

"Shopping Around": An Experiment in Preferences and Incentives for Placing Long-term Patients

VINCE BARTLE, Department of Information Science, Cornell Tech, USA

NICOLA DELL, Department of Information Science, the Jacobs Institute, Cornell Tech, USA

NIKHIL GARG, Operations Research & Information Engineering, the Jacobs Institute, Cornell Tech, USA

Hospitals and care homes devote significant resources to placing post-acute patients from hospitals into long-term care. This paper describes a two-phase experiment over SMS, conducted with a hospital in Hawai'i, in which care homes express preferences, indicate availability to accept patients, and express interest in patients. In the first phase, the treatment asks care homes to reconsider their stated preferences to better support matching. The second phase measures whether resulting changes in preferences increased how often homes express interest in patients that match newly stated preferences. First, to motivate and inform experiment design, we explore factors contributing to extended hospital stays for patients, uncovering how care homes' preferences play a major role in what patients they consider accepting. Second, we conduct a 16-week randomized controlled trial with 960 homes, where we experimentally probed, via SMS messages, homes' willingness to change their preferences to improve potential patient match recommendations. We show that inducing homes to reflect on their preferences increased the number of homes who changed their preferences by over 50%: 9.8% of homes who received our treatment changed their preference compared to 6.0% of homes in the control group (p-value = 0.0421). Third, followup interviews with 22 home operators highlight how preference malleability is shaped by a combination of design constraints and on-the-ground realities, such as load-balancing existing patient rosters. Finally, we discuss implications for real-world systems like ours that must balance constrained communication with situational complexity towards improving outcomes.

CCS Concepts: • **Information systems** → **Collaborative and social computing systems and tools**; *Decision support systems*; • **Applied computing** → **Health informatics**; • **Human-centered computing** → **Empirical studies in HCI**.

Additional Key Words and Phrases: Empirical Methods, Mixed Methods, Information Seeking, Medical Support, Clinical Health, Older Adults, Nursing Homes, Hospitals, Empirical Study, Field Study

ACM Reference Format:

Vince Bartle, Nicola Dell, and Nikhil Garg. 2025. "Shopping Around": An Experiment in Preferences and Incentives for Placing Long-term Patients. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW348 (November 2025), 31 pages. <https://doi.org/10.1145/3757529>

1 Introduction

In the US healthcare system, a small fraction of patients account for a substantial portion of post-acute hospital stays: although 2% of patients spend more than 21 days in hospital, this proportion accounts for approximately 15% of all patient days in hospitals [18, 34]. Over half these days are post-acute, where a patient has been treated for their acute care needs and transitional planning is completed [3, 70]. Research suggests that a leading cause of post-acute discharge delays is finding available beds in community-based care homes that offer nursing-home level of care [10], with

Authors' Contact Information: Vince Bartle, Department of Information Science, Cornell Tech, New York, NY, USA; Nicola Dell, Department of Information Science, the Jacobs Institute, Cornell Tech, New York, NY, USA; Nikhil Garg, Operations Research & Information Engineering, the Jacobs Institute, Cornell Tech, New York, NY, USA.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2573-0142/2025/11-ARTCSCW348

<https://doi.org/10.1145/3757529>

decision making among stakeholders, including care home operators, contributing significantly in time spent discharge planning [65, 67].

In this paper, we discuss how patient placement can be viewed as a digital marketplace [24, 40]: hospital staff use a computer-mediated system to find care homes that match patients in need of placement, and care home operators use the same system to find patients to fill vacancies. Based on this framing, we describe an experiment to collect and refine care home operators' preferences, with the goal of facilitating better matching and improving patient placement.

Our study is based on a three-year collaboration with a major hospital in Hawai'i, deploying an SMS-based system to facilitate long-term patient placement. This system aids care homes in filling vacancies and assists hospital staff in finding suitable accommodations for post-acute patients.

We first analyze historical data from our system to understand patient placement challenges. We begin with a regression analysis on data from 760 patients, correlating the length of hospital stay with their mobility levels (assessed by a nationally recognized ambulation score). Results indicate that homes often shy away from less mobile patients despite higher insurance payouts, suggesting that current incentives are insufficient. We subsequently contextualize our regression analysis findings with how care homes express preferences on our platform. We explain how a quarter of homes explicitly state preferences, concerning both patient conditions and general characteristics like sex and weight. Other homes express no preference, ostensibly accepting any patient; however, in practice, they often do not accept any patient but in fact, are evaluating multiple options, a phenomenon described as "shopping around." We categorize these behaviors as "underconstrained," where the expressed lack of preference does not reflect genuine interest, undermining efficient placement. Conversely, some homes are "overconstrained," specifying rare ideal patient profiles, leading to prolonged vacancies, when they would accept other patients. Our research aims to refine preference assessment methods to improve data quality and system efficacy in patient placement.

We conducted a two-phase randomized controlled trial (RCT) to determine whether expressed preferences reflect ground truth or are malleable. This study involved 960 care homes, equally split into control and treatment groups, with the treatment group further divided based on homes' preferences being potentially overconstrained or underconstrained.

Phase 1 of our experiment evaluated preference malleability. Homes in the control group received routine messages confirming their current preference status and prompting updates if necessary. By contrast, the overconstrained homes in the treatment group received tailored messages encouraging them to reconsider stringent preferences, suggesting that they might receive more matches if they do; underconstrained homes were asked whether they were interested in a patient that matched their preferences but who matched the least with other homes, which we describe as 'hardest to match' patients. **Phase 2** presented all homes with profiles of their hardest to match patients, to gauge interest and whether Phase 1 treatment messages increased matches.

Results from Phase 1 indicated preference malleability; homes in the treatment group were over 50% more likely to change their preferences (p -value = 0.04), suggesting some flexibility when confronted with alternative scenarios. Homes were more willing to adjust their weight preferences than sex preferences; the rate of sex preference changes was not statistically significantly different between control and treatment, while weight showed borderline significance (p -value = 0.054).

Phase 2 did not reveal significant differences in patient interest between the control and treatment groups: both groups demonstrated similar levels of engagement when presented with actual patient profiles. However, engagement among homes was notably higher when homes were provided with specific patient scenarios that closely matched, or otherwise challenged their stated preferences. We thus emphasize the importance of precise and contextual communication in eliciting meaningful responses from care homes. Overall, our RCT highlights the complexity of preference management in patient placement systems and suggests that while homes may be open to modifying their

preferences, significant shifts are contingent on clear, compelling communication that directly relates to their operational realities and patient care capacities.

To deepen our understanding of our experiment results, we conducted interviews with 22 participating care homes from both the control and treatment groups. We show how homes' preferences are shaped by a combination of economic need, load balancing, and operational risk. For example, we discuss how a home may have a given roster of patients around whom they must align new admits (e.g., a patient may not want to share a room with a patient of another sex). We also see how long-standing biases around patient aggression further contribute to differing malleability between sex and weight preferences.

We conclude by discussing how fundamental marketplace characteristics are facilitated through the preference solicitation and refinement mechanisms. We also highlight challenges and implications for systems that aim to collect and manage fine-grained, changing preferences to facilitate matchmaking in high-stakes contexts like patient placement.

Table 1. Summary table of our research questions and findings for the different stages of our study.

Stage	Research Question	Finding
Motivational Analysis (Section 3)	How does ambulation impact duration to placement? (Section 3.2)	The more complex the patients' needs, the harder to place. (Section 3.2)
	What preferences do care homes indicate and how do they impact patient matches? (Section 3.3)	Some homes have preferences, but many others do not state their preferences. Most homes' preferences shared on the system involve sex or weight, though some are more complex. (Section 3.3.1)
	Would homes benefit from suggestions to change preferences? (Sections 3.3.2 and 3.3.3)	Homes may be over- or under-constrained, and would benefit from changing preferences. (Sections 3.3.2 and 3.3.3)
Randomized Controlled Trial (Section 4)	To what extent are homes' expressed preferences around patient sex and weight malleable? (Section 4.1.1)	Homes change their preferences when probed to do so, and are more open to changing weight preferences than sex. (Section 4.2.1)
	Do refined preferences on our system improve homes' interest in patients that match their preferences? (Section 4.1.3)	Probing did not show a significant effect on patient interest, though it increased engagement. (Section 4.2.2)
Interviews (Section 5)	What factors might explain the differences we observed in how weight and sex preferences change? (Section 5.2)	Economic needs significantly influence preferences related to ambulation and weight, but not sex. (Section 5.2.1) Care homes must balance staff capabilities, patient needs, and roster composition, which influences weight-related preferences more than sex-related ones. (Section 5.2.2) Social desirability affects preferences, particularly for patient aggression, which strongly influences sex preferences over long time horizons. (Section 5.2.3)
	What factors might explain the absence of observable differences in patient interest between treatment and control groups? (Section 5.3)	SMS design constraints, including expectations about texting and the inability to capture conditional preferences, may explain why patient interest remains unchanged over SMS. (Section 5.3)

2 Related Work

Our work investigates the management of complex, changing preferences at scale in residential matching settings between hospitals and care homes. Here, we discuss relevant literature on marketplace design, healthcare operations, and in managing user preferences in digital marketplaces.

2.1 Marketplace Design and Healthcare Operations

Our setting can be considered a digital marketplace: care home operators use a computer-mediated communication system to find patients to fill vacancies, and hospital staff use the same system to find care homes that match patients they need to place. Digital marketplace design has long been an area of interest for CSCW [24, 37, 47], with research exploring design implications within sharing and gig economies [39, 41, 49, 57]. Research in residential marketplaces has also studied, for example, how discrimination might be amplified by marketplace design [66], opportunities for platforms to encourage boundary setting [35], and challenges platforms face in representing diverse residents with varying preferences [38].

Lampinen and Brown [40] propose an analytic vocabulary for understanding why markets may succeed or struggle, based on work by Roth [61]. Drawing on this vocabulary, the paper shows the importance of participants being able to share honest preferences without fear of consequence (*safety*). They also show how participants may prefer matches other than the ones they are assigned by a system (e.g., AirBnB), which undermines *stability*, requiring platforms to introduce mechanisms to prevent this behavior, such as binding contracts and financial penalties. Similarly, our work is motivated by observed behaviors where, although homes are offered patient matches that fit their expressed preferences, they continue to “shop around” for potentially better matches.

Further work has sought to understand the impact of marketplace design in healthcare settings. Research has shown that dynamic algorithms that account for psychological factors can stabilize patient-doctor relationships and improve satisfaction [45, 71]. Experimental evidence in mental healthcare settings similarly shows how algorithmic matching of therapists to patients might improve therapeutic rapport and treatment outcomes [15]. In the realm of patient discharge, Chan et al. [13] find significant improvements to readmission rates by utilizing algorithmic matching to identify matches between patients and facilities. Broadly, operations and healthcare researchers have developed decision-support frameworks that guide both *when* and *where* to route patients under uncertainty. Bowles et al.’s ‘DIRECT’ tool lowered 30-day readmissions by weighting patient characteristics closely tied with varying levels of care, to produce post-acute placement referral recommendations [7]. Similarly, authors have modeled Markov Decision Processes to forecast discharge timing and patient length of stay, by accounting for a variety of clinical patient characteristics [63] and by accounting for hospital bed congestion and long-term care bed availabilities [14]. Our work expands this body of literature by specifically highlighting the unique challenges in matching long-term care placement in rural settings with unique affordances around, e.g., accessibility of marketplace information. Among our findings, we highlight, for example, how preferences served over free-form text yield a significant signal for placement efforts and, further, that preferences can be shaped by the platform.

Substantial work has examined digital marketplaces that enable care work [31, 37, 46, 47]. For example, Petterson et al. [55] discuss how platforms supporting care work exist in the context of changing patient needs, and that those caring for patients must similarly adapt to these changing needs. Bellotti et al. [6] discuss the phenomenon of “out-of-market transactions” on platforms like Care.com, where participants often stop using the platform and instead coordinate directly. They suggest such behaviors are due to the formation of relationships outside the platform. Li et al. [44]

investigate older adults as active users of digital marketplace platforms and discuss their concerns and needs, pointing to a need for platforms to be designed in ways that they say are "age friendly."

Our work on placing patients from hospitals into care homes connects with this literature: a high proportion of care home operators are older adults [1, 19, 52], with prior work [67] discussing how accessible communication channels (i.e., SMS) may fit well with older adult' needs, although exchanging complex preferences over these channels is an ongoing challenge. Recent work has also explored accessible communication channels in the form of phone calls to support high-stakes work in rural settings, such as improving maternal health in India [68] or supporting caregivers of Syrian refugee children [56]. Our work explores how an SMS-based system might improve downstream call referrals for patient placement, contributing a better understanding of how seemingly coarse communication channels like SMS might enable supportive phone calls.

