Faster Information for Effective Long-Term Discharge: A Field Study in Adult Foster Care

VINCE BARTLE, Department of Information Science, Cornell Tech, USA
ASHLEY SHEARER, Care Coordination, The Queen's Health System, USA
ALEXANDRA WROE, Operations, The Queen's Health System, USA
NICOLA DELL, Department of Information Science, the Jacobs Institute, Cornell Tech, USA
NIKHIL GARG, Department of Operations Research, the Jacobs Institute, Cornell Tech, USA

As the US population ages, a growing challenge is placing hospital patients who require long-term post-acute care into adult foster care facilities: small long-term nursing facilities that care for those unable to age in place because their care requirements exceed what can be delivered at home. A key challenge in patient placement is the dynamic matching process between hospital discharge coordinators looking to place patients and facilities looking for residents. We designed, built, deployed, and maintain a system to support decision making among a team of 6 discharge coordinators assisting in the discharge of 127 patients across 1,047 facilities in Hawai'i. Our system collects vacancy and capability data from facilities via conversational SMS and processes it to recommend facilities that discharge coordinators might contact. Findings from a 14 month deployment provide evidence for how timely, accurate information positively impacts matching efficacy. We close with lessons learned for information collection systems and provisioning platforms in similar contexts.

CCS Concepts: • **Information systems** → **Collaborative and social computing systems and tools**; • **Applied computing** → *Operations research*; *Health care information systems.*

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, Ashley Shearer, Alexandra Wroe, Nicola Dell, and Nikhil Garg. 2024. Faster Information for Effective Long-Term Discharge: A Field Study in Adult Foster Care. *Proc. ACM Hum.-Comput. Interact.* 9, CSCW2, Article 85 (April 2024), 29 pages. https://doi.org/10.1145/3710983

1 Introduction

As the US population ages, a growing body of CSCW and HCI research explores how technology might support the provision of care for older adults. Much of this research focuses on enabling older adults to remain at home or "age in place" [18, 25, 43], often with the help of formal [50, 60] or informal caregivers [8, 46]. However, 19% of adults over 50 believe their home is insufficient to facilitate aging in place [67], and over 1.5 million US adults over 65 currently live in nursing, residential, or adult foster care homes [16, 30, 74].

One prevalent avenue to no longer being able to remain at home is after a hospital stay: older adult patients often cannot be discharged back home even if they no longer need acute hospital

Authors' Contact Information: Vince Bartle, Department of Information Science, Cornell Tech, New York, NY, USA; Ashley Shearer, Care Coordination, The Queen's Health System, Honolulu, HI, USA; Alexandra Wroe, Operations, The Queen's Health System, Honolulu, HI, USA; Nicola Dell, Department of Information Science, the Jacobs Institute, Cornell Tech, New York, NY, USA; Nikhil Garg, Department of Operations Research, the Jacobs Institute, Cornell Tech, New York, NY, USA.

ACM 2573-0142/2024/4-ART85
https://doi.org/10.1145/3710983

care, and must instead be *placed* in a long-term care home. Such placement is a major challenge: in November 2022, 35 medical organizations, including the American Medical Association, wrote to the Biden administration detailing "gridlock" in Emergency Departments, and a "public health emergency" in part due to the inability to discharge patients [56]. Such gridlock leads to increased length of patient stays in hospital and substantial extra expense: an occupied bed can cost thousands of dollars per day, which is often paid by the hospital [56].

This paper contributes findings from a 14 month deployment of an intervention to improve the process of discharging patients in need of long-term care, from *hospitals* to *adult foster care (AFC)*. We focus on patients whose hospitalization has lasted longer than 21 days; prior work has shown that such long-term stays are important; while they are only 2% of hospitalizations, they account for roughly 15% of overall hospital bed-days [28].

At hospitals, discharge is managed by *Discharge Coordinators* who work to find an appropriate home to discharge patients to. Coordinators often have limited control over a patient's care, including insurance coverage and care home availability; thus, a large part of their work is information discovery and communication. They may need to call hundreds of homes before finding one that can accept a given patient. Coordinators also seek to identify placements that are less likely to lead to readmission, which may increase costs and worsen patient outcomes—motivating an inclusion of social and cultural factors in the placement process. These factors combine to make coordinators' job important and challenging that, although fulfilling when a successful placement happens, often leads to high levels of stress and burnout.

On the other side are *Adult Foster Care Homes*, who must find suitable residents. In contrast to better-known nursing homes, AFC homes are smaller scale (often five or fewer residents), less institutionalized homes that are often run by, for example, a retired nurse who is also an older adult [29, 51]. AFC homes are often more local to where their patients used to live—and thus more integrated within the community. However, finding suitable residents can be challenging. Homes have specific capabilities and preferences, for example with respect to the medical severity of a patient. Empty beds or ill-matched patients may have disproportionate impacts on small homes and their patients, compared to larger nursing homes.

