Evaluating Prioritization Policies for Scarce Permanent Supportive Housing

A discrete-event microsimulation of adults experiencing homelessness in Santa Clara County

Yupeng Chen, MS • Margaret L. Brandeau, PhD

Department of Management Science and Engineering, Stanford University

Stanford | ENGINEERING
Management Science & Engineering

informs HEALTHCARE CONFERENCE

Background

Adults experiencing homelessness face elevated risks of premature mortality, overdose, and avoidable acute-care use. Permanent supportive housing (PSH) improves housing stability and health, but capacity remains far below need.

HUD requires each Continuum of Care to operate a Coordinated Entry System that assesses, prioritizes, and refers households. Prioritization has relied on assessed vulnerability; newer proposals rank by predicted mortality or expected public cost. These criteria encode different objectives and select different people.

Most prior work studies assignment across program types. We study the decision that remains after eligibility: among adults already waiting, who receives the next opening?

Objective

To compare PSH prioritization rules within one dynamic system and characterize their effects on who receives PSH; mortality, QALYs, and time unhoused; queue performance; and the distribution of access and waiting across subgroups.

The aim is not a universally optimal rule, but the trade-off surface a coordinated-entry system faces.

Model Overview

We extend a published continuous-time microsimulation of homelessness, housing, opioid use disorder (OUD), health outcomes, costs, and QALYs, adding an open homeless population, PSH eligibility assessment, a coordinated-entry queue, stochastic openings, and alternative allocation rules.

Sheltered or unsheltered homelessness → **Eligibility and queue entry** → **PSH opening and policy selection** → **PSH placement and outcomes**

Each placement changes the system the next decision faces: who remains queued, the composition of those waiting, subsequent deaths and background housing exits, and returns to homelessness. This feedback is why a dynamic model, rather than a static scoring exercise, is required.

Data and Calibration

Calibration targets include initial queue and population composition, the sheltered/unsheltered split, homelessness inflow, PSH inventory and turnover, queue inflow, background housing exits, placements, returns, and mortality.

Sources include county Point-in-Time and Housing Inventory Counts, HMIS system performance measures, and published transition estimates. Background exits to housing are modeled separately from PSH placements to prevent double counting.

Methods

Simulation population and events

Individuals carry age, sex, housing state, chronic homelessness, chronic condition count, OUD and treatment status, and queue status. Events occur in continuous time: arrivals, queue entry and administrative exit, housing transitions, PSH placement, OUD treatment, overdose, return to homelessness, and death.

Eligibility sets who may enter the queue; prioritization sets who is selected when an opening arrives.

Prioritization policies

On the poster
First-come-first-served (FCFS)
Unsheltered first
Active OUD first
Highest mortality risk (oracle)

In the experiment
Random allocation
Chronic homelessness first
Age 55+ first; under 55 first
2+ chronic conditions first
Highest expected public cost (oracle)

Selection within a rule
Priority rules serve their tier in queue-entry order, so a targeting rule versus FCFS isolates targeting rather than confounding it with a change of service discipline.

When no queued individual meets a rule's criterion the opening is filled FCFS. Openings would go unused only if the eligible queue were empty, which never occurred.

Oracle rules observe model-implied risk exactly: upper bounds, not implementable.

Outcome measures

Health. Deaths before PSH, total deaths, overdose, QALYs.
Housing and exposure. Unhoused person-months, PSH retention, returns.
Queue performance. Placement rate, share still waiting, time without PSH, unused openings.
Distribution. Placement share and unhoused exposure by pre-allocation subgroup.

Experimental design and analysis

10-year horizon; 1:10 representation of county stocks and flows
25 paired replications per policy under common random numbers
identical initial populations, arrival streams, and opening schedules across policies within each replication
separate streams for mortality, overdose, transitions, queue entry, allocation
costs and QALYs discounted at 3% per year

Contrasts are paired differences against FCFS with 95% simulation intervals. Directionally robust means the interval excludes zero: a nonzero difference, not a beneficial one.

Per-person outcomes use the FCFS queue-cohort size in the same replication as a common denominator: the queue cohort is downstream of the allocation rule, so each policy's own cohort would put an outcome in the denominator.