2.2 Managing Preferences in Digital Marketplaces

Research in market mechanism design has explored how participants can build, represent, and communicate preferences [5, 20], and how platform design affects participant preferences, incentivizes flexibility, and affects participant outcomes [11, 22, 23, 32]. Budish and Kessler [9] explore how students asked to provide ranked class preferences might accurately convey their preferences, calling for work on "optimal language design" to better collect accurate preferences.

Researchers, including at CSCW, have long studied mechanisms that balance accuracy-diversity trade-offs in recommendation systems—how systems should balance showing matches that best align with estimated preferences versus showing diverse content to prevent insular "bubbles" [21, 30, 53, 54, 59]. Research also raises concerns that recommendations may push people towards preferences that are only partly of their choosing; work has explored the roles that systems play in shaping preferences through recommendations, and cautioned that systems may lean more toward persuasion than recommendation [26, 27]. This phenomenon may be exacerbated by systems in which people are not able to report their preferences accurately, or where the system is more a contributing factor for what is recommended than the users' true preferences [9, 26, 27, 67].

Research has also examined user preferences for self-curation of recommendations [12, 28, 43, 60, 64], as well as ideas around "context-aware" preferences that facilitate social matching [50], discussing the potential utility of incorporating users' levels of passion around their specific interests to achieve better matches for individuals. These works highlight how users may strategically engage with recommender systems (for example, choosing which items to respond to), so as to steer their future recommendations. Our work engages with a similar problem in a novel and high-stakes context: adult foster care, which may increase the importance of such behavior with respect to operational and financial risk; homes may further submit preferences with certain expectations of how those preferences are used to match them with patients.

Although our work (to our knowledge) is the first to examine the management of preferences in adult foster care, several studies have explored preference management in other high-stakes domains. For example, authors have explored better managing preferences of adults with mental disorders re-entering the workforce [2]. More closely related, work studying refugee placement explores how to incorporate family preferences around location [4], showing a need to focus on a small set of preferences. Huber and McClelland [33] investigate patient and caregiver preferences around participating in the discharge planning processes, showing how stakeholders' preferences are often incongruent and developing ways to aggregate preferences to ease information frictions.

Finally, research has explored managing stakeholder preferences to improve outcomes in children's foster care. For example, researchers have examined matching and search mechanisms in children's foster care broadly, and design considerations for platforms to improve outcomes for children with disabilities in particular [16, 17]. Dierks et al. [17] note that a significant portion of

matches are not due to a lack of parents willing to adopt, but rather difficulties in finding specific matches, and how systematically taking into account disability preferences could prove beneficial in improving match speed and quality. Dierks et al. [16] similarly point to the utility of managing heterogeneous preferences to improve outcomes in the foster placement process.

To this literature, we contribute a preference aggregation experiment in the context of adult foster care that explores the utility of text-based probes, which in-turn yields more accurate preferences, while simultaneously showing that market participants may also explicitly express preference rigidity. Our work further details how some preferences may be more accurately expressed than others, and how systemic and on-the-ground factors play important roles in preference malleability.

3 Research Context and Motivational Analyses

Our work is a collaboration with a major hospital in Hawai‘i, where our system has been deployed for three years to help place patients who require long-term care. Our system helps care homes find patients to fill vacancies in a timely manner. A detailed system description is provided in [67].

Every three weeks, via automated SMS messages, our system asks homes for their vacancies and associated patient preferences. Homes may also be sent de-identified patient descriptions; if a home expresses interest in the patient, hospital staff are alerted to call about a potential placement.

Hospital staff use this information to identify homes willing to accept patients in need of placement. When a patient is medically ready for placement, hospital staff add patients to our system using an online form; patient information is strictly de-identified. For ease of use, we aligned the form’s design (e.g., attribute ordering) with existing forms used by staff. Form attributes further correspond to preferences shared by homes. This approach helped us avoid the need for complex EHR integration and provided safeguards for staff to prevent including identifiable patient information, as regulated by HIPAA in the US. We provide further detail on measures taken to safeguard data in Appendix A.4. We also note that the focus of our work here is on the SMS interactions; we leave the design of improved interfaces for hospital staff to future work.

Our system uses SMS since care home operators, who are often older adults, find SMS to usable and accessible. The following is an example of a text message sent to a care home, asking them to confirm or change their previous status in terms of vacancies and patient preferences:

Happy Monday, [NAME]! We are messaging to confirm your status: [STATUS, eg:] 0 medicaid, 0 private. No weight limit. Male OK, Female OK. Accepting: HMSA, UHC, OHANA, ALOHA, KAISER. If this is still accurate, please ‘confirm’, or please respond with changes, so we can find the best fit if a vacancy arises. Thank you.

In the above, “0 medicaid, 0 private” indicates that the home previously indicated they have 0 available beds, for either patients on Medicaid or using private insurance; the remainder of the message indicates their preferences for patient sex, insurance accepted, and weight limits.

Prior work [67] suggests that systems like this might reduce hospital staff workload by enabling up-to-date vacancy information, thereby reducing unnecessary phone calls between hospital staff and care homes. However, this work also highlights a need to better understand and refine the patient-care home matching process, which is the focus of our study.

3.1 The importance of care home preferences

As suggested above, an important system component is the communication of care home operator preferences. Preferences encompass a wide variety of characteristics, from the operator’s capability to work with specific conditions, e.g., supporting dialysis or tracheotomy, to more general preferences like the patient’s sex or weight. These preferences are essential to both hospital staff and

home operators, as they help determine which homes will be contacted for an attempted placement. In this work, we investigate these preferences and their role in patient placement.

In our system, 59% (562/960) of homes have expressed at least one preference over SMS, while the others have not expressed any preference. Both *not expressing* preferences and expressing *too many* preferences are potentially detrimental. For homes that do not express a preference, and hence are a potential match for *any* patient, one challenge is understanding whether these homes truly would take any patient. Hospital staff report that some homes may express that they will take any patient; subsequently, when they visit the hospital to assess a challenging patient, they ask to see other patients, without genuine interest in the patient they were called for. Hospital staff call this "shopping around", where a care home prioritizes identifying all their options over a timely placement with a patient matching their stated preferences. This strategic behavior presents a significant challenge, as hospital staff spend their limited time coordinating in-person visitations for homes that are not interested in the patient. In this study we refer to these homes as "underconstrained" in their stated preferences.

On the other hand, for homes that do express preferences, one potential challenge is the unintended consequence of being "overconstrained." In this case, a home may express an idealized patient description, for whom there rarely ever is a match and, in turn, the home must wait an exceptionally long time for a match, if ever there is one. Prior literature [67] has shown that care home operators who might not be aware of the consequences of their preferences may express frustration at never receiving calls about potential matches. Upon examining the homes' stated preferences, these homes were found to be overly constrained, with no patient matches. In such cases, a home may indeed be willing to take a patient that does not exactly match their preferences, since caring for a patient slightly different than their ideal patient may be preferable to a long-standing vacancy, which has economic costs for the home.

In both over- and under-constrained cases, there may be a gap between a home operator's stated preferences and the patients they truly would take into their home. This gap has both patient care and economic consequences for the hospital and care homes. In Section 4, we describe an experiment that probes home preferences to understand this gap and analyze ways to systematically improve data quality in our system.

In the remainder of this section, we provide motivational quantitative analyses regarding care home preferences and economic incentives. We begin by (a) understanding the relationship between patient complexity and duration to placement, which then informs (b) our motivation for prompting homes to validate or change their preferences.

3.2 The relationship between patient complexity and duration to placement

One goal of patient placement, among others such as delivering a high quality of care, is reducing the amount of time patients spend in the hospital once they are no longer receiving acute care (reducing post-acute length of stay).

We focus on the patients who require long-term care placement and are supported by long-term care insurance. These insurance payouts are designed to incentivize homes to take on more challenging patients [29, 51], largely informed by patient ambulation (e.g., can they get out of bed and move independently). Homes are paid less for highly ambulatory patients, who may be easier to care for; these patients are often called "Level One." Conversely, homes are paid more for low to immobile ambulation patients, who may be harder to care for; these patients are called "Level Two." Within this context, we ask: **how does ambulation impact duration to placement?** If low-ambulation patients have longer post-acute length of stay durations, then that may suggest that insurance payouts do not sufficiently incentivize care for those patients.

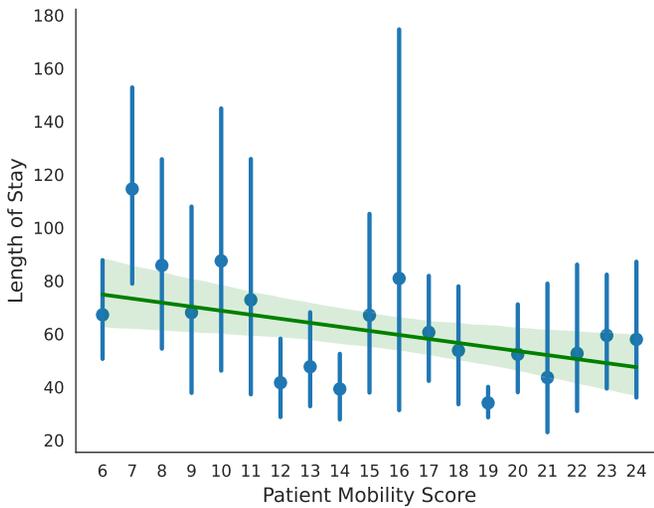


Fig. 1. Regression analysis of AM-PAC Mobility Score. The plot shows a negative relationship between patient mobility score and length of stay (regression coefficient -1.52 , and p -value < 0.011); see Appendix A.1 for details. As patient mobility increases, their length of stay at the hospital decreases. The shaded green region represents the 95% confidence interval of the mean length of stay for each mobility score value. The plot further contains intervals for each patient mobility score: the dot in the middle of each error bar is the average length of stay at that mobility score, and the length represents the variance, for that mobility score.

3.2.1 Regression Method. We use patient placement data across 760 patients who the hospital deemed as having or being at risk for an extended hospital stay between January and September of 2023. For each patient, we analyze two values: (1) an aggregate ambulation score, which is a nationally recognized standard for assessing patient mobility [25]; a higher ambulation score indicates a more mobile and/or independent patient and a lower score indicates a less mobile and/or more dependent patient. (2) A post-acute length of stay, representing the number of days spent in hospital after first being ready for placement.¹ We use an ordinary least squares regression to predict the length of stay from the ambulation score.

3.2.2 Managing risks and incentives. Figure 1 illustrates that patients who are less mobile and independent are harder to place, i.e., spend more time at the hospital. We note that, because care homes receive higher payments for accepting more challenging patients, this imbalance should, in theory, not exist. However, our regression suggests this is not the case: we see a negative relationship between mobility and length of stay. We further note that the greatest length of stay variance is at an ambulation score of 16, which is the beginning of an insurance cutoff.² Our preliminary analysis thus suggests hospitals must bear the added burden of placing patients who

¹We note that patient care statuses vary naturally, and so patients may start receiving acute care again after initially being deemed ready for placement. We start counting days to placement after *first* being ready for placement; hospital staff report that this length of stay metric is largely post-acute and a significant indicator used to self-evaluate placement efforts.

²In other words, some patients are placed quickly at the insurance boundary, suggesting that an increase in payment helps their outcomes. We reserve for future work a more rigorous treatment of an initial regression discontinuity analysis that we share in Appendix A.1.2.

are not appropriately incentivized by payors: incentives are not high enough to motivate homes to take on more challenging patients.

3.3 Motivation for prompting homes to validate or change their preferences

Understanding that, despite insurance incentives, less ambulatory patients are more challenging to place, we sought to understand how homes express preferences for patients, guided by the question: **What preferences do care homes indicate and how do they impact patient matches?**

3.3.1 Match Count Method. To answer this question, we analyzed care home-patient match counts in our system guided by finding how many patients satisfy each care home's preferences. For each home, we filter for only patients at the hospital whose features intersect with that home's preferences. At any given time, there are 30-40 patients, and homes may match with none, some, or all patients. We take these match counts, i.e., how many patients match with a home, and contrast them with potential match counts, *if the home were to relax a given preference*, e.g., around patient weight or sex. The number of homes and patients in our system fluctuates over time: homes shut down, new homes enroll, and patients are placed or otherwise newly eligible to be placed. For this analysis, we analyzed 960 homes and 38 patients.