In our setting, coordinators historically rely on State Department of Health data to decide which homes to contact about placement. This data principally serves to convey that a home is licensed to operate, to prevent placements into unlicensed homes, but also includes whether a home has a vacancy. As we'll show, this data is insufficient for effective matching: we find it is updated on average every 105 days, but actual vacancies change much faster. We show how the task of calling homes about patient placement is challenging, which can in-part be alleviated by systems designed to align with existing workflows, that enable faster information update intervals.

We designed and deployed a conversational SMS system in Hawai'i, in partnership with a local hospital, to exchange patient availability and home vacancy data. As a working system in a high-stakes domain, our system needs to continuously engage with stakeholders, balance their constraints and priorities, and fit into existing workflows. We worked closely with coordinators and care homes for over two years to understand the daily workflows and challenges associated with placement, and over multiple iterations of system design, to create tools that support their work. This paper describes a portion of our work which sends personalized surveys periodically to care homes via SMS. We use human-in-the-loop labeling to compile these messages into status updates for homes, which informs ranked call recommendations provided to a team of 6 coordinators, who use this list to decide which homes to call.

Findings from a 14 month study in Hawai'i provide evidence of the impact that timely, accurate information has on increasing matches between patients and homes. By reducing informational and coordination frictions, our system improves patient placement workflows. We show substantial engagement with the system throughout the population, both for individual surveys and over the duration of our study, indicating the potential for such systems to effectively collect and disseminate information. Homes frequently updated their information in our system, demonstrating the need for such systems and the desire from homes to participate.

We conclude by discussing challenges and opportunities associated with the long term viability of such a system. We highlight the importance of centering improvement in everyday workflows, rather than measuring success via specific end metrics. We also discuss how our system constitutes a coordination mechanism that aids coordination between a team of hospital discharge coordinators and a large, distributed network of independently-owned and operated care homes. Finally, we discuss questions of sustainability and potential scaling across contexts.

Author Positionality. We acknowledge that researchers' backgrounds and experiences shape their perspectives and that our research may include our personal biases [72]. Our team consists of interdisciplinary collaborators with expertise in HCI, Operations Research, and healthcare, and years of experience in long-term care placement in Hawai'i. Our work embraces a reflective approach necessitated by the sensitive nature of foster care placement of vulnerable patients.

2 Related Work

In this section, we review relevant literature on (1) patient placement in adult foster care, (2) coordination work in healthcare, and (3) algorithmic decision support for vulnerable populations.

2.1 Adult Foster Care Placement

Adult foster care (AFC) is an important, but critically understudied component of the healthcare infrastructure in the United States. State-funded AFC started in the late 1970's in Oregon and Washington, and since then all but four states have adopted state-supported AFC [31]. Research has examined the challenges faced in AFC [38]. For example, Mollica et al. [51] provide guiding frameworks for state implementation and state-specific idiosyncrasies. Researchers have also studied resident perceptions around placement in AFC [65], the differences between nursing homes and AFC homes [66], and the need for different policy provisions between larger assisted-living facilities and smaller AFC homes [53]. To our knowledge, our work is the first to address the placement challenges faced by home operators and discharge coordinators or to discuss potential interventions (let alone computational interventions) aiding placement.

However, there is literature on effective *children's* foster care placement. Zeijlmans et al. [81] detail, for example, major gaps in the "Decision-Making Ecology" of foster care placement, noting a clear need for effective matching into foster care. Other work has examined the risks of placement instability and the implementation of a web-based matching system to inform placement decisions [52]. Saxena et al. [68] review the use of algorithms in the US Child Welfare System, and emphasize the importance of utilizing salient social worker knowledge and employing context-relevant social work literature. Researchers have also examined the role of trust in existing apprehensions of algorithmic decision-making in child welfare services [14]. Beyond foster care placement, Sun et al. [77] detail the reality that data-driven decision making in long-term care delivery and management is ultimately contingent on a growing burden of work on frontline care workers, for example to aggregate sensor data and digitize field notes about a patient's level of care.

Our work centers concerns of how data-driven and algorithmic systems affect the work and burdens placed on stakeholders. To this literature, we contribute insights into how out-of-date data leads to a phenomenon we call "runaway labor"—the amount of work a discharge coordinator needs to do to place a patient grows sharply with the data's age. We further contribute the design and deployment of a system to collect timely data from care homes in ways that aid coordination via integration with stakeholders' existing workflows.

2.2 Coordination Mechanisms in Healthcare

CSCW has long sought to understand how technology might support coordination and reduce friction in complex workflows [2, 34, 45]. Much of this research explores coordination in healthcare settings, including among formal and informal caregivers (e.g., family members) as older adults seek to age in place [10, 12, 78], within community health programs in the Global South [26, 27, 36, 79], and among formal care teams in hospital settings [1, 71, 75].

In our context, our placement data provision system can be viewed as an example of a *coordination mechanism*. Coordination mechanisms in cooperative work have been formalized since 1996 [70], and have been especially studied in hospital settings. For example, prior work studied the use of whiteboards and digital wall displays shared between internal hospital teams to call attention to important matters across teams of varying purpose [9]. Further work has sought to understand non-physical coordination mechanisms, for example in supporting emergency communication teams in hospital settings with a variety of computer systems and information panels [83]. With the rapid standardization of electronic health records and the growing adoption of clinical decision support, more recent work has studied for example the design of hospital alerts as coordination mechanisms to support decision making [49].