Table 1. Queue-cohort outcomes, 25 paired replications

Policy	Still waiting at year 10, %	Deaths before PSH per 1,000	QALYs per person	Unhoused months per person
FCFS	11.6 (reference)	84.0 (reference)	1.697 (reference)	16.3 (reference)
Unsheltered-first	11.1 -0.5 [-1.0, +0.0]	88.4 +4.3* [+0.4, +8.3]	1.717 +0.020* [+0.003, +0.038]	15.9 -0.4* [-0.6, -0.1]
Active OUD-first	12.1 +0.5* [+0.1, +0.8]	64.3 -19.8* [-23.5, -16.0]	1.731 +0.034* [+0.016, +0.052]	16.5 +0.2 [-0.1, +0.5]
Mortality-risk oracle	12.3 +0.7* [+0.4, +1.0]	46.0 -38.0* [-40.4, -35.7]	1.737 +0.039* [+0.022, +0.057]	17.3 +1.0* [+0.8, +1.3]

Each cell: mean, then paired difference versus FCFS with its 95% simulation interval. * marks an interval excluding zero.

Results

Openings and placements are identical under every policy, so all differences are compositional. No rule dominates.

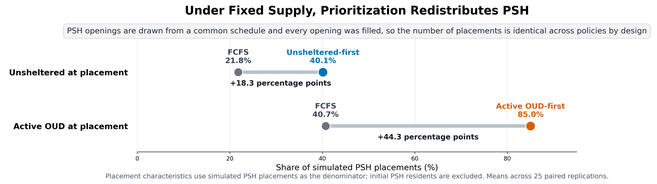

Figure 1. Under fixed supply, prioritization reallocates access rather than expanding it.

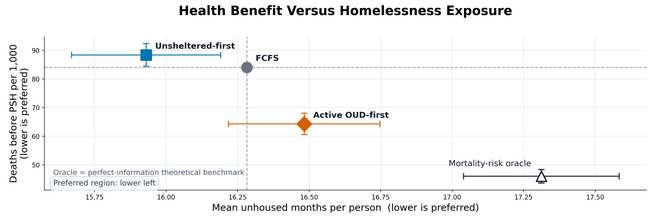

Figure 2. Health benefit versus homelessness exposure. Bars are paired 95% simulation intervals against FCFS.

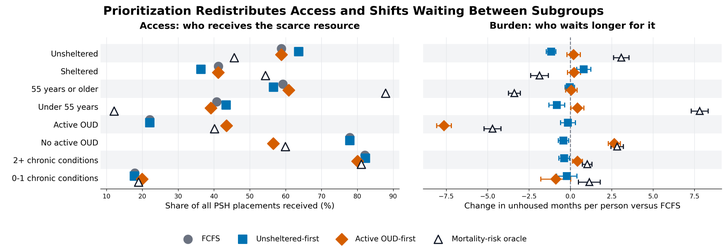

Figure 3. Access and burden by subgroup, defined at the pre-allocation assessment. Race and ethnicity are not represented in the model.

Discussion

Capacity is exogenous and the eligible queue never emptied, so throughput is invariant across rules: prioritization determines who is served and when, not how many.

No rule was dominant. Active OUD-first and the mortality-risk oracle reduced deaths before placement but left more households waiting at year 10; Unsheltered-first reduced homelessness exposure while raising deaths before placement (+4.3 per 1,000).

Visible unmet need and short-horizon mortality hazard identify distinct populations. Active OUD status, requiring no scoring instrument, captured 52% of the oracle's mortality reduction and 85% of its QALY gain.

Conclusion

Under fixed supply, prioritization determines who is served rather than how many. No rule dominated: reducing time unhoused and reducing death before placement selected different people and traded against each other, and every targeting gain was paid for by a subgroup that waited longer. A simple observable clinical indicator recovered much of the benefit of perfect risk information.

Limitations and future work

Race, ethnicity, and physical disability are absent from the model. Vulnerability is modeled characteristics observed without error, not a scored instrument. Supply is exogenous; single adults and PSH only.
Future work: prioritization under imperfect assessment, including measurement error and group-specific bias in vulnerability scoring, score-threshold rules, benefit-based allocation, and hybrid and wait-time-adjusted rules.

References and Acknowledgments

Gande et al., JAMA Health Forum 2026;7:e260624. Singham et al., Socio-Econ Plan Sci 2025;99. Shinn & Richard, Am J Public Health 2022;112:378-382. Data: HUD PIT and HIC, HMIS system performance measures, published transition estimates. Supported by NIDA Grant R01-DA064386.