Our system collects a wide array of preferences, varying from general (e.g., sex) to highly specific, (e.g., supporting dialysis). Importantly, most preferences can be categorized as sex (319 homes have stated a sex preference), weight (353), ambulation (102), or insurance (552) preferences. In our analysis of preference relaxation, we do not consider insurance because it is a relatively rigid constraint: a home may not freely certify with a new insurance provider. We note that, in future work, it may be important to better understand insurance constraints, as our findings indicate economic factors play a significant role (detailed in Section 5.2 and Appendix A.1.2).

Thus, in our analysis, we only consider sex and weight, the two largest remaining preference types. In our qualitative interviews, we detail the interaction between ambulation and sex or weight. In particular, ambulation can be captured to some degree by weight preferences, which is also reflected in the relative frequency of expressed weight versus expressed ambulation preferences. Nonetheless, ambulation and weight are not completely overlapping preferences, which highlights potential limitations of using SMS for capturing nuanced preferences, e.g., that users may find it easier to express a numerical weight than a degree of ambulation. We also note that preferences involve only a maximum weight (i.e., there is no minimum weight), though some homes may desire to use minimum weight preferences as a proxy for ambulation and monetary incentives. Lastly, we note that sex preferences are set as a binary for male and female. This binary representation has major limitations, which we discuss in Section 6.

3.3.2 Overconstrained Preferences. Our match count analysis (Figure 2) shows that 271 homes had preferences that restricted their match count to less than three patients, with 53 homes having zero possible matches. We view these 271 homes as potentially overconstrained.

We conducted a preference relaxation analysis, simulating what would happen if we (a) relaxed sex preferences so homes were open to patients of either sex, or (b) relaxed weight preferences so homes were open to patients of any weight. Of 271 homes that were overconstrained, i.e., had 2, 1, or 0 matches, we found that 260 would have at least doubled or increased from zero their match count from one of these relaxations (see Figure 3).

3.3.3 Underconstrained Preferences. On the other extreme, many homes who use our system either do not express preferences, e.g., 398 homes match every patient; (see Figure 2) or provide very few preferences, suggesting that they would accept any—even the most challenging—patient as long as they are adequately compensated. Yet prior work [67] suggests that when actually making

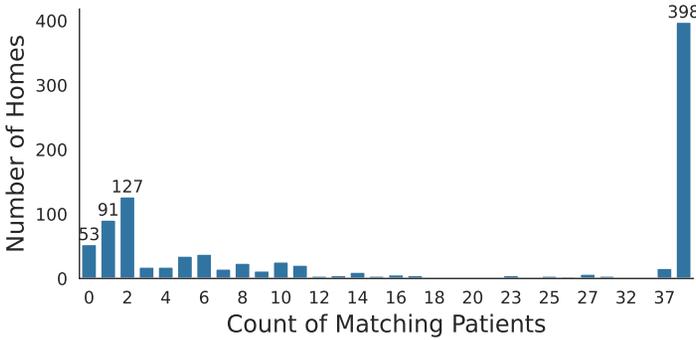


Fig. 2. Analysis of patient match counts for care homes in our system: 53 homes had 0 matches, which we call "overconstrained"; 398 homes matched with all 38 patients, indicating no patient preferences. For the purposes of our study, we consider all homes with 3 or more matches to be "underconstrained".

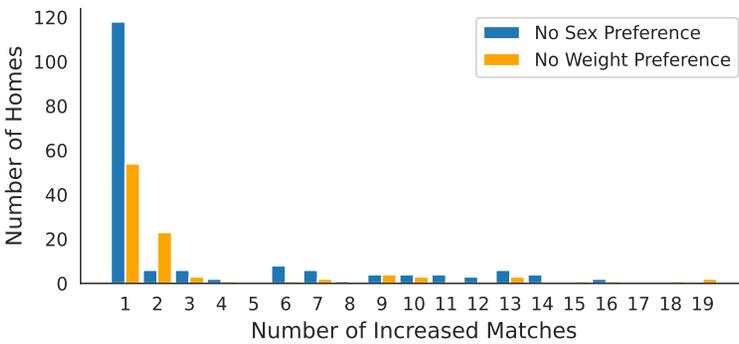


Fig. 3. A graph showing histograms for how many additional matches a home can receive if they fully relax their sex or weight preferences. For each of the 271 overconstrained homes (initially having less than 3 matches), we calculate which relaxation (out of sex or weight) would most increase their match count, and include the home only for that corresponding preference’s histogram. For example, about 20 homes would gain 2 matches if they relaxed their weight preference (and no more than 2 by relaxing their sex preference). 11 out of the 271 overconstrained homes do not benefit from either a sex or weight relaxation, and are omitted.

placements, the challenge of care a patient poses can deter homes: they do not end up accepting such a patient, and hold out for another patient, perhaps one that is still Level 2 but easier to care for. We consider these homes underconstrained, because the true number of viable matches is less than their saved preferences imply. If a home is underconstrained it can be hard for the hospital to decide what patient to call a home about.

3.3.4 *Experiment motivation and hypotheses.* The above analyses simply count the number of matches based on each home’s stated preferences; they do not establish that these preferences are not, in fact, the homes’ true preferences. Based on these analyses, we hypothesized that some homes are heavily under- or over-constrained and that these stated preferences do not match true preferences, and, by improving preference accuracy, we might also improve matching patient interest. We thus designed an experiment to help refine homes’ preferences and understand how

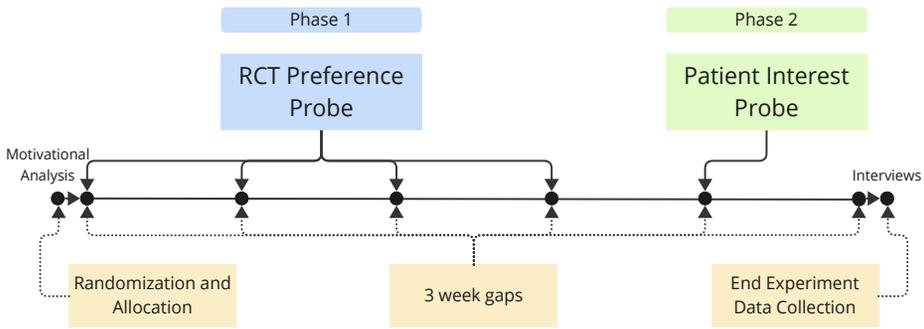


Fig. 4. Overview of our experiment timeline. Phase 1 envelopes four, three week intervals with one SMS-based probe of homes' preferences per interval. Phase 2 envelopes one SMS probe of patient interest, followed by a three week gap, after which we end experiment data collection.

they might change: whether specifically questioning a preference might shape that preference, and whether preferences were more likely to change than others.

4 Randomized Controlled Trial

With the hypothesis that homes may benefit from changing or otherwise validating their preferences, we sought to better understand whether homes' preferences are malleable. We conducted a two-phase randomized controlled trial to probe homes' preferences. In Phase 1, we explored **to what extent are homes' expressed preferences around patient sex and weight malleable?** Then, in Phase 2, we investigated **do refined preferences on our system improve homes' interest in patients that match their preferences?**

4.1 Experiment Methods

Figure 4 summarizes our experiment timeline. In Phase 1, homes were randomized to either a treatment or control group, after which they received SMS messages every three weeks, for a total of 12 weeks. The control group received the default standard message (as shown above), and the treatment group had their preferences probed (as detailed below). Then, in Phase 2, all homes received an SMS message with a potential matching patient's information and were asked if they were interested in assessing the patient. All procedures were IRB-approved.

4.1.1 Phase 1 Design. Randomization and Allocation. We begin our experiment allocation by accounting for responsiveness across our 960 homes. We order homes by how frequently they have responded in our system. We then conduct stratified sampling such that each of treatment and control has the same number of homes from the top half (most responsive) and bottom half (least responsive), with 480 homes in each of the control and treatment groups.

We further segment the treatment group to allocate messages that best align with potential preference changes for a home. As discussed in Section 3.3, homes can be characterized as either potentially overconstrained or potentially underconstrained based on previously expressed preferences. For our experiment, we defined potentially overconstrained as any home with match counts less than three (2, 1, 0), and all other homes (i.e., homes with matches greater than or equal to three) as potentially underconstrained. The selection of three matches as a cutoff is guided by the downstream heuristic shared by hospital staff: for any visitation or phone call with a prospective

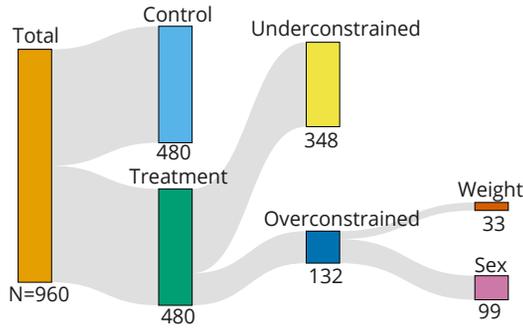


Fig. 5. Sankey diagram showing allocation of homes in our experiment into treatment and control groups.

home about a patient, hospital staff will strictly show or discuss three or fewer patients due to time constraints and concerns about homes that may “shop around.”

As shown in Figure 5, our treatment allocations break down into 348 homes who receive an underconstrained preferences message and 132 who receive an overconstrained preferences message. Overconstrained homes further break down into 99 homes who receive a message suggesting a sex relaxation, and 33 homes who receive a message suggesting a weight relaxation—determined by which relaxation would lead to more matches. For homes that would have gained the same amount of patient matches from relaxing either sex or weight, we randomly allocated a relaxation. We note that 11 homes among all 271 potentially overconstrained homes did not stand to benefit from any single relaxation in sex or weight; of these, 6 were in the treatment group; their message indicated that they would not receive a benefit from relaxing any preference by itself.

We now detail the specific messages homes received across treatment and control groups, with message templates in Appendix A.1.2.

Control Messages. For the past three years, our system has been sending text messages to homes every three weeks. The control group received the same baseline service that homes have received since they began using our system. During our experiment, an example exchange for the control group during Phase 1 might have looked like this:

Phase 1 | System: *"Good morning ... We are messaging to confirm your status: 0 medicaid, 0 private. No weight limit. Male OK, Female OK. ..."*

Phase 1 | Home_{control}: *"Thanks I have open bed looking for Medicaid female thanks"*

Treatment Messages. In addition to the message sent to the control group, homes in the treatment group were sent a separate message depending on if they were potentially over- or underconstrained.

Overconstrained homes. For a home that was overconstrained and, for example, stood to benefit from accepting patients who were one 10-pound weight difference away³, they received:

"In the last week you would have been eligible to match with 2x more patients if you accepted patients instead of at most 120 pounds, at most 130 pounds. Are you interested in assessing these additional patients? If not, please let us know, thank you."

³Our initial analysis looked at relaxation gains for no weight preference, however our eventual deployment revealed almost all homes would benefit from only a 10-pound relaxation, and some from less than 10-pounds. Nonetheless, the choice of a 10-pound relaxation for weight in particular is largely chosen on the ease of understanding this conveys, as opposed to a 7 or 13 pound relaxation. Nonetheless, we view more specific suggestions as a potential avenue for future work.

Similarly, homes who stood to gain more matches from relaxing a sex preference were suggested that their match count would have doubled if they accepted patients of both sexes. For the few homes who would not have benefited from a relaxation but were nonetheless allocated to the treatment group as overconstrained, our probe simply stated they would not gain on match count for any relaxation, e.g., "No relaxation benefit", and asked if the preferences we reflected were still correct. Lastly, for homes who would not have doubled their match count, and instead have gone from 0 to 1 match gain from a relaxation, they received a message stating "You would gain 1 match from no matches.", rather than the usual "You would double your match count."

Underconstrained homes. For the 348 underconstrained homes, we sent a message with a de-identified description of a patient that matched their expressed preferences, and that also matched with the fewest other homes (i.e., the most challenging-to-place patient that satisfied the home's preferences).⁴ For example, if a patient has only one other home they match with, they are considered more challenging to place than a patient who matches with every home. In brief, for all homes we calculate all their patient matches, then for all matches for all homes, we create a tuple such that the resulting list of matches for each home is of the form: (Patient ID, N homes matching patient,...). We then sort by least N matches with other homes.

We asked the home if they were interested in assessing the patient described and, if not, what characteristic did not work for them. In this way, we probed homes to refine their expressed preferences in our system in ways that better matched their true preferences.

For all treatment homes, once a home responded to the question probe we sent, they were no longer asked the treatment question and were reverted to the control message for the remainder of the experiment. If a home did not respond, we repeated the question in the next messaging cycle until either the home responded or the experiment ended.