However, much of this work focuses on coordination among workers in a single organizational context (e.g., a hospital) and among internal teams. By contrast, studies that explore how to achieve coordination across multiple *different* organizational contexts are relatively rare. A notable example is the Care Hotel [11] which proposed ways to support actors from various organizational settings as they provide rehabilitative services to older adults transitioning from a hospital to their own home. The authors emphasize how successful coordination requires a shared communication platform that is flexible enough to support standardized and ad-hoc communication and coordination [11].

Our paper contributes an example of how technology can support coordination between an external network of over a thousand independently operated community care providers coordinating with a team of hospital discharge coordinators. Our system uses a communication channel (SMS) as a coordination mechanism, where natural language is used to enable data updates, which ensures sufficient flexibility to facilitate coordination with a large, heterogeneous set of AFC homes.

2.3 Algorithmic Decision Support for Vulnerable Populations

The challenge tackled by our system can be understood as a *matching* one: we wish to efficiently connect patients to compatible care homes that currently have vacancies. A long line of work in dynamic matching—for organ exchange [5], housing [80], ridesharing [23], etc.—develops algorithmic matching approaches that take in data and provide an answer. One particularly relevant example is refugee resettlement, wherein social workers use decision support tools to facilitate optimal placement of refugees across locations, to maximize objectives such as employment, burden sharing, and personal preference [3]. This literature describes a progression from a greedy batch matching process, that primarily seeks to burden share across locations, and eventually implements dynamic matching using predictions of placement success.

The CSCW community has similarly sought to understand the challenges faced in human-AI integration, ranging from supporting work in government [4, 33] to healthcare [17, 35]. In particular, work has highlighted a need to calibrate AI in service of mitigating over-reliance and loss of equity [19, 33], and a need to enable augmenting decision making rather than fully automating it [4]. Work to better understand supporting frontline health workers has especially emphasized the importance of successful integration with daily workflows [35, 36], and a need for AI to be accompanied with concise and comprehensive information regarding their capabilities and limitations [17].

To this literature, the present work establishes that even absent sophisticated matching algorithms, there are substantial gains from accurate, timely information provision, as a decision-aid to participants engaging in a human-driven matching process. This work thus serves as a foundation to develop collaborative human-algorithmic approaches to matching in complex, sensitive contexts—prioritizing *information discovery* as well as matching, given the data.

A growing body of research has also examined the impact of algorithmic decision support tools deployed in practice [40–42, 62]. One salient cluster of studies focuses on child welfare [39, 68, 69]. This work has in part centered on studying the impact that the Allegheny Family Screening Tool—a decision support tool that calculates a score representing the risk of child maltreatment allegations—on helping social workers decide appropriate next steps. Findings from these studies point to misalignments between algorithmic predictions and stakeholders' decision-making objectives [39] and call for algorithmic decision-support tools that *augment* human discretion and incorporate holistic assessments, rather than focusing on deterministic scoring [69]. Further, these studies discuss how deployments of algorithmic decision support tools often fail to account for people's existing workflows, calling for research that explores "how best to support complementary human-AI performance in practice, in real-world organizational contexts" [39].

We answer these calls by contributing the design and deployment of a system in an understudied domain: adult foster care. In contrast to prior approaches, our system seeks to improve matching via dynamic information discovery, rather than computing scores or predictions given a static dataset. In doing so, we prioritize improvements in stakeholders' day-to-day workflows and job satisfaction, rather than end metrics such as decision quality. We also pay close attention to the needs and priorities of stakeholders on *both* sides of the matching process—in our context, discharge coordinators and care home operators.

3 Intervention Design

3.1 Research Context and Background

Before describing our intervention, we explain how patient discharge processes currently operate and discuss the status-quo state-provided data. To our knowledge, the state has offered data to support the ecosystem of foster care homes since June 2014. The primary purpose of the data is to verify legal operation of a home; nonetheless, discharge coordinators historically used this data to identify homes to discharge patients. The most recent dataset lists 1,240 homes as of August 2023.

Discharge process. Once a patient has been financially and medically assessed as being ready for discharge from the hospital and is in need of long term care,[1] a discharge coordinator makes referrals to shelters, rehabilitation facilities, case management agencies, and other community partners, seeking to place the patient with one of these resources.

