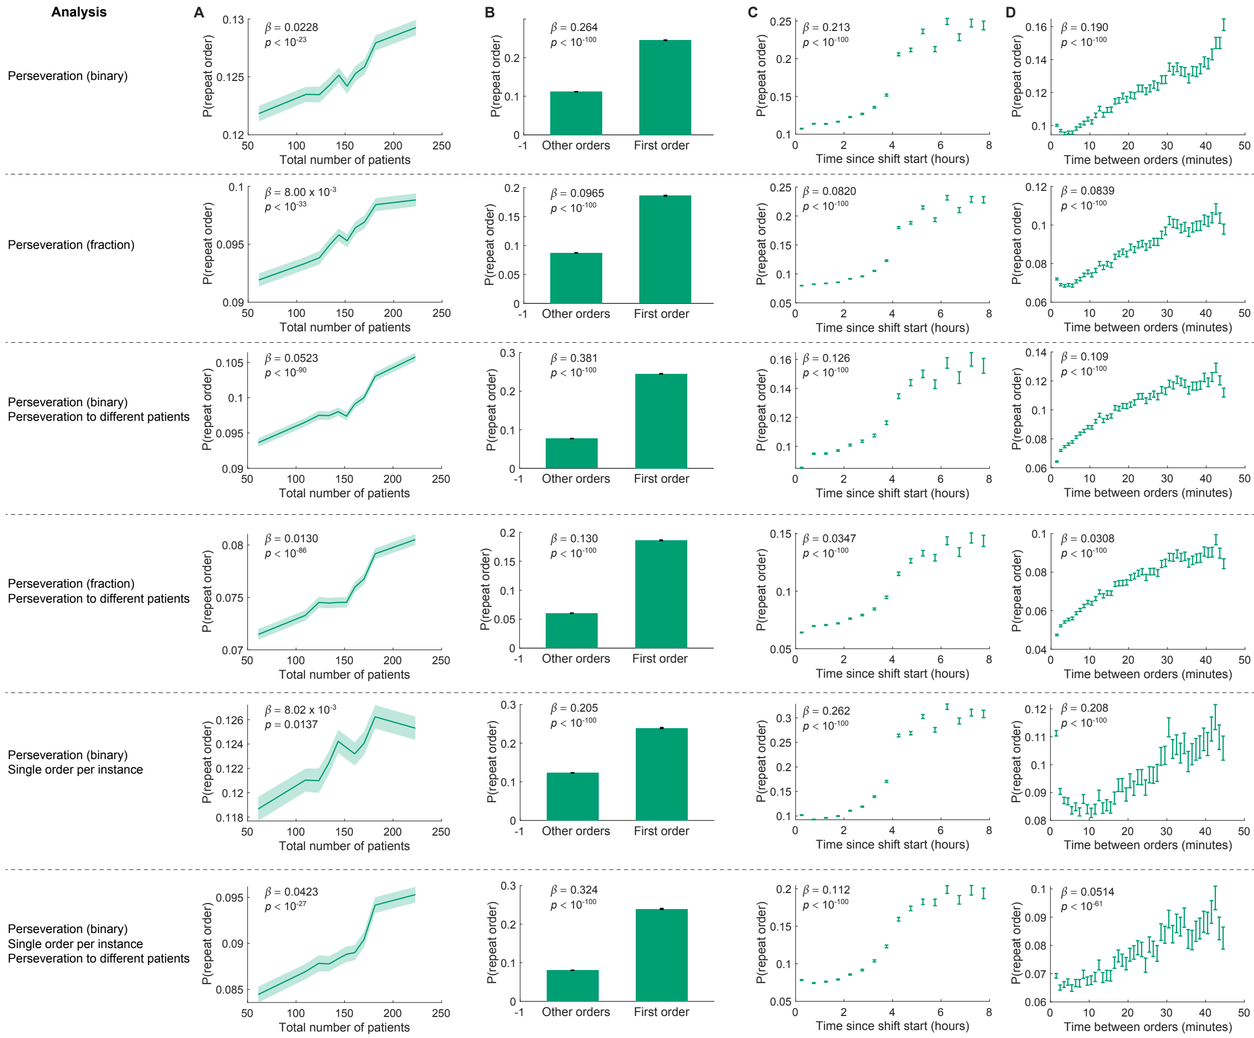

Extended Data Figure 1: Control analyses for perseveration results. The left row summarizes the control analyses for that row. See Materials and Methods for full details. Top row is the same analysis as in Figure 4. For each analysis, we fit mixed effects models (logistic for binary, linear for fraction) as a function of the graphed predictor and chance levels of perseveration. The inset lists the regression coefficient and p value for the graphed predictor. (**A**) Control analyses for total number of patients. (**B**) Control analyses for first order. (**C**) Control analyses for time since shift start. (**D**) Control analyses for time between orders.