4.1.2 Response Labeling. Our experiment aims to measure whether homes' stated preferences are their true preferences. We do so by probing homes' preferences to determine if they are rigid or if there is any malleability to change. The response variable we seek is thus: how many homes changed their preferences during the experiment. As the system's interaction with homes is via free-form SMS messages, measuring preference changes requires human verification of labels.

In total, we received 2,182 messages from homes during the experiment: 1,689 during Phase 1 and 493 during Phase 2. We detail how we semi-automatically labeled homes' SMS responses in Appendix A.2. In short, for each SMS message received, we labeled whether the home was changing a sex or weight preference. For the underconstrained treatment group in Phase 1, we also labeled whether the home was expressing interest in the specific de-identified patient sent to the home.

After labeling individual SMS messages, we measure our final experiment outcomes as follows. For our primary analysis, we determine for each home whether they changed *any* preference anytime during the experiment, i.e., whether any of their SMS messages were labeled with "changed weight preference" or "changed sex preference." For secondary analyses, we also count the number of homes that changed each type of preference separately. This data is aggregated in Table 2 and discussed further in Section 4.2.

4.1.3 Phase 2 Design and Method. The second phase consisted of a single wave survey where all homes in treatment and control groups were sent a matching patient description. This patient description was generated the same way matches were generated for underconstrained homes as described in Section 4.1.1. For homes who had already received a patient description, i.e., among the underconstrained homes, we (1) recalculated matches and (2) if the least-matched patient is

⁴As described below in the Phase 2 description, we randomize among the *two* patients with the fewest matches with other homes, to avoid biasing Phase 2 results.

Table 2. Summary of homes who responded across Phase 1 and 2, homes who made at least one sex or weight preference change, and homes who expressed patient interest. Columns include: (1) N, representing the number of homes in each of the control, treatment, and sub-treatment groups; (2, 3) Number of responding homes in Phase 1 and Phase 2; (4, 5) Δ Sex and Δ Weight in Phase 1, where weight is the maximum weight the home is willing to accept, and represent the count of homes that changed either sex or weight preferences at least once; (6) Homes Δ , indicating how many homes changed any preference at least once in Phase 1; and (7) how many homes expressed interest in the provided patient in Phase 2.

Condition	N	# Responding homes Phase 1	# Responding homes Phase 2	Δ Sex ≥ 1 Phase 1	Δ Weight ≥ 1 Phase 1	Homes $\Delta \geq 1$ Phase 1	Exp. Interest in Phase 2
Control	480	312	201	16 (3.3%)	17 (3.5%)	29 (6.0%)	40
Treatment (Overall)	480	297	219	21 (4.4%)	31 (6.5%)	47 (9.8%)	41
– sub-treatment (underconstrained)	348	211	150	7	22	25	20
– sub-treatment (overconstrained)	132	86	69	14	9	22	21

new, we send this patient, otherwise, if the least-matched patient is still the same as their initial patient, we send the second matching patient, thereby avoiding sending a duplicate patient.⁵

Our intent in this phase was to measure if, after potentially refining preferences, there might have been a downstream impact on patient interest. In other words, after changing preferences, did homes find the suggested patients better suited? As such, we label responses from homes who expressed interest, identically to Phase 1 labeling for the underconstrained treatment group who were sent de-identified patients.

4.2 Experiment Results

We now discuss the experimental results. Table 2 contains the number of changes of each preference type in Phase 1 and the number of homes who expressed interest in the given patient in Phase 2, for each of treatment and control.

4.2.1 Phase 1 Results. Running a χ^2 test on the number of homes who changed any preference during the experiment period, we find our treatment group was overall more likely to change preferences than our control group (p-value = 0.0421); this is largely contributed by weight preference changes with borderline significance (p-value = 0.0542), and that sex preference changes were not significantly impacted by our treatment (p-value = 0.5024).

In particular, as Table 2 shows, 6.0% of homes express a preference change in our control group during our experiment. By contrast, 9.8% of homes in our treatment group expressed a preference change during our experiment—inducing homes to reflect on their preferences increased the number of homes who changed their preferences by over 50%. The majority of the difference between control and treatment is being driven by changes to weight preferences, rather than sex preference. For sex: we see 3.3% change in the control group against 4.4% in the treatment group, where for weight we see 3.5% in the control group versus 6.5% in the treatment group. It’s important to note that these values do not sum to the prior percentages because, some homes may have changed both types of preferences, but are only counted once for the overall treatment effect.

Overall, our Phase 1 results to indicate that (a) our treatment did make a significant impact on preference changes, and (b) weight preference changes were more malleable than sex preference changes. We explore factors contributing to this difference via qualitative interviews in Section 5.2.

⁵We note that to ensure that this did not bias results, in Phase 1 we had chosen a patient description randomly among the top two hardest to match with other homes, for the overconstrained treatment group. This ensures that the patients received by the overconstrained group were not systematically easier to match in Phase 2, than patients received by other treatment groups.

4.2.2 Phase 2 Results. The second phase of our experiment sought to understand if more refined preferences improved homes' interest in patients. Our experiment yielded no statistically significant difference in patient interest counts between control and treatment. As detailed in Table 2, 40 homes from the control group and 41 homes from the treatment group expressed interest in the patient description they received in Phase 2.

One reason that the experiment may not have increased patient interest is that, while treatment homes *relatively* changed their preferences more (by 50%), overall preference change rates are still low, and many homes remain underconstrained. For example, only 25 out of the 348 underconstrained homes expressed preference changes during the experiment. Thus, most homes remain underconstrained, with no given preferences. It may be that changing downstream patient interest outcomes requires a stronger, more frequent, or longer-running intervention. A bigger sample size would also lead to an experiment with the power to detect more modest effect sizes; doing so would require either sending more patients to each home to gauge interest or a platform with more homes. We further discuss factors that may explain these findings, surfaced via interviews, in Section 5.3—especially regarding how homes view and expect to use our system. It may also be the case that preference management that leverages multi-modality or finer interactions over, e.g., a smartphone application may yield different results, in particular with respect to changing patient interest.

Despite not finding evidence of an increase in patient interest, exploratory analyses suggest the value of sending patient descriptions to homes over SMS. Interestingly, we found heightened engagement for Phase 2 wherein all homes were provided a patient match description: in Phase 1 the average response rate was 398 responding homes per wave; in Phase 2, which was a single wave, we received 420 responses. This difference suggests a demand from homes to receive patient descriptions over SMS. We further discuss the need to balance preference probing alongside sending patient descriptions in Section 5.3.

We also find several instances where preference changes led to specific patient matches. In the example below, a home was offered and rejected a hard-to-match patient description, and changed their preferences; subsequently expressing interest in a better-matching patient in Phase 2.

System_{phase1}: "Good morning ... Your preferences match the following patient: 46y/o MALE, 311 pounds. Ambulation is ..."

Home_{treatment}: "Female 120lbs below"

System_{phase2}: "Good morning ... Your preferences match the following patient: 72y/o FEMALE, 118 pounds. Ambulation is ..."

Home_{treatment}: Yes, I would like to assess her

5 Interviews

To shed light on our experimental findings and understand homes' responses, we conducted interviews with care home operators after the experiment ended. Our interviews sought to explore: **(1) what factors might explain the differences we observed in how weight and sex preferences change?** (Section 5.2); and **(2) what factors might explain the absence of observable differences in patient interest between treatment and control groups?** (Section 5.3).

5.1 Interview Methods

5.1.1 Recruitment and Participants. We compiled a list of homes, stratified by allocation to control and treatment groups. We called homes on this list and invited them to participate in an interview, continuing until we reached data saturation. In total, we interviewed 22 care home operators (Table 3). Our participants included 11 homes from each of the control and treatment groups.

Table 3. Care homes contacted during the interview stage of our study.

Condition	N	Total Responses Ever	Responses During Experiment	RCT Preference Changes	Expressed Patient Interest
Control	11	263	38	4	5
Treatment	11	278	26	5	7
Total	22	541	64	9	12

5.1.2 Interview Procedures. We conducted semi-structured phone-based interviews that explored participants' use of our system and experiences with patient placement. Participants were told that participation was optional and non-participation would not adversely impact their involvement with the system. All homes were asked for feedback about the system; for example, message frequency, if there was information they would like to receive more of, and if they were aware the system was used to surface patient matches. Homes in the treatment group were additionally asked about the specific probes they received, for example, if they rejected a suggestion to relax their preferences, why they kept their preference, and what they thought of the suggestion. Interviews lasted 15-20 minutes and were audio-recorded with consent.

5.1.3 Interview Analysis. We collected a total of 6 hours of interview data across the 22 participants, which was transcribed and analyzed using thematic analysis, adapted from Braun and Clarke [8], wherein we took multiple coding passes through each transcript. The first author conducted the analysis and synthesis of these transcripts, which first resulted in 57 codes, which were consolidated into 41 codes. Example codes included "reject male because aggressive", "patient cognition", and "sex because cognition concern." We then clustered our codes into five themes, for example, the prior three codes helped construct "Operational Risk, Sex" as a theme.

5.2 Factors that might explain differences in changes to weight and sex preferences

Our interview findings suggest that changes to homes' sex and weight preference are shaped by many factors, including (a) their fiscal needs, (b) their infrastructural capabilities, and (c) managing unknown risks, while contending with the reality that mismatched patients can be dangerous for the patient and themselves. In each of these areas, we find weight preferences are subject to a greater breadth of factors and at comparatively shorter timescales. By contrast, additional constraints, including shared rooms and managing operational risk, lead to sex preferences being less likely to change, with changes happening over longer periods.

5.2.1 Economic Need. Home operators emphasized that caring for patients is the core of their work and it's *"not about the pay, it's more about the service"* (P10); however, operating a care home is nevertheless an expensive endeavor for small residential homes who have undergone significant retrofitting measures to accommodate complex patients, or otherwise must account for *"the [high] cost of living here"* (P7); thus, economics play a significant role in decision making: *"It is a business, you know"* (P9). Homes must therefore balance patients who (1) might be easier to care for, against patients who (2) might be more challenging but come with greater financial incentives.

In our setting, there are two types of patients: "level one" and "level two". A *"level one is easier and level two is hard, but the pay is better"* (P17). For some homes, as one operator said: *"if it's a level one, we cannot survive as a [business]"* (P18). These classifications are important enough that homes may hold out on accepting a patient in the hope of finding a well-matched level two patient. However, holding out also results in lost income, so level one patients are sometimes placed because: *"[Homes] have been waiting for too long and just need a placement"* (P18). One participant

explained this preference as *"Level One would be like, I can ... but if there's a Level Two, then I would take that instead"* (P11). Another explained:

"If I'm tired working, then I like relax, then maybe better to admit a level one ... but for now, I'll wait it out, you know, I'll wait one month. If I have a choice, I rather [take a] level two, and if I don't have a choice, I'll take the level one" (P7).

This distinction between patients, contingent on economic incentives (i.e., insurance payouts), is relevant to our study because these insurance levels are in part informed by patient weight, as an extension of general ambulation; in short, heavier, less ambulatory patients are more likely to be classified as level two. Thus, homes' preferences for level two patients imply a preference for heavier patients.

In parallel, there are competing factors that result in preferences for lower weight patients. In many cases, if the patient can move on their own, the caregiver's work is easier: *"If I put [the patient] in the shower, can they stand up [independently], and maybe go two or three steps, and go inside the shower alone. So long as I don't [have to] lift them up, since I cannot"* (P20). The need to help a patient get into the shower or out of bed—which depends on a caregiver being able to lift, carry, or otherwise hold the patient—provides a basis for setting a weight preference. Thus, weight can also act as a determinant in rejecting a less ambulatory level two patient, even though a home might receive more pay for them, *"Like ... I wrote on [SMS], as long as they are under the weight. But otherwise, no. Not really considering [even if] level two"* (P11).

Although participants discussed a general preference for more ambulatory patients who might weigh less, they also highlighted challenges with such patients. In particular, homes described how ambulatory patients with greater independence may come with wandering or flight risks, *"I'm not looking for those that can walk around, because it's too hard, they can go away, they do go away"* (P20). Such patients may end up creating more work for caregivers: *"You know, then it's a 24 hour watch, it's kind of harder."* (P12). One participant explained:

"They think it's always like the same pay for the same difficulty of patient, so less pay for a level one, but it's actually more difficult for the level one, because they move, they walk, they can escape, you keep on guarding them. We have had several clients before that went to the street, I didn't know which way they went, I had to drive around to find him." (P18)

In other cases, home operators were more ambivalent about patient wandering: *"Wanders is okay, we can just close the gate and all that."* (P1).