Figure 1 details the coordinator workflow that begins with: (a) placing calls to homes (if they are able to reach them on the phone, itself a several-step process). Homes with potential vacancies are historically identified using the state-provided data; (b) during a conversation, the home may express interest in a patient and plan a visitation; (c) not all planned visitations turn to actual visitations (there is substantial drop-off); (d) an actual visitation, in which the home operator visits the patient in person, may result in a placement. There are numerous reasons why each stage may not proceed to the next. Discharge coordinators have indicated that, using state data, it might take upwards of 200 calls to get to a placement for a single patient. The system described in this paper

[1]This process includes medical risk assessments and securing Medicaid Long Term Care insurance. [20]

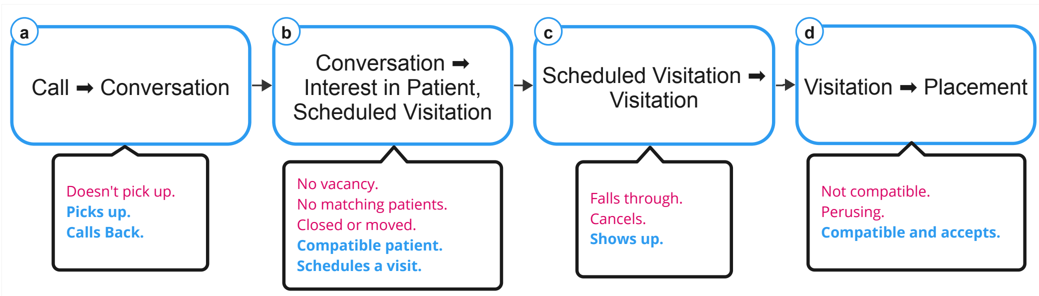

Fig. 1. A four part workflow for discharging a medically-ready, insurance cleared, long-term care patient: (a) a call which may lead to a conversation, (b) a conversation may lead to interest in a patient and scheduled visitation, (c) a scheduled visitation may lead to a visitation, and (d) a visitation may lead to a placement.

Fig. 2. Recorded Updates between 01/2014 and 12/2022. Mean gap: 105 days, SD: 98, minimum: 29, mode: 31, outliers: 365, 395.

focuses on improving results in the first two stages, (a) and (b), since better, more up-to-date data may result in more success in calling the correct homes. We leave to future work the important challenges faced in stages (c) and (d).

Description of historical state-provided data. To understand whether a data collection intervention is required and to inform system design, we conduct an analysis of state-provided historical data, collected via a state data portal. All homes, when first licensed to operate, fill out forms which are sent either by email or mail, describing their capabilities; these are published on an online data portal that hospital staff can access. Then, the onus is on care homes to update portions of this form on a case-by-case basis via various other pdf forms; for example, a home must either email or mail a "Remove Vacancy Form" to a government agency, who updates the database behind the state data portal. Participants discussed how the state data portal is frequently out of date, with updates to the data being slow and unreliable, occurring on average every 105 days.

To further illustrate this point, we analyze changes to the state data portal, including 29 data refreshes identified between January 2014 and December 2022. We collect this data primarily via website archives on Archive.org; from October 2021 through December 2022, we further verify in real-time the Archive updates, via weekly website polling for changes. While it is possible that Archive.org may have missed a posted update by the state, this seems unlikely as the archive recorded 93 snapshots over this timeframe, and we uncovered a significantly lower count of only 29 refreshes. Figure 2 shows the set of recorded snapshots, depicting when updates were provided. The gaps between refreshes amount to: a) an average refresh rate of 105 days, b) a standard deviation of 98 days, and c) a minimum of 29 days, mode of 31, and two outliers of 365 and 395 days.

Over time, this data has been refined to indicate more precisely where homes are located and what specific capabilities they have. In turn, this data has served not only in the interest of verifying legality of operation, but also in finding suitable long term care patient discharge locations, as the

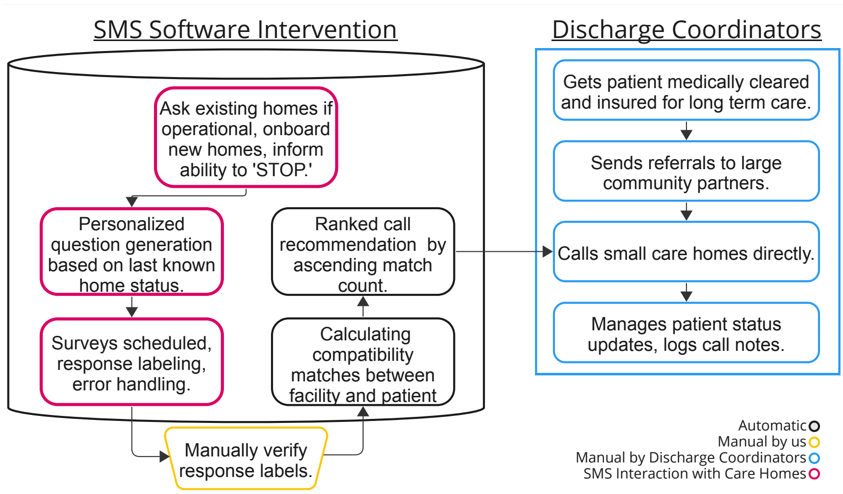

Fig. 3. Our system design and information flow, detailing the initiation of survey participation, survey roll-out, manual verification of responses, and ranked call recommendations fitting into a workflow coordinators at a hospital use to discharge patients. We highlight in pink, the stages of SMS participation by care homes we automate, in yellow what is manually done by us, in black what is automated and shared with coordinators, and in blue the manual work carried out by coordinators.

coordinators we worked with explained. We further observed several format changes in the data throughout the years, suggesting qualitative changes in its utility for placement.

