RESEARCH
Inside IISE Journals

This month, we spotlight two articles from the April 2026 issue of *IISE Transactions* (Vol. 58, No. 4) that demonstrate how rigorous analytical modeling and optimization techniques can enhance decision-making performance. The first article tackles a common challenge in distributing treatment resources during an unfamiliar infectious disease outbreak: Should treatments be given immediately to everyone who requests them, or should some be saved for later? By leveraging a large-scale optimization model to consider highly uncertain demand and disease trajectories, the authors show that when treatments benefit certain age groups more, a carefully paced release policy is crucially needed to maximizing lives saved.

The second article tackles a key challenge in sequential decision-making under uncertainty: How to make robust decisions when only partial or incomplete information is available. The researchers introduced a novel DR-POMDPs, a framework that guards against worst-case scenarios, delivering reliable, high-performing policies. Their approach addresses real-world ambiguity in complex systems while providing actionable insights and measurable performance guarantees.

When should we release scarce medical resources during an outbreak?

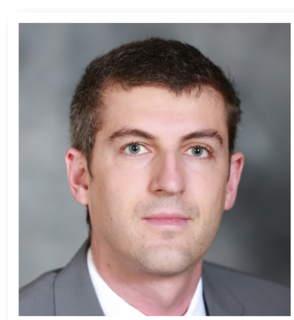

Steffen Rebennack

Bismark Singh

When a new infectious disease emerges, public health leaders face an uncomfortable dilemma: Should limited lifesaving treatments be released immediately or should they be rationed and held for later stages of the outbreak?

At first glance, releasing everything right away seems compassionate; if people are anxious and seeking help, why wait? But recently published research shows that early release can unintentionally work against long-term public health goals. In the early phase of an outbreak, it is hard to distinguish truly infected individuals from the "worried-well" – people who are not sick but seek treatment out of fear. When diagnostic testing is limited or unreliable, these worried-well users can consume scarce antivirals, respirators or other therapeutics that could later save lives.

The research comes from Bismark Singh, an associate professor in the School of Mathematical Sciences at the University of Southampton, and Steffen Rebennack, a professor at the Institute for Operations Research at the Karlsruhe Institute of Technology, as presented in their article, "Release Immediately or Sequentially? Strategies For Allocating Scarce Therapeutic Resources During Disease Outbreaks."

The authors built a large-scale optimization model to determine how a central stockpile distributes treatments to regions over time. The model incorporates three realities of emergency response: resources cannot be pulled back once shipped, demand and disease trajectories are highly uncertain and, ethically, no one who requests a treatment can be turned away if supplies are available. Using scenarios based on six historic pandemics and Texas-specific simulations, the authors compared two strategies: releasing all stockpiled resources immediately versus distributing them gradually.

Their results reveal a clear pattern. If a treatment has the same benefit for everyone regardless of age or risk, then releasing everything early is best. But when certain groups, such as older adults or those with chronic conditions, benefit more than others, holding some stock back until demand peaks can save thousands of additional lives. Sequential release helps ensure that high-benefit groups receive treatment "at the moment" it matters most.

For practitioners, their research conveys a straightforward message: timing matters. When treatments affect groups differently, carefully paced release policies – supported by rigorous analytics – can outperform simplistic direct approaches and strengthen the overall response to an unfamiliar disease outbreak.

CONTACT: Bismark Singh, b.singh@southampton.ac.uk, B54 School of Mathematical Sciences, University of Southampton, UK

Navigating the 'curse of ambiguity' in decision-making under partial information

Partially observable Markov decision processes (POMDPs) have been widely used as invaluable tools in sequential decision-making. They are well-suited for a wide variety of real-world problems in engineering, economics and healthcare among others, where only partial or incomplete information is available but there exist observations or signals which yield probabilistic belief about the hidden state.

However, the state dynamics and signal distributions are often estimated from limited or noisy data and are therefore subject to large statistical errors. This leaves the decision-maker with inevitable model uncertainty, sometimes referred to as the curse of ambiguity. For instance, in medical decision-making, the progression of the disease can only be inferred from symptoms and/or lab tests and is inherently ambiguous. A critical question posed by the curse of ambiguity is: How does uncertainty in model parameters translate into uncertainty in decision performance?

To answer this question and address the robustness of the decisions challenged by data inadequacy, Tong Li, a fifth-year Ph.D. student, and Yisha Xiang, Associate Professor, both from the Department of Industrial and Systems Engineering at the University of Houston, developed a new dynamic decision framework. In their article, "Distributionally Robust Partially Observable Markov Decision Processes (POMDP) with Distance-


Tong Li


Yisha Xiang

Based Ambiguity Sets," they use statistical distance metrics to construct ambiguity sets that contain possible distributions, formulate a distributionally robust POMPD (DR-POMDP) model, and seek the optimal policies hedging against the worst-case distributions within these sets.

The researchers first establish the convexity of the value function, which enables the development and use of efficient optimal search algorithms, thereby directly handling the "curse of dimensionality." They further provide a performance guarantee for the proposed approach by developing a probabilistic lower bound on out-of-sample performance, offering meaningful managerial insights.

The researchers demonstrate the utility of the proposed DR- POMDP through a comprehensive computational study in the context of a machine replacement problem. Their results show that the proposed new decision framework consistently achieves better out-of-sample rewards and higher reliability of the performance guarantee, compared to the conventional model that ignores parameter uncertainty.

CONTACT: Yisha Xiang, yxiang4@Central.UH.EDU, Department of Industrial and Systems Engineering, University of Houston, 4222 Martin Luther King Blvd, Houston, TX 77204-4007