These findings underscore how a patient's weight may connect to numerous variables that impact a home's preferences. Overall, setting a weight preference is in part an economic decision: a home must weigh insurance incentives against their ability to handle ambulatory edge cases, where they must physically bear the load of their decision. By contrast, there are no differing financial incentives for accepting patients of a particular sex. Thus, homes' sex preferences are less influenced by economic factors, and more influenced by factors like load balancing and operational risk, as discussed next.

5.2.2 Load Balancing. Although a home's physical infrastructure is relatively fixed—the home is built or otherwise retrofit to be licensed to operate as a care home—its supporting staff and the current roster of patients at the home are variable. Home operators must make patient placement decisions that balance the make-up of their current patients with what their staff can handle.

Having staff available to help is critical for heavier patients, especially because a significant portion participants self describe as, for example: *"I'm a petite person"* (P16) which leads them to set a weight preference, *"because I cannot lift the [patient] up"* (P17). With more staff available, managing a heavier patient becomes more viable:

"I want someone ambulatory because then for weight, I don't mind if they are 180 pounds, because I have an [assistant] around, so just in case I have somebody to help me" (P14).

Nonetheless, in the context of these small residential care homes, staff are rarely full-time at the home, and thus staff availability is variable:

"I have other caregivers, but not all the time. They [work other jobs]. So I like patients that I can manage by myself ... otherwise I'm not going to [be able to] carry him." (P20)

Even if staff were available to help, operators stressed the need for careful assessment of a patient: *"I have caregivers to help, but, you know, if too heavy...it depends, I need to assess" (P19).* By contrast, when asked if the availability of supporting staff informed patient sex preferences, homes generally said: *"No, for sex it's no difference, just if they are aggressive" (P4).* We further discuss challenges related to expected patient behaviors and sex preferences in Section 5.2.3.

Home operators discussed the need to balance their roster of patients according to their staff availability and the characteristics of other patients in their home. For example, homes may be more amenable to accepting a level one patient if they already have one or more level two patients, and thus their economic needs are met. Moreover, although homes are paid more for level two patients, having all their patients require multiple staff on hand to be removed from their beds might be a significant hazard for a home:

"I'm not supposed to have three maximum assist patients in case of emergency, it's impossible to remove them from the house quickly when you have to like lift everybody up. So just for emergency reasons, I have two level two patients [now], so I need [a patient] that's not too, you know, complicated, because I already have hard clients right now. So I am just looking for someone to fill the bed, and it's a shared room." (P14)

Load balancing by current patient roster is further complicated in shared room situations. A home operator may host a shared room to make the most of a relatively small space that provides a nursing-home level of care. Our interviews suggested that shared rooms were one of the few factors that played a big role in influencing homes' sex preferences. If a care home operator has a shared room with two beds, and the first bed is filled with some sex, they are inclined to fill the second bed with a patient of the same sex, in the hope of mitigating roommate differences:

"I sent a message two days ago that I need a medicaid male because I have male. It is a shared room, because ... the room is big, so it has to be shared, and because I have one male there already, I need male. It cannot be female because the other bed is male." (P19)

Moreover, because patients are placed for long-term stays, it is unlikely that the sex preference associated with a shared room will change frequently. If one patient is no longer at the home, the other is still there to uphold the sex preference associated with that room. By contrast, we saw that for private rooms sex does not play a role: *"In the private [room] either female or male is OK" (P18).*

While balancing with current patients in a room largely influenced sex preferences, weight preferences were also involved. A home operator noted concerns about a patient who is:

"Over 180 pounds because they might not fit well in the bathroom ... my other patients I have a toilet guard, and a [heavier] patient ... even if they are ambulatory, I have to take out the toilet guard, in and out, in and out." (P18).

In sum, we saw that home operators must make patient placement decisions that are informed by how much support they can expect to receive from staff, the present difficulty of their patient roster, and the layout of patients across their rooms. We find that these factors have a stronger influence on changing weight preferences over time; where sex preferences are influenced, the timescale of change is significantly longer than for weight.

5.2.3 *Operational Risk.* Another factor we discovered that influenced homes' preferences around patient sex was the need to balance operational risk. In particular, our interview findings suggest that many home operators use patient sex as a proxy for the likelihood that a patient may be aggressive. In short, home operators perceived female patients were less likely to be aggressive, and so expressed a preference for female patients. Although some homes said that they do not have a sex preference, other than that the patient is not, for example: "*combative and stuff like that*" (P14), for some care home operators this played an important role in their safety:

"Because sometimes [patients] kick, kicking up, you know, if it's a lady or female, it's kinda light, but if it's a man, very strong, especially violent, we cannot hold them. If it's female at least we can handle it" (P2).

For some homes, this was more of a general preference, such as for P17 who said, "*I'm just more comfortable with the ladies.*" Other homes shared serious examples of risks they were averse to:

"We had a guy who was demanding, he was looking for a refrigerator, looking for food. And in another case [in the area] you might have heard, a lady was preparing food, and the guy went to the kitchen, and he grabbed the knife from her, and she died. He was a level one – you know, maybe if he was level two, even if they want a knife, at least he would be stuck in bed." (P3)

While stories like these were relatively uncommon in our data, they underscore the challenges of this frontline care work and the important role preferences play for homes seeking appropriate patient placements. Moreover, we observed that using patient sex as a proxy for managing operational risk again resulted in sex preferences that changed less frequently, since these types of cultural beliefs may span long windows of time.

5.3 Factors that might explain patient interest across treatment and control groups

In Phase 2 of our RCT, we sent all homes a patient description that matched their preferences and asked if they were interested in assessing that patient. The responses we received did not show a detectable difference between treatment and control groups, suggesting that having more accurate preferences does not necessarily lead to higher rates of patient interest. Although this finding may be due to a lack of statistical power and the large number of homes that remain underconstrained, we also note several factors that could explain why the treatment and control groups may exhibit similar levels of patient interest. First, homes' interest in assessing patients is already relatively high: if a home has a vacancy, they will often want to assess any available patient. Second, many homes simply do not have any vacancies, and so would not be interested in a new patient, even if the patient matches their preferences, simply due to not having space.

We initially expected that homes in the treatment group, who may have better matches after our experiment, would have been more likely to express interest in a matching patient than homes whose preferences may not be as refined. Instead, our experiment resulted in similar levels of patient interest regardless of condition. Our interviews suggest that while a match might have better connected with homes' preferences, homes felt there was still enough ambiguity about the patient that they were unsure about ruling the patient out. As a home operator in the control condition said: "*When I assess, then I know. Like last I received a text [about a patient], I didn't want to [reject the patient] without more information ... then I called about them and we placed them*" (P10).

More broadly, it was clear from our interviews that the complexities involved in overall patient placement processes mean that responding to an SMS inquiring if a home is interested in a patient is only the first of many steps in an eventual placement. Care home operators emphasized the need to "*talk to the social worker over [the phone]*" (P9) about a patient, and that closing the loop on a placement would require that they "*look at the patient to assess her, come and actually see.*" (P18).

Care home operators further explained how some patient characteristics are difficult to discern through plain text; for example, one home wanted to assess *"if the [patient] can hold the [hospital bed] rail when I'm caring for them"* to determine just how much assistance they would have to provide the patient, because *"That's a big help, if she can't do that it's gonna be a problem, because she will need more people to assist"* (P9).

We also found many homes had conditional preferences that could not necessarily be captured by our system, further contributing to homes' desire to assess a patient regardless of improved preferences. For example, several homes who were sent a message about relaxing a preference, for example *"Would you accept patients above your weight preference?"*, responded with more questions that assumed the system was asking about a specific patient, such as: *"Is this a private client?"* (SMS). When we asked home operators to clarify, they explained that their preference was conditional on other details, such as if the patient was a level one or level two patient, and if they were self-funded or otherwise private-pay funded: *"If they are private pay, then level one is ok"* (P8).

Regardless of their expressed interest in any given patient, homes appreciated the connection with the hospital that our system provides, *"It's good that you guys are always checking our [preferences], and asking if there is space, because then we save time"* (P12). Home operators also noted that it was valuable to receive information about potential patients, even if they ultimately did not respond: *"I don't have time to make comments, but it's helpful for me as a caregiver to know about what patients [there] are"* (P8). The constraints of communicating via SMS also led to an expectation that message responses would be brief and relatively simple: *"I just say 'yes' or 'no' ... I don't really have enough time to be putting in a message, or leave [an explanation] on my phone"* (P12).

Nonetheless, other homes we interviewed asked if we might be able to send additional clinical details over SMS to help decision-making. Homes requested that they be able to: *"See the information like the diagnosis, and if you can text some of the medications he's under, or he's taken before"* (P11). Other homes similarly asked: *"What's his diagnosis, you can text me the diagnosis, and then I'll call you."* (P6). Of course, sending such information via SMS would require caution due to patient privacy; nevertheless, these findings suggest potential opportunities to explore ways to send more detailed information to homes to better aid patient assessment.

6 Discussion

Our findings explore the role of preferences in the early stages of patient placement. We show how imperfect preferences may inhibit match-making by being either under- or over-constrained. To ameliorate these challenges, we conducted an experiment to explore the malleability of homes' stated sex and weight preferences. We demonstrate how preferences may be malleable, that malleability may vary across preference types, and that accessible, easy-to-use mechanisms like ours may help reveal these preferences. Finally, we discuss on-the-ground factors that impact homes' preferences, including economic incentives, load balancing, risk, liability, and more. Here, we synthesize these findings to discuss (1) how fundamental marketplace characteristics are facilitated through mechanisms we design in service of marketplace success; (2) challenges and opportunities for managing preferences, and (3) implications for similar work in high-stakes contexts.

6.1 Patient Placement as a Digital Marketplace

Placing patients into appropriate care homes is an enormous challenge [18, 34]. Prior work has stated how in an ideal world, patient placement might "be pretty objective": any home that is licensed to provide a needed level of care simply accepts a patient who needs that level of care as soon as possible [67]. However, in reality, as we discuss in Section 5.2, accepting a patient who will live at the home for ostensibly a long period of time is an important and complex decision influenced by economic factors, care staff capacity, current patient needs, the operational risk of

patients who may be aggressive or wander away, and more. Thus, home operators have preferences for what patients they accept and want to do their own assessments.

As such, patient placement processes and home operators' preferences function as a computer-mediated marketplace, and our work connects closely with research on managing preferences and understanding marketplace dynamics in sensitive contexts. Lampinen and Brown [40] provide a framework for analyzing how marketplaces may or may not succeed and call for the deliberate design of marketplaces as a collection of social artifacts rather than naturally occurring phenomena. They describe a marketplace as characterized by: volume of active users (*thickness*), crowding which inhibits match finding (*congestion*), participants being content with their match outcomes (*stability*), sharing honest preferences without fear of consequence (*safety*), and transactions being rejected due to moral or ethical concern (*repugnance*) Lampinen and Brown [40].

Our study details the design and efficacy of new mechanisms that engage with these factors. For example, our probing messages enable refinement of preferences, especially for homes that are potentially over- or under-constrained, which encourages engagement and improves *thickness*. In addition, care homes' "shopping around" behaviors, in which they are constantly looking for better matches than those suggested, may both undermine *stability* and increase *congestion*. Our work to refine homes' preferences may help, by providing homes with better matches that lead to increased interest, thereby reducing homes' need to "shop around."

Our work also contributes a nuanced perspective on *safety*—the challenge of enabling participants to honestly share preferences. Lampinen [38] discuss how homes that share hosts struggle to accurately represent a home with multiple people on a platform where an entire home has a single profile. Similarly, our findings in Section 5.2.2 show how home operators must balance what new patients to take in, depending on what patients are presently in the home, who themselves might have strong preferences, such as for new patients to be of the same sex, especially if sharing a room. These findings suggest a need for future systems in similar contexts to be able to accommodate granularity in preference tracking, where homes may be able to better express more specific preferences in ways that allow preference changes to be a natural part of the placement process, rather than treating preferences as discrete, let alone an altogether ignored factor.

Finally, our work provides a different angle from which to consider marketplace *repugnance*, where monetary transactions may be avoided due to moral or ethical concerns. In our setting, scrutiny may be deserved around the monetization of patient care. Our findings in Section 5.2.1 show numerous complexities around monetary incentives for accepting more challenging patients. Resulting preferences may further reflect cultural biases on who is considered a challenging patient. On the other hand, these preferences play a significant role in safeguarding continued operation as a care home taking on high operating expenses in a context with a high cost of living. Altogether our work builds on prior CSCW work discussing challenges in mixing financial compensation with hosting people in private homes [35]. Our system is situated as part of a set of supportive tools in an ecosystem that tackles these monetary challenges. Insurance incentives play complex and predefined roles that we seek to help marketplace participants navigate. In line with ethical challenges associated with mixing monetary transactions with patient placement, we also note that most actors in this system are paid, including social workers. The system presented here is also funded by the hospital, whose management aims to reduce burnout, improve workflows, and improve placement outcomes.