The first change in data format we note is between June 2014 and October 2016; prior to this time, license expiration was not included; we interpret the subsequent inclusion of this data as the intent to facilitate license verification. In the early data updates, we see 7 columns for home location and contact information, and a remaining 11 columns intended for placement such as "How many patients do you need?", “Will you accept all Medicaid?” and “How many Male,Female?” We then note that from 2017 onward, there were 12 columns for home location and contact, one of these additional columns being an alternate phone number to contact the home, and two being for date of certificate expiration. However, columns related to assisting with patient placement were reduced to one—simply indicating home capacities. We explore why this change may have taken place in Section 5, though it is clear that the granularity of data provided for patient placement was reduced. These changes mark a potential shift in burden on discharge coordinators and social workers who may have relied on this data before these updates. It is also worth noting that this change corresponds with a decreased rate of the data being updated at all, as apparent in Figure 2.

Effective matching requires access to accurate and up-to-date information about homes’ capabilities and vacancies. We describe below how this can be achieved through the use of data collection and sharing platforms, such as the conversational SMS system discussed in this paper. Providing timely and accurate information can increase the efficiency of the matching process and help to ensure that patients are placed in homes that are well-suited to meet their needs.

3.2 Design Choices

Our system (flowchart in Figure 3) uses automated conversational SMS messages to communicate with care home operators and request status updates, which are manually verified before being passed on to discharge coordinators. Curated data and facility recommendations for each patient

Table 1. An example list of ranked call recommendations provided to discharge coordinators for a patient. This is the information a discharge coordinator received from our system.

Total Matches ↕	Facility	Recent Messages
1	Facility A	"male medicaid hospice less than 150 lbs"
3	Facility B	"good morning [...] i have 1bed open right now ..."
18	Facility C	"confirm"
...	...	...

are shown to discharge coordinators as a list (see Table 1); coordinators use this information to help decide which facility to contact regarding a patient. The right-most column of Figure 3 shows how these call recommendations fit into discharge coordinators' workflows, who work both to: a) get patients medically cleared, and b) get their insurance ready. Lastly, referrals are also sent to community partners, such as shelters or case management agencies (CMAs), to request that they consider accepting the patient. We detail their involvement further in Section 5.

The system was designed via collaboration among the first three authors (researchers and hospital staff) between February and December of 2021, in response to hospital stakeholders who asked about building a system to improve data quality and in turn support patient placement. Subsequently, the first author had multiple discussions with relevant stakeholder groups, including hospital leadership, discharge coordinators, and care home operators, to understand how they operated. These initial conversations, for example, surfaced that SMS communication would be a preferred mechanism for care homes and that sending SMS surveys roughly every 21 days would support coordinators' workflows. Based on the information obtained, the first author developed and proposed an initial system design to stakeholders. After developing an initial system design, we engaged in multiple rounds of iterative feedback and prototyping with stakeholders, including discharge coordinators and care home operators, starting with low-fidelity mockups and slowly transitioning to a functional system that was deployed. In these prototyping sessions, hospital discharge coordinators were involved with the design of the ranked list of matches, while both discharge coordinators and care home operators were involved in the design of the SMS exchanges. This incremental prototyping and feedback stage unfolded between June and December of 2021, with stakeholder feedback involving approximately four discharge coordinators and nine care home operators. After settling on a functional system, we tested it by sending pilot SMS surveys to a subset of 57 care homes, with the responses routed to a single discharge coordinator. The system functioned as intended, and we received a high SMS response rate of 64% over a 24 hour period. Following this pilot phase, we then deployed the system with all 1,240 care homes and progressively on-boarded five more discharge coordinators from the hospital team. The system has been fully operational since December 2021 and continues to operate at the time of writing. We now discuss important design choices that shape our intervention: (1) why we use SMS, (2) what questions we ask and why, and (3) how status updates are manually verified and integrated into ranked call recommendations.

Why SMS. SMS-based systems are a convenient and low-cost communication channel in the United States [63], especially for communities where smartphone apps may be a barrier. A wealth of prior research has shown the potential for SMS to support community health programs (e.g., [55, 58, 59, 76, 82]). In our setting, many home operators are themselves older adults who may be uncomfortable using new technologies [47, 54]. However, initial discussions with stakeholders suggested that SMS was an actively used form of communication among the population and preferred over smartphone apps or email. We use an external service to identify SMS-able phone numbers from the roster of care home operators[2]. Notably, (1) our SMS messaging is designed in compliance with federal consent and opt-in standards, as illustrated in Figure 3; and (2) the phone numbers shared by care home operators are provided for professional work use to the state data portal. In initial surveys with 57 homes, we received a high SMS response rate of 64% over a 24 hour period. In our study findings (Section 5) we show similar high response rates, received quickly, and sustained over the duration of our study.

Question Finding. Our initial stakeholder discussions also sought to identify questions being asked by stakeholders about one another. We identified three sources of questions: what *care home operators* want to know, what *discharge coordinators* are most often asked for, and what level of care is determined by a risk evaluation (regularly conducted by *social workers*) to secure health insurance. Each of these sources of questions have their own set of forms that each party fills out, which discharge coordinators and care home operators helped us rank in importance. For the purposes of this study we examine vacancy, weight and sex as changing factors.

