

How Data-Driven Personalized Encouragement Guidelines Can Optimize Efficiency and Equity

Angela Zhou, University of Southern California

The technological advancements of the digital era have transformed Americans' lives in myriad ways. The provision of government services is one area where digital design and technology has made its mark: Consider how a service like GetCalFresh.org, an online application-assistance tool and resource hub co-created by the nonprofit Code for America and the government of the state of California, has streamlined the onerous process of applying for SNAP benefits. Using technology to make government benefits more accessible is a worthy project; increasing accessibility, however, does not necessarily ensure that those benefits are accruing to those whose lives would be most improved by their take-up.

In tandem with the rise of technological improvements in service provision, recent research documents the burdens that prevent all eligible individuals from learning about, applying for, and receiving public services to which they are entitled. It is crucial to evaluate the impact of service delivery interventions on equity: How do we make sure government benefits reach those who need them the *most*? My research demonstrates that personalized encouragement guidelines can help optimally distribute limited resources and maximize the impacts of resource allocation. These personalized guidelines can consider factors such as income level, geographical location, and historical disadvantage to target potential benefit-receivers whose lives would be most improved by using benefits.

High-Touch vs Low-Touch Support in Benefit Take-Up

Government and nonprofits have long partnered in interventions to ensure eligible populations receive benefits, such as employing counselors who support benefit-seekers navigating complex application processes. This type of "high-touch" assistance can be very effective, but also resource-intensive and therefore less sustainable. On the other hand, "low-touch," more universal support (such as mailing mandatory eligibility determination notices) may reach more people with less resource strain, but can perpetuate inequity. If an agency only uses eligibility determination notices and no counselors, then they may technically have conducted universal outreach but will have left behind benefit-seekers who do not have consistent access to email or a permanent address, or who require translation help.

For example, consider pilots of personalized letters and email reminders for enrollment in ACA zero-premium plans. California's individual ACA marketplace, Covered California, reached out to not-yet-enrolled low-income households that were eligible for low- or zero-premium plans to inform them of their eligibility. In addition to the eligibility notice sent to all, an experimental group received additional personalized letters and reminders. The researchers found that this low-touch personalized outreach could increase enrollment by 1-5% points, depending on the plan—but low overall rates of take-up suggested that higher-touch interventions, such as more costly phone calls with enrollment counselors, may be necessary for further increases.

Optimal Encouragement Guidelines Help Manage Costs and Reduce Disparities

My recent research focuses on how best to allocate resources to increase both access and equity in government service provision. I studied how *personalized* encouragement guidelines can leverage data on past individuals, their take-up of benefits (or rate at which others enroll them), and final causal effects of interventions to balance *efficiency* in resource allocation with *equity* in service access. I examined encouragement guidelines that optimized efficiency while ensuring fairness, defined as the reduction of disparities in benefit take-up; even if disparities are not equalized, this methodology can be used to find similar eligibility guidelines that trade some amount of efficiency to achieve greater disparity reduction. The Covered California example illustrates this equity-efficiency tradeoff: The optimal encouragement guideline

would be a cost-effective rubric for allocating low-cost interventions, such as personalized mailings, and highcost interventions, such as counselor phone calls, to result in the greatest measurable impact of increased insurance enrollment among the total targeted population. The methods I employ can allow benefit administrators to use previous randomized data to assess the utility of alternative encouragement assignment rules without having to run costly new pilots. Two case studies illustrate different applications:

The 2012 Oregon Health Insurance Study investigated what happens when more Oregonians got public health insurance through a lottery that randomly granted access to eligible applicants. Instead of complete randomization, a hypothetical personalized eligibility guideline might determine certain subgroups to prioritize based on their unmet healthcare needs; a fair and optimal guideline could determine which subgroups were most likely to improve their healthcare utilization due to being granted access while reducing disparities in further healthcare take-up.

Another study examined how judges decided if individuals were eligible for supervised release, a form of postprison or pretrial monitoring. A previous decision-making matrix recommended certain groups for release based on a general risk score quantifying likeliness to commit another crime, rather than actual potential varying causal effects of supervised release itself. Individual judges exercised a lot of discretion, and communities were concerned about racial disparities in outcomes. Here, my method could be employed to find a fair and optimal recommendation for supervised release, taking into account typical judicial decisionmaking patterns (for example, perhaps judges are more lenient for younger defendants) to find rules that balance expected efficiency of allocation against expected disparities in rates of who receives the intervention.

Modern digital technologies have vastly increased the potential for government services' data-driven optimization. We know that humans' varied life conditions and trajectories can affect the efficacy of policy interventions. When this is the case, personalized eligibility guidelines that target services towards those who benefit the most can improve the efficient use of resources. Personalization increases degrees of freedom, allows comparisons of equity-efficiency trade-offs, and can ultimately reduce disparities in benefit take-up. Methods like those I have developed are a promising tool for improving access to services while reducing administrative burden.

Read more in Angela Zhou, "Optimal and Fair Encouragement Policy Evaluation and Learning," arXiv, Cornell University, November 2023.