6.2 Challenges and Future Directions for Preference Management

Our work takes initial steps towards tracking the true state of a diverse set of changing preferences. One challenge is balancing the complexity of collecting and possibly probing preferences with the practical need for simplicity in communication channels like SMS. Corroborating prior work [67],

our findings suggest that participants benefit from the ease of use of SMS, where homes expect quick and easy communication that enables the system to fit into their workflows. This is also in line with Mayer et al. [50], who discuss how people might respond to match recommendations in social settings with a simple "yes" or "no." During interviews, the authors found similar findings to ours (Section 5.3), where more granular conditional preferences exist than the platform currently captures, such as "yes, but not now."

More broadly, although homes find SMS usable, the information that can be exchanged quickly is limited. Thus, by prioritizing ease of use, the system currently does not capture more conditional or detailed preferences, and participants might be deterred from sharing preferences if they find that their specific preferences are not being well captured. Our findings in Section 4.2 also show that many homes still do not express preferences, and some homes that do may still be under-constrained. This highlights a need to balance capturing more fine-grained preferences while also accommodating participants with *few* stated preferences.

Our findings also show how preferences may be proxies for more latent concerns. For example, Section 5.2 shows how homes often expressed a preference for female patients as a way to manage concerns about potentially aggressive patient behaviors. Such proxy preferences present a challenge, since there may be patient matches that might be prematurely filtered out (e.g., male patients who do not present behavioral challenges). Our work to refine preferences (Section 4.2.1) may help by probing homes to understand if their preferences are malleable or rigid, which may also avoid later mismatches of preferences during a visitation. Moreover, the difference in malleability of sex and weight preferences that we uncover suggests the capacity for mechanisms like ours to facilitate the preservation of important preferences while also revealing malleability.

Importantly, our findings in Section 4.2.2 indicate that homes are still interested in assessing matches regardless of preference refinement: homes in both control and treatment groups had roughly equal interest in the patients we ultimately recommended as matches. One possibility is that home operators did not yet trust the system to predict accurate matches, which might also change if, for example, the system consistently sends reliably strong matches for a long time. Another possibility is that, for home operators, assessing patients who are ultimately not strong matches may provide useful opportunities for them to reassess and reinforce their preferences. Future work is needed to further explore these possibilities.

A central design choice in our system is the use of natural language to facilitate the exchange of information about preferences; as such, our system makes use of a natural language agent to manage matching, which has proven successful in terms of convenience and usability for participants. Several components of our system are automated to facilitate the scale at which we perform this matching and manage preferences, such as some automatic parsing of preference updates and tagging messages for human labeling (discussed in Section 5.3); however, managing increasingly specific preferences over long periods remains a challenge. To this end, future work could leverage advancements in language modeling to manage long chains of state management, or otherwise long contexts, in a conversational setting, to better refine and manage user preferences [58, 62, 69, 72]. Where classic conversational systems may have been constrained to managing conversational state using a decision tree, and a user is presented with options to choose from, our findings suggest the cumbersome nature of this may be detrimental to engagement. Meanwhile, advancements in language-based technologies such as large language models suggest potential benefits in tracking changes throughout a conversation and may offer salient improvements to systems like ours. Future work might, for example, leverage such a model to parse text and maintain a structured state where the model is prompted in how to interpret increasingly detailed preferences.

Finally, we note that our SMS-based approach, the specific language we use, and frequency of messages may influence preference refinement in ways that our analysis does not capture.

Further work is needed to understand the effect that varying these factors may have in changing preferences. For example, repeated messages and stronger language may lead to more homes expressing preferences, leading to fewer underconstrained homes. Similarly, informed by our study, it may be possible to shift stated preferences if we ask homes whether such preferences are proxies for other characteristics, such as ambulation or aggression. However, as discussed above, these other characteristics may be more difficult to express succinctly over SMS.

6.3 Implications for Matching Systems in High-Stakes Contexts

Our work sheds light on the challenges faced in placing patients, which may also be useful in other marketplace settings concerned with preference management. For example, trust and privacy also play central roles in communication channels in caregiving or children's foster care. Here we highlight how similar high-stakes contexts may benefit from improved communication channels and preference aggregation mechanisms. We note that (a) there may be similarities in metrics with other high-stakes contexts, (b) our work suggests benefits in enabling reciprocal exploration of supply, for example in our case: hospitals are enabled exploration of potential home matches, and homes benefit from exploring potential patient matches and (c) the value in aiming to support stakeholders rather than aiming to fully automate marketplace activity.

Similar to children's foster care, stakeholders in our work have incentives to reduce staff burnout, reduce placement recidivism, and make visitation and assessment logistics as efficient as possible due to resource constraints and potential negative emotional consequences [16, 42]. Metrics may also be similar; for example, length of stay at a hospital is analogous to duration in foster care before adoption. Furthermore, although our setting deals principally with older adults, there may be similarities in the changing characteristics of children early in life, akin to changing care needs at the end of life. As such, there may be similar load-balancing, as discussed in Section 5.2.2, that adoptive parents and home operators alike take into account in preference building. We also note that it is possible for similar systems to ours—that continuously collect and aggregate family preferences, instead of treating them as static and immutable—to help improve placement outcomes.

Prior work [16] has compared caseworker-driven matching and family-driven matching systems. They critique family-driven matching as tending to result in many families expressing an outsized interest in a few children, and advocate for caseworker-driven matching that takes into account families' preferences and children's specific needs. In our setting, home operators are analogous to families, and many homes might be interested in only a few highly desirable patients. While hospital staff may act as third-party matchmakers in our setting, our work suggests it is also important to facilitate exploration by home operators—who ultimately decide whether to accept a patient—and enable them to change their preferences over time in response to real-world constraints. Further, we see benefits in both tracking homes' changing preferences and sending homes candidate placements.

Finally, we note that both our work and work in child welfare point to the need for ultimate face-to-face assessment. Our findings show that participants want to conduct their own assessments, and the high number of underconstrained homes both before and after our experiment further illustrates this demand. These findings emphasize that systems in sensitive contexts like foster care should augment the immediate daily needs and workflows of workers carrying out placement efforts. Thus, rather than automating a marketplace, it is important for systems like ours to instead surface valuable up-to-date data that supports human decision-making by marketplace participants.

6.4 Limitations

Our study has several limitations, one of which is the use of binary categories, male or female, to classify patients by sex. The heteronormative assumptions underlying this design do not adequately account for the variety of non-binary gender identities. Future research should provide more

inclusive approaches that better reflect the complexity of gender identity and better support care decisions. One benefit of our approach to communicate with homes through SMS is that homes can, and do, communicate more complex preferences, as we discuss. One challenge for future work is in better understanding how to encode such preferences in our system to inform patient matches.

Lastly, our study took place in a single state and our results may not directly generalize to contexts with differing economic or cultural dynamics. Although our work was conducted with the largest hospital system in the state and influenced by federal economics systems (e.g., long term care insurance), the specific influences detailed in our findings relating to preference management, for example as driven by cost of living and cultural assumptions, may differ greatly in other regions.

Acknowledgments

This work would not have been possible without the extensive support from Queen's Medical Center. We also thank funding organizations supporting our researchers: VB is supported by the Gates Millennium Scholar Program, NG was supported by NSF CAREER IIS-2339427, and Cornell Tech Urban Tech Hub, Google, and Amazon research awards.

References

- [1] AARP Public Policy Institute. 2015. Caregiving for 75+ Population: Caregiver Profile. https://www.aarp.org/content/dam/aarp/ppi/2015/AARP1001_75orOlder_CGProfileAug26.pdf Accessed: October 23, 2024.
- [2] Christine Adamus, Dirk Richter, Kim Sutor, Simeon Joel Zürcher, and Sonja Mötteli. 2024. Preference for Competitive Employment in People with Mental Disorders: A Systematic Review and Meta-analysis of Proportions. *Journal of occupational rehabilitation* (2024). <https://api.semanticscholar.org/CorpusID:269360571>
- [3] Marc Afilalo, Xiaoqing Xue, Nathalie Soucy, Antoinette Colacone, Emmanuelle Jourdenais, and Jean-François Boivin. 2017. Patient Needs, Required Level of Care, and Reasons Delaying Hospital Discharge for Nonacute Patients Occupying Acute Hospital Beds. *Journal for Healthcare Quality* 39 (2017), 200–210. <https://api.semanticscholar.org/CorpusID:4721392>
- [4] Narges Ahani, Tommy Andersson, Alessandro Martinello, Alexander Teytelboym, and Andrew C. Trapp. 2021. Placement Optimization in Refugee Resettlement. *Oper. Res.* 69 (2021), 1468–1486. <https://api.semanticscholar.org/CorpusID:53067139>
- [5] Itai Ashlagi, Mark Braverman, Yashodhan Kanoria, and Peng Shi. 2017. Communication Requirements and Informative Signaling in Matching Markets. *Proceedings of the 2017 ACM Conference on Economics and Computation* (2017). <https://api.semanticscholar.org/CorpusID:51685685>
- [6] Victoria Bellotti, Daniel Turner, Kamila Demkova, Alexander Ambard, and Amanda Waterman. 2017. Why Users Disintermediate Peer-to-Peer Marketplaces. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017). <https://api.semanticscholar.org/CorpusID:26691809>
- [7] Kathryn H. Bowles et al. 2019. Using a Decision Support Algorithm for Referrals to Post-Acute Care. *Journal of the American Medical Directors Association* 20, 4 (2019), 408–413.
- [8] V Braun and V Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* (2006).
- [9] Eric Budish and Judd B. Kessler. 2016. Can Market Participants Report Their Preferences Accurately (Enough)? *ERN: Bargaining Theory (Topic)* (2016). <https://api.semanticscholar.org/CorpusID:221678633>
- [10] Mark R. Carey, Heena S. Sheth, and Dr. R. Scott Braithwaite. 2005. A prospective study of reasons for prolonged hospitalizations on a general medicine teaching service. *Journal of General Internal Medicine* 20 (2005), 108–115. <https://api.semanticscholar.org/CorpusID:32627896>
- [11] Francisco Castro, Peter Frazier, Hongyao Ma, Hamid Nazerzadeh, and Chiwei Yan. 2020. Matching queues, flexibility and incentives. *arXiv preprint arXiv:2006.08863* (2020).
- [12] Sarah H Cen, Andrew Ilyas, Jennifer Allen, Hannah Li, and Aleksander Madry. 2024. Measuring strategization in recommendation: Users adapt their behavior to shape future content. *arXiv preprint arXiv:2405.05596* (2024).
- [13] B. Chan et al. 2019. A Decision Support Algorithm for Referrals to Post-Acute Care. *Journal of the American Medical Directors Association* 20, 4 (2019), 414–421.e3. doi:10.1016/j.jamda.2019.01.001
- [14] Yi-Ting Chuang et al. 2023. From Prediction to Decision: Optimizing Long-Term Care Placements among Older Delayed-Discharge Patients. *Production and Operations Management* (2023). In press.
- [15] M. J. Constantino et al. 2021. Effect of Matching Therapists to Patients vs Assignment as Usual on Adult Psychotherapy Outcomes: A Randomized Clinical Trial. *JAMA Psychiatry* 78, 9 (2021), 960–969. doi:10.1001/jamapsychiatry.2021.1234