We use this question priority order to progressively request updates via SMS. As shown in Figure 3 (left object, top box), all surveys began with an opt-in and explanation of how to opt-out: *"Welcome to [this] SMS service for care homes. At any moment you can unsubscribe by replying 'stop'. Please text '1' to begin."* This line of messaging is in accord with well established practices in automated messaging [24]. We subsequently ask in every survey to confirm the status of their vacancy or lack thereof. Each iterative survey contains a question specific to their status with respect to the queue of questions. Homes, for example, may receive messages such as:

> *"Good morning from QMC Care Coordination. The last update we have is your facility is full: 0 medicaid 0 private vacancies, and you have no patient preferences. Please respond 'confirm' if these are the same. Otherwise, respond, for example, with: '0 medicaid, 0 private, prefer male less than 180 [pounds].' So we can find the best fit when you do have a vacancy."*

The majority of responses receive a standardized thank you message to indicate their response or update was received, and that we will follow up via phone call if a vacancy and capability match is possible. However, if a home indicates a vacancy, and also has preferences which reduce their matching patient pool to less than three, we send a de-identified patient description for the home to assess if they might be interested in such a patient, for example: *"Please respond with '1' if interested, or '0' if not: 60y/o MEDICAID: UHC Female, 130-140 lbs."* In practice, it is rare that a home has constrained their patient pool to between 1 and 3 potential matches, and therefore many homes may not receive this message. Still, we found homes continue to participate, because they will nonetheless receive a phone call from a coordinator if they do have a vacancy. Section 5.3.3 discusses challenges around how to enable this information exchange, while heeding prior cautions on risks of 'gaming the system' and pitfalls of attempting to mitigate such gaming [39].

Hybrid Computer-Human SMS and Manual Labeling. While a significant portion of our survey deployment is automated, a similarly significant portion is enabled by manual, hybrid computer-human SMS. As described in prior work [58], hybrid computer-human SMS refers to messaging that combines automated and manual processes. In our setting, population-wide outbound surveys are automatically triggered, with personalized messages being generated based on a finite set of home capabilities; responses to these messages are then analyzed by a human to help determine, for example, the next home capability question to ask. The degree of automation and manual intervention can vary depending on the specific use case. For example, in some instances, the

[2]Homes that are not SMS-able communicate largely via landline phone; of the 1,240 care homes in the state data in August 2023, for example, 165 homes were not SMS-able.

Table 2. Summary of Interview Participant Demographics

	Discharge Coordinator (n=6)	**Care Home Operators** (n=4)	**Other participants** (n=2)
Age (years)	30–35: 2, 35–40: 2, 61+: 1, Did not say: 1	35–40: 1, 51–55: 1, 61+: 1, Did not say: 1	30–35: 1, Did not say: 1
Gender	Female: 5 Male: 1	Female: 3 Male: 1	Female: 1 Male: 1
Experience (years)	1–5: 2, 6–10: 3, 10+: 1	1-5: 1, 6-10: 1, 10+: 2	6-10: 2

human may only be involved in a few stages of the conversation, while in others they may be constantly monitoring and guiding the conversation. Updates received are processed and analyzed automatically, but with the option for human interaction to verify or provide additional information as needed. This hybrid approach afforded us efficient data collection, while still maintaining the accuracy and reliability of traditional human-led data collection methods.

We employ hybrid computer-human SMS primarily as a means of verifying that updates to a home's status are being conveyed to discharge coordinators correctly, and in some cases where error handling was insufficient to process a home's response, a human would follow up with a clarifying question. We differentiate between 'automated' and 'manual' in Figure 3, to delineate how we facilitate this hybrid approach. The manual work then feeds back to update home status and subsequently matching and call recommendations in the automated steps.

Labeled responses in turn drive automated compatibility matching and ultimately ranked call recommendation. Discharge coordinators are provided a list of call recommendations (Table 1), with homes ranked in order of ascending match count for each patient. This means that a home with no stated preference is most likely to be ranked at the bottom of the list, while homes stating specific preferences that match with, e.g., only one patient, are at the top of the list. The selection of which home to call is ultimately left to the discharge coordinator.

We also underscore here the manual work conducted by coordinators surrounding their use of our call recommendations. This includes both seeking out placement referrals from other community partners, such as shelters, rehabilitation facilities and case management agencies, updating patients and their needs, and keeping internal call notes on our system with respect to calling homes sourced through our SMS surveys. We discuss further in Section 5 how, throughout our deployment, 127 patients were uploaded and subsequently discharged in the system.

4 Field Study

We deployed our system with a local hospital in O'ahu, Hawai'i in February 2022 and continue to run the system and collect data at the time of writing. Here, we report on 14 months of study data, from February 2022 to April 2023. As described in Section 3.2, care homes were initially sent SMS messages that informed them the system was being operated by the hospital for the purpose of patient placement and instructed them on how to use it.