- [16] Ludwig Dierks, Nils Olberg, Sven Seuken, Vincent W. Slaugh, and M. Utku Ünver. 2021. Search and Matching for Adoption from Foster Care. *ArXiv abs/2103.10145* (2021). <https://api.semanticscholar.org/CorpusID:232269865>
- [17] Ludwig Dierks, Vincent Slaugh, and M. Utku Ünver. 2024. Child Welfare Platform Design to Improve Outcomes for Children with Disabilities. *SSRN Electronic Journal* (2024). <https://api.semanticscholar.org/CorpusID:268870375>
- [18] Lauren Doctoroff, Douglas J. Hsu, and Kenneth J. Mukamal. 2017. Trends in Prolonged Hospitalizations in the United States from 2001 to 2012: A Longitudinal Cohort Study. *The American journal of medicine* 130 4 (2017), 483.e1–483.e7. <https://api.semanticscholar.org/CorpusID:4523617>
- [19] Family Caregiver Alliance. 2021. Caregiver Statistics: Demographics. <https://www.caregiver.org/resource/caregiver-statistics-demographics/> Accessed: October 23, 2024.
- [20] Apostolos Filippas, Srikanth Jagabathula, and Arun Sundararajan. 2023. The limits of centralized pricing in online marketplaces and the value of user control. *Management Science* 69, 12 (2023), 7202–7216.
- [21] Mingkun Gao, Hyo Jin Do, and Wai-Tat Fu. 2018. Burst Your Bubble! An Intelligent System for Improving Awareness of Diverse Social Opinions. *Proceedings of the 23rd International Conference on Intelligent User Interfaces* (2018). <https://api.semanticscholar.org/CorpusID:3803773>
- [22] Nikhil Garg and Hamid Nazerzadeh. 2022. Driver surge pricing. *Management Science* 68, 5 (2022), 3219–3235.
- [23] Dawn G. Gregg and Steven Walczak. 2006. Auction Advisor: an agent-based online-auction decision support system. *Decision Support Systems* 41, 2 (2006), 449–471. doi:10.1016/j.dss.2004.07.007
- [24] Jungpil Hahn. 2001. The dynamics of mass online marketplaces: a case study of an online auction. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2001). <https://api.semanticscholar.org/CorpusID:91848>
- [25] Stephen M. Haley, Wendy Jane Coster, Patricia L. Andres, Larry Ludlow, Peng Sheng Ni, Tamara L. Y. Bond, Samuel Justin Sinclair, and Alan M. Jette. 2004. Activity Outcome Measurement for Postacute Care. *Medical Care* 42 (2004), 1–49. <https://api.semanticscholar.org/CorpusID:2214729>
- [26] Gerald Häubl and Kyle B. Murray. 2001. Preference Construction and Persistence in Digital Marketplaces: The Role of Electronic Recommendation Agents. <https://api.semanticscholar.org/CorpusID:995037>
- [27] Gerald Häubl and Kyle B. Murray. 2001. Recommending or Persuading? The Impact of a Shopping Agent’s Algorithm on User Behavior. In *Proceedings of the 3rd ACM Conference on Electronic Commerce* (Tampa, Florida, USA) (EC ’01). Association for Computing Machinery, New York, NY, USA, 163–170. doi:10.1145/501158.501176
- [28] Andreas Haupt, Dylan Hadfield-Menell, and Chara Podimata. 2023. Recommending to strategic users. *arXiv preprint arXiv:2302.06559* (2023).
- [29] Health Services Advisory Group. 2018. Nursing Facility Level of Care Presentation. <https://www.hsag.com/contentassets/be6241c592274b8b9ed776e95492e3c4/nf-loc---presentation.pdf> Accessed: October 23, 2024.
- [30] David Holtz, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. 2020. The engagement-diversity connection: Evidence from a field experiment on spotify. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 75–76.
- [31] Jana S. Hopstaken, Lynn Verweij, C. J. H. M. van Laarhoven, Nicole M. A. Blijlevens, Martijn W. J. Stommel, and Rosella Hermens. 2021. Effect of Digital Care Platforms on Quality of Care for Oncological Patients and Barriers and Facilitators for Their Implementation: Systematic Review. *Journal of Medical Internet Research* 23 (2021). <https://api.semanticscholar.org/CorpusID:237615031>
- [32] Yufeng Huang. 2022. Pricing frictions and platform remedies: the case of Airbnb. *Available at SSRN 3767103* (2022).
- [33] Diane L. Huber and Eleanor McClelland. 2003. Patient preferences and discharge planning transitions. *Journal of professional nursing : official journal of the American Association of Colleges of Nursing* 19 4 (2003), 204–10. <https://api.semanticscholar.org/CorpusID:5622101>
- [34] Halah Ibrahim, Thana Harhara, Syed Athar, Satish Chandrasekhar Nair, and Ahsraf M Kamour. 2022. Multi-Disciplinary Discharge Coordination Team to Overcome Discharge Barriers and Address the Risk of Delayed Discharges. *Risk Management and Healthcare Policy* 15 (2022), 141 – 149. <https://api.semanticscholar.org/CorpusID:246536978>
- [35] Tapio Ikkala and Airi Lampinen. 2015. Monetizing Network Hospitality: Hospitality and Sociability in the Context of Airbnb. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (2015). <https://api.semanticscholar.org/CorpusID:207220249>
- [36] Guido Imbens and Karthik Kalyanaraman. 2012. Optimal Bandwidth Choice for the Regression Discontinuity Estimator. *The Review of Economic Studies* 79, 3 (2012), 933–959. <https://EconPapers.repec.org/RePEc:oup:restud:v:79:y:2012:i:3:p:933-959>
- [37] Maya A Kaneko, Caitlin Lustig, Daniela Rosner, and Audrey Desjardins. 2024. Care Layering: Complicating Design Patterns. *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (2024). <https://api.semanticscholar.org/CorpusID:270822678>
- [38] Airi Lampinen. 2014. Account sharing in the context of networked hospitality exchange. *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (2014). <https://api.semanticscholar.org/CorpusID:20255816>

- [39] Airi Lampinen, Victoria Bellotti, A. Monroy-Hernández, Coye Cheshire, and Alexandra Samuel. 2015. Studying the "Sharing Economy": Perspectives to Peer-to-Peer Exchange. *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing* (2015). <https://api.semanticscholar.org/CorpusID:20882562>
- [40] Airi Lampinen and Barry A. T. Brown. 2017. Market Design for HCI: Successes and Failures of Peer-to-Peer Exchange Platforms. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (2017). <https://api.semanticscholar.org/CorpusID:39769664>
- [41] Airi Lampinen, Vilma Lehtinen, Coye Cheshire, and Emmi Suhonen. 2013. Indebtedness and reciprocity in local online exchange. *Proceedings of the 2013 conference on Computer supported cooperative work* (2013). <https://api.semanticscholar.org/CorpusID:9378789>
- [42] Lindsay Lanham. 2022. Predictors of Adoption Disruption and Dissolution: A Literature Review. <https://adoptioncouncil.org/publications/predictors-of-adoption-disruption-and-dissolution-a-literature-review/>
- [43] Angela Y. Lee, Hannah Mieczkowski, Nicole B. Ellison, and Jeffrey T. Hancock. 2022. The Algorithmic Crystal: Conceptualizing the Self through Algorithmic Personalization on TikTok. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 543 (Nov. 2022), 22 pages. doi:10.1145/3555601
- [44] Lin Li, Vitica Arnold, and Anne Marie Piper. 2023. "Any bit of help, helps": Understanding how older caregivers use carework platforms for caregiving support. *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (2023). <https://api.semanticscholar.org/CorpusID:258217278>
- [45] J. Liu and L. Jiang. 2022. Stable Two-Sided Satisfied Matching for Hospitals and Patients Based on the Disappointment Theory. *Frontiers in Public Health* 10 (2022), 9470515. doi:10.3389/fpubh.2022.9470515
- [46] Ana Llena-Nozal. 2020. Who cares?: attracting and retaining care workers for the elderly. <https://api.semanticscholar.org/CorpusID:226531985>
- [47] Caitlin Lustig, Maya A Kaneko, Meghna Gupta, Kavita Dattani, Audrey Desjardins, and Daniela Rosner. 2024. Porous by Design: How Childcare Platforms Impact Worker Personhood, Safety, and Connection. *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (2024). <https://api.semanticscholar.org/CorpusID:270822852>
- [48] Evan Magnusson. 2024. rdd: Regression Discontinuity Design Package. <https://github.com/evan-magnusson/rdd/>. Accessed: 2024-10-22.
- [49] David B. Martin, Benjamin V. Hanrahan, Jacki O'Neill, and Neha Gupta. 2014. Being a turker. *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (2014). <https://api.semanticscholar.org/CorpusID:13511589>
- [50] Julia M. Mayer, Starr Roxanne Hiltz, Louise Barkhuus, Kaisa Väänänen, and Quentin Jones. 2016. Supporting Opportunities for Context-Aware Social Matching: An Experience Sampling Study. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (2016). <https://api.semanticscholar.org/CorpusID:5315656>
- [51] Med-QUEST Division, Hawaii State Department of Human Services. 2015. *Medicaid Provider Manual: Long Term Care, Chapter 12*. <https://medquest.hawaii.gov/content/medquest/en/archive/PDFs/Provider%20Manual/PMChp12.pdf> Revised edition, originally issued in November 2008.
- [52] Robert L. Mollica, Kristin Simms-Kastelein, Michael Cheek, Candace Baldwin, Jennifer Farnham, Susan Reinhard, and Jean Accius. 2009. Building adult foster care: What states can do. *AARP Public Policy Institute* (2009). <https://eadn-wc03-6094147.nxedge.io/cdn/wp-content/uploads/sites/default/files/Building%20Adult%20Foster%20Care.pdf> Accessed: 05/02/2023.
- [53] Ruchi Ookalkar, Kolli Vishal Reddy, and Eric Gilbert. 2019. Pop: Bursting News Filter Bubbles on Twitter Through Diverse Exposure. *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing* (2019). <https://api.semanticscholar.org/CorpusID:207960113>
- [54] Kenny Peng, Manish Raghavan, Emma Pierson, Jon Kleinberg, and Nikhil Garg. 2024. Reconciling the accuracy-diversity trade-off in recommendations. In *Proceedings of the ACM Web Conference 2024*. 1318–1329.
- [55] Adrian Petterson, Isabella Jaimes Rodriguez, Olivia Doggett, and Priyank Chandra. 2024. Networks of care in digital domestic labour economies. *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2024). <https://api.semanticscholar.org/CorpusID:269748402>
- [56] Joyce Rafla, Kate Schwartz, Hirokazu Yoshikawa, Dennis Hilgendorf, Anaga Ramachandran, Mohammad Khanji, Rawan Abu Seriah, Mohammad Al Aabed, Ragheb Fityan, Phoebe Sloane, Ayat Al Aqra, Razan Mousa, Tareq Sharawi, Andrés Molano, Kimberly Foulds, Jere Behrman, and Alice Wuermli. 2024. Cluster randomized controlled trial of a phone-based caregiver support and parenting program for Syrian and Jordanian families with young children. *Early Childhood Research Quarterly* 69 (2024), 141–153. doi:10.1016/j.ecresq.2024.07.004
- [57] Noopur Raval and Paul Dourish. 2016. Standing Out from the Crowd: Emotional Labor, Body Labor, and Temporal Labor in Ridesharing. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (2016). <https://api.semanticscholar.org/CorpusID:18810608>
- [58] Zhaochun Ren, Zhi Tian, Dongdong Li, Pengjie Ren, Liu Yang, Xin Xin, Huasheng Liang, M. de Rijke, and Zhumin Chen. 2022. Variational Reasoning about User Preferences for Conversational Recommendation. *Proceedings of*

- the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2022). <https://api.semanticscholar.org/CorpusID:248914154>
- [59] Paul Resnick, R. Kelly Garrett, Travis Kriplean, Sean A Munson, and Natalie Jomini Stroud. 2013. Bursting your (filter) bubble: strategies for promoting diverse exposure. In *Conference on Computer Supported Cooperative Work*. <https://api.semanticscholar.org/CorpusID:20865375>
- [60] Anna Marie Rezk, Auste Simkute, Ewa Luger, John Vines, Chris Elsdon, Michael Evans, and Rhianne Jones. 2024. Agency Aspirations: Understanding Users' Preferences And Perceptions Of Their Role In Personalised News Curation. *Proceedings of the CHI Conference on Human Factors in Computing Systems* (2024). <https://api.semanticscholar.org/CorpusID:269745451>
- [61] Alvin E. Roth. 2018. Marketplaces, Markets, and Market Design. *American Economic Review* 108, 7 (July 2018), 1609–58. doi:10.1257/aer.108.7.1609
- [62] Qi Shen, Lingfei Wu, Yiming Zhang, Yitong Pang, Zhihua Wei, Fangli Xu, Bo Long, and Jiangsen Pei. 2023. Multi-Interest Multi-Round Conversational Recommendation System with Fuzzy Feedback based User Simulator. *ACM Transactions on Recommender Systems* (2023). <https://api.semanticscholar.org/CorpusID:261077393>
- [63] Ping Shi et al. 2021. Timing It Right: Balancing Inpatient Congestion vs. Readmission Risk at Discharge. *Operations Research* 69, 5 (2021), 1500–1518.
- [64] Ellen Simpson, Andrew Hamann, and Bryan Semaan. 2022. How to Tame "Your" Algorithm: LGBTQ+ Users' Domestication of TikTok. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 22 (Jan. 2022), 27 pages. doi:10.1145/3492841
- [65] Woan Shin Tan, Wai Fung Chong, Karen Sui-Geok Chua, Bee Hoon Heng, and Kay Fei Chan. 2010. Factors associated with delayed discharges after inpatient stroke rehabilitation in Singapore. *Annals of the Academy of Medicine, Singapore* 39 6 (2010), 435–41. <https://api.semanticscholar.org/CorpusID:18095586>
- [66] Jacob Thebault-Spieker, Loren G. Terveen, and Brent J. Hecht. 2015. Avoiding the South Side and the Suburbs: The Geography of Mobile Crowdsourcing Markets. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (2015). <https://api.semanticscholar.org/CorpusID:257103>
- [67] Anonymous Article to Appear in CSCW 2025. 2025. Faster Information for Effective Long-Term Discharge: A Field Study in Adult Foster Care. <https://drive.google.com/file/d/1kRjua2nH3X1cZPNZVMq2Hk0306o4A1Hu/view?usp=sharing> Shared Document on Google Drive; Accessed: September 25, 2024.
- [68] Shresth Verma, G Singh, Aditya Mate, Paritosh Verma, Sruthi Gorantla, Neha, Madhiwalla, Aparna Hegde, Divy Thakkar, Manish Jain, Milind Tambe, and Aparna Taneja. 2023. Increasing Impact of Mobile Health Programs: SAHELI for Maternal and Child Care. In *AAAI Conference on Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:253400315>
- [69] Kerui Xu, Jingxuan Yang, Jun Xu, Sheng Gao, Jun Guo, and Ji rong Wen. 2021. Adapting User Preference to Online Feedback in Multi-round Conversational Recommendation. *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (2021). <https://api.semanticscholar.org/CorpusID:232126028>
- [70] Emma J. Zhao, Apurva Yeluru, Lakshman Manjunath, Lei Ray Zhong, Hsiao-Tieh Hsu, Charles K Lee, Anny C Wong, Matthew Abramian, Haley Manella, David Svec, and Lisa Shieh. 2018. A long wait: barriers to discharge for long length of stay patients. *Postgraduate Medical Journal* 94 (2018), 546 – 550. <https://api.semanticscholar.org/CorpusID:52947436>
- [71] M. Zhao, Y. Wang, X. Zhang, and C. Xu. 2023. Online Doctor–Patient Dynamic Stable Matching Model Based on Regret Theory under Incomplete Information. *Socio-Economic Planning Sciences* 87 (2023), 101615. doi:10.1016/j.seps.2023.101615
- [72] Jinhang Zuo, Songwen Hu, Tong Yu, Shuai Li, Handong Zhao, and Carlee Joe-Wong. 2022. Hierarchical Conversational Preference Elicitation with Bandit Feedback. *Proceedings of the 31st ACM International Conference on Information & Knowledge Management* (2022). <https://api.semanticscholar.org/CorpusID:252211868>

A Appendix

A.1 The relationship between ambulation and hospital stay duration

Here, we extend the analysis from Section 3.3, establishing the relationship between the patient's ambulation score and their duration in the hospital past first being ready to be discharged.

A.1.1 Initial analysis. Table 4 shows the regression table corresponding to the analysis in Figure 1. Higher ambulation (mobility) scores correspond to lower lengths of stay, suggesting that homes may be incentivized to take more mobile patients more readily.

A.1.2 Regression discontinuity to understand the effect of insurance brackets. Next, we extend our regression analysis to better understand economic incentives driving preference strategy in our

Table 4. Regression table for ambulation versus post-acute hospital stay duration

Dep. Variable:	Enc Duration	R-squared:	0.009			
Model:	OLS	Adj. R-squared:	0.007			
Method:	Least Squares	F-statistic:	6.679			
		Prob (F-statistic):	0.00994			
		Log-Likelihood:	-4572.0			
No. Observations:	760	AIC:	9148.			
Df Residuals:	758	BIC:	9157.			
Df Model:	1					
Covariance Type:	nonrobust					
	coef	std err	t	P > t 	[0.025	0.975]
const	83.9337	8.931	9.398	0.000	66.402	101.465
AM-PAC Mobility Score	-1.5188	0.588	-2.584	0.010	-2.672	-0.365
Omnibus:	803.345	Durbin-Watson:	1.792			
Prob(Omnibus):	0.000	Jarque-Bera (JB):	41706.699			
Skew:	5.025	Prob(JB):	0.00			
Kurtosis:	37.872	Cond. No.	37.8			

Notes:

[1] Standard Errors assume that the covariance matrix of the errors is correctly specified.

patient placement setting. In particular, we are motivated by the observation in Figure 1 that there is a high variance of outcomes for patients near an insurance bracket change; intuitively, one may expect that patients just below an insurance bracket change are easier to place than patients just above the insurance bracket change—care home operators are reimbursed more for the patients just below the bracket cutoff than just above, even though the patients have similar mobility scores. Such discrete cutoffs may distort care home preferences and the length of stay for patients.

To begin to quantify this effect, we run a *regression discontinuity* to analyze whether patients just above the insurance bracket cutoff indeed are harder to place. Figure 6 illustrates a discontinuity between the two regressions between 15 and 16, where the first ends at 47.0 and the second begins at 60.9. For a more formal analysis, we leverage the regression discontinuity package [48]. The discontinuity estimate generates an optimal bandwidth within which to measure the discontinuity, according to Imbens and Kalyanaraman [36]. Results are in Table 5. The analysis shows a statistically significant difference ($p = 0.011$) for the cutoff at ambulation level 15, the cutoff between level one and level two patients. In other words, there is a statistically significant increase in the length of stay at the hospital, once the patient crosses the ambulation threshold such that the insurance payments are lower. We note that further work is needed to better corroborate the robustness of these findings and understand potential confounders.

A.2 Additional details regarding SMS labeling

During our experiment our system carried out vacancy collection as usual: we asked homes if they have a vacancy or not, and saved this information on our system. Much of this process our system conducts automatic labeling for: if a home's entire response is: 'confirm' or '0', for example, then no status updates are saved, or their vacancies are set to 0. Similarly for preferences, if a home states "male only", their sex preference is overridden accordingly. For more complicated free-form messages, such as "No space but accepting male," messages are tagged for human verification. For

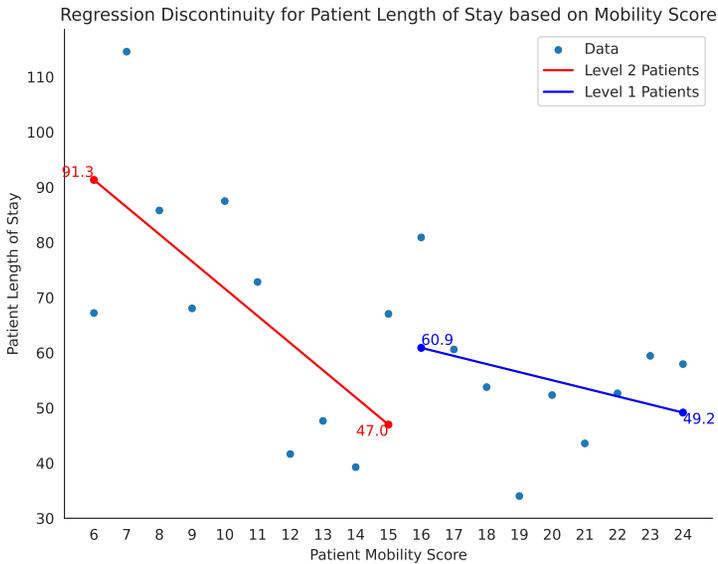


Fig. 6. A scatter plot with two regressions, with patient mobility score on the x-axis ranging from 6 to 24, and hospital length of stay on the y-axis from 30 to 115. Each scatter point represents the average length of stay for a patient with the associated mobility score. The first regression is in red from 6 to 15 and represents level one patients, the second regression is in blue from 16 to 24 and represents level two patients.

the purpose of our study, each message received throughout the study was further labeled as one of: "no preference change", "expressed patient interest", "changed weight preference", and/or "changed sex preference", all of which were human-verified.

If a home stated they are interested in assessing a patient, for example by responding with "Yes" to our prompts, or if they ask a clarifying question to gain more information about the patient, for example, "is the patient paralyzed?", they are labeled with "expressed patient interest." If a home stated a preference that was different from the one we had saved and sent them, the message is labeled as a preference change. In other words, restating the same preference we already saved did not count as a preference change. A message could be tagged with multiple labels, for example, indicating both a sex preference and weight preference change.

A.3 Message Templates Used

A.3.1 Control.

Happy Monday! We are messaging to confirm your status: [STATUS, eg:] 0 medicaid, 0 private. No weight limit. Male OK, Female OK. Accepting: HMSA, UHC, OHANA, ALOHA, KAISER. If this is still accurate, please 'confirm', or please respond with changes, so we can find the best fit if a vacancy arises. Thank you.

A.3.2 Treatment: Underconstrained.

Good Morning! We are messaging to confirm your status. 0 medicaid, 0 private 145 lbs limit. Male OK, Not Female. Please 'confirm' or update, e.g. '1 medicaid'.

Table 5. Regression discontinuity analysis based on the work of Imbens and Kalyanaram [36].

Dep. Variable:	los	R-squared:	0.642			
Model:	WLS	Adj. R-squared:	0.552			
Method:	Least Squares	F-statistic:	7.173			
		Prob (F-statistic):	0.0164			
		Log-Likelihood:	-40.901			
No. Observations:	11	AIC:	87.80			
Df Residuals:	8	BIC:	88.99			
Df Model:	2					
Covariance Type:	nonrobust					
	coef	std err	t	P> t 	[0.025	0.975]
Intercept	32.4497	8.487	3.824	0.005	12.879	52.020
TREATED	46.7603	14.154	3.304	0.011	14.120	79.400
ampac	-8.4408	2.229	-3.787	0.005	-13.580	-3.301
Omnibus:	0.824	Durbin-Watson:	2.159			
Prob(Omnibus):	0.662	Jarque-Bera (JB):	0.605			
Skew:	-0.043	Prob(JB):	0.739			
Kurtosis:	1.854	Cond. No.	14.8			

Notes:

[1] Standard Errors assume that the covariance matrix of the errors is correctly specified.

Further, if you were to have a vacancy, your preferences currently match this patient. ‘[Patient Age, Weight, Insurance]. [Patient ambulation descriptors: Ambulation is immobile. Bed transfers are maximum. Bathing is maximum. Dressing is maximum.]. [Patient cognition descriptors, eg: does not wander, has difficulty communicating, has short term memory loss, has long term memory loss.]. [Further patient notes, e.g.: Also requires support for: hospice.] Please respond ‘Yes’ if you would be interested in assessing such a patient if you were to have a vacancy at some point. Otherwise, please respond with ‘No’, and which characteristic does not work and we will work to find a better match. For example, ‘No, cannot take insurance that is ...’ so we can find the best fit when a vacancy arises. Mahalo.

A.3.3 Treatment: Overconstrained-Sex.

"In the last week you would have been eligible to match with 2x more patients if you accepted [EXCLUDED SEX] patients in addition to [EXISTING SEX] patients. Are you interested in assessing these additional patients? If not, please let us know, thank you."

Here the rationale is that by simply accepting patients of any sex, a home who may at least double their preferences, will potentially be open to changing their preference and assessing these additional patients.

A.3.4 Treatment: Overconstrained-Weight.

"In the last week you would have been eligible to match with 2x more patients if you accepted patients instead of at most [CURRENT WEIGHT LIMIT] pounds, at most [CURRENT WEIGHT LIMIT+10] pounds. Are you interested in assessing these additional patients? If not, please let us know, thank you."

Here the rationale is that by simply accepting patients one step of 10 pounds above their current expressed limit, knowing they would double their current match count, they would then potential relax these preference and assess the additional patients.

A.3.5 Treatment: Overconstrained-Ambulation.

"In the last week you would have been eligible to match with 2x more patients if you accepted patients instead of at most [CURRENT AMBULATION LIMIT], any degree of ambulation. Are you interested in assessing these additional patients? If not, please let us know, thank you."

Here the rationale is that homes who have expressed an ambulation limit simply state they only want ambulatory patients, so if by knowing that also considering non-ambulatory patients they would at least double their match count, they might assess these additional patients.

A.4 Additional details regarding security practices implemented in software stack

Our work is integrated within a hospital system, and as such numerous security measures are in place in alignment with hospital security practices, US HIPAA regulations, as well as standardized System and Organization Controls (SOC) required of our software. Among these are the use of industry best-practices in authentication for web application users, ranging from Google OAuth2 with multifactor authentication, to annual penetration tests by third party auditors, live vulnerability scanning and encryption in transit and at rest. For the purposes of this study, SMS communication was informed strictly by de-identified data which ultimately act as descriptions of patients.

Received October 2024; revised April 2025; accepted August 2025