4.1 System Data Collection

From our analysis of state data, we find 1,840 unique homes ever listed, growing steadily since the July 2014 update which contained 300 homes. For our data analyses, we primarily consider the 1,047 homes that were operational and that had SMS-able numbers (verified by an external service) during our study period of February 2022 and April 2023. We sent out 16 surveys over our 14 month deployment, totalling 37,402 messages and 8,014 received messages from 16 surveys. We then manually verified status updates such as vacancies and preferences, as we detail in Section 3.2, for each responding home. Surveys were sent out approximately every 21 days (mode 21, median 23, standard deviation 4.1, with no outliers). This survey period of 21 days was determined in coordination with discharge coordinators, who felt that 21 days provided a conservative estimate of how frequently they contact care homes and would effectively balance their need for fresh data while not overburdening care homes. We further corroborate the utility of this rate in Figure 9, and discuss opportunities for future work to explore survey periods in Section 6. Our analyses omit the first 5 surveys, which were early prototypes of our messaging system, not sent to every home. We analyze the remaining 11 surveys: documenting response time, whether the home updated their vacancy status, and whether they updated their patient preferences.

4.2 Participant Interviews

To complement our system data, we conducted qualitative interviews to understand stakeholders' perspectives. Care home participants were recruited via randomized phone calls to reach both active and inactive participants. Hospital participants were recruited via our existing partnership with the hospital system in Hawai'i. All study procedures were IRB-approved.

Participants. We recruited 12 interview participants (see Table 2): six were staff on the hospital care coordination team using our system (identified as H1-6), and four were care home operators (CH1-4). All hospital participants had experience interacting with or using data from our system over the full deployment period before being interviewed, and all care home operators received SMS surveys, although one had been unresponsive (CH1).

The other two participants had different backgrounds: one was a case management agency owner (CM1) and one was a home licensing staff working with the State Department of Health (S1). Although these two participants had not used our system, they were able to provide valuable perspectives on how such a system might affect the broader landscape of patient placement.

Three participants were male, nine were female, and all had worked to help facilitate placement for more than five years, with an average of 12 years of professional experience. We also include in our analysis messages obtained over SMS, labelled as 'SMS.'

Interview Procedures. We conducted semi-structured interviews that explored participants' experiences and opinions regarding long-term placement broadly and with respect to our system. Interviews lasted roughly 45 minutes and participants were offered a $25 giftcard. Participants were asked what challenges they face in patient placement, work they have historically done to surface information, and for those using our system, successes or challenges they had with the system.

For the two participants who were not system users (the CMA owner and State interviewee), we did not ask about system usage or efficacy. Instead, we provided a system description and demo, after which we limited our questions to the implications of such a system for the broader landscape of patient placement (e.g., impact on non-hospital stakeholders like CMAs, accessibility of technology, etc.)

Data Analysis. In total we collected nine hours of audio recordings, which were transcribed before being analyzed via thematic analysis [13], wherein we analyzed transcripts with multiple passes reading each transcript, allowing an initial set of codes to emerge. The first author conducted the analysis and synthesis of these transcripts, which took place after quantitative data analysis. This process initially resulted in 88 codes, and after iterating on feedback with other authors eventually evolved into 34 codes ranging from "idiosyncracies of long term care" to "outdated or inaccurate knowledge." Codes were clustered into 6 high level themes such as "process efficacy" containing codes e.g. "measurement challenges", "redundancy of work", or the theme "placement fairness" containing codes e.g. "perceived rush to discharge", "patient difficulty", "re-admissions."

Table 3. Summary of our study findings.

Theme / Finding	Section
Why Discharge Coordinators Use Our System	**5.1**
Coordinators find our system surfaces data more accurate than state data and helps them with placement.	5.1.1
Prior to our system, inaccurate data led to excessive repetitive labor.	5.1.2
Coordinators find faster, more accurate information may help address burnout.	5.1.3
Why Care Homes Use Our System	**5.2**
Care homes exchange fast changing information with our system consistently.	5.2.1
Our system enables a gain of information for care homes.	5.2.2
Design of an easy-to-use communication channel, that aligns with existing workflows, increases its use.	5.2.3
Uncovering Fast-Changing Information and Preferences	**5.3**
Our system shows that home vacancies change rapidly over time.	5.3.1
Homes communicate preferences through our system, but may over-restrict their patient pool.	5.3.2
Preferences help avoid miscommunication about a patient, stakeholder incentives remain a challenge.	5.3.3
Opportunities and Limitations of Automating Placement Processes	**5.4**
Participants perceived an opportunity to improve our system to help with scheduling and assessments.	5.4.1
Data from our system enables coordinators to give their perspective to homes about patients.	5.4.2

These themes were then combined with our system usage data to form the mixed methods analysis presented below.

5 Findings

This section interweaves quantitative data collected via our system with qualitative insights from our interviews. Table 3 provides an overview of our findings.

5.1 Why Discharge Coordinators Use Our System

We begin by discussing how coordinators use our system, primarily replacing use of state provided data. We detail how the system reduces repetitive labor, decreasing burnout and leaving more time for non-automatable tasks.

5.1.1 Discharge coordinators use our system. We begin by demonstrating that our system is used by coordinators. For a relatively new system deployed in practice, we saw a substantial fraction of such long-term stay patients added to the system: 155 patients out of a total of approximately 400-500 such patients that the hospital would receive over the 14 month deployment. Although the total number of long-term stay patients is relatively small (i.e., 400-500), we note that this group is especially hard to place and makes up the single largest proportion of post-acute bed days at most hospitals: 1 in 7 nationally, and closer to 1 in 5 in our context [28]. Over the deployment period, 155 patients were added to our system, and coordinators have confirmed that 127 of these patients were discharged. Coordinators confirmed that at least one-third of these placements were to homes first contacted using our data and stated, *"We found several accepting [homes] from using [this system]"* (H1). The original source of the remaining patients is unknown, and some home placements may have been first contacted using other channels, such as community referrals. Although these other channels may still contribute to sourcing potential placement locations, coordinators shared that our data has replaced their use of state data. H6 told us, *"[with state data] you would get a lot of 'I'm*

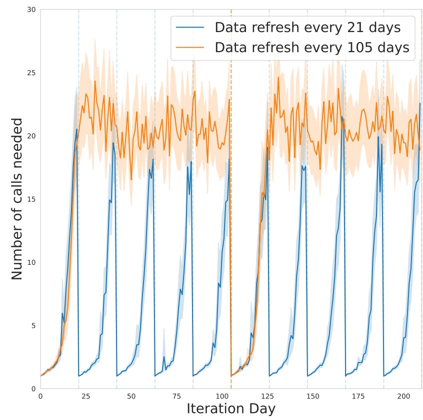

Fig. 4. Via simulation, the expected number of phone calls as a function of data refresh rate. Overall, the expected number of calls with a refresh rate of 21 days is 6 calls/patient. With a data refresh rate of 105 days, it is 18 calls/patient. We show two 105 day refresh cycles, and ten 21 day refreshes overlapped, with significant peaking of needed calls between 11 and 21 days, this sustains for the remainder between each 21 and 105 day cycle, which constitutes the 'runaway labor' we seek to alleviate, and sharply drops after a 21 day refresh.

full, I don't have any vacancy', so it was really time consuming ... [now we] mainly go off [this data], this at least narrows it down to people who actually have an opening."

5.1.2 Addressing runaway labor. As suggested above, the state data often led to coordinators calling homes that did not have a vacancy, even though the state data does contain vacancy information. To provide further intuition for the challenge tackled by our system, we formalize and quantify this effect, through a *calibrated simulation* of how coordinators call homes to match patients. We demonstrate that slow data refresh rates lead to "runaway labor", where relying on stale data means that coordinators may eventually call homes essentially at random.

This simulation was constructed using shared expertise with respect to current call workflows. At a high level, simulated coordinators call homes each day. They prioritize calling homes who, based on the latest available data, have vacancies and would be a match for the patient. However, each home's vacancy and capability status may change over time. Coordinators only have data from the last *data refresh* and from previous calls, so their data may no longer be accurate if a home's status has changed. Details are available in Appendix 9.1. We calibrate the parameters of our simulation to our findings regarding the state provided data and our system's data.

Figure 4 shows the average number of calls made (across multiple simulation runs) over time, at different refresh cadences: either every 105 days, representing historical state data refresh rates, or every 21 days, representing our system data. With fresh data (right after a data refresh, as indicated by the striped vertical lines), the coordinators only have to make a few calls, as the homes who had a vacancy at the last refresh still likely have a vacancy. As the data grows stale, however, the number of calls required to successfully place a patient grows rapidly, which we call "runaway labor." Eventually, with stale data the coordinators are calling essentially at random (which would have a call success rate of about 5%, or 20 calls until a home with a vacancy is found). Overall, more timely data in our simulation leads to a factor of three reduction in the number of phone calls needed between coordinators and care homes.

Coordinators we interviewed noted that, in practice, they can experience far greater call counts with state data. For example, one participant described how, with state data, *"it would take them hundreds and hundreds of calls before they could get any kind of hit."* (H2) As such we note that our simulation likely underestimates ground truth, which we understand is largely due to not incorporating other factors discussed in the background such as having a visitation fall through or other complex home preferences that may prevent matches.

5.1.3 Addressing Length of Stay and Burnout. A key long-term goal of our system is to reduce the length of a patient's post-acute hospital stay. We initially believed that being able to measure a reduction in length of stay would determine whether the system should continue to be used: if it was not measurably reducing length of stay, there was no use in continuing the system. However, patient length of stay is a function of many other factors, such as time varying effects, patient needs, home availability, among many others, meaning measurements, for example, of a counterfactual to a patient placement that reduced length of stay are highly complex. While reducing length of stay remains a central goal, we learned from interviews that there are immediate challenges workers face in their current day-to-day workflows, and we can do a lot of good to support their decision making processes. In particular, we find that our system helps reduce workloads, enabling coordinators to carryout more patient-focused work. Coordinators emphasized that the efficacy of our system—in reducing the number of calls needed—helps to mitigate burnout among staff, and this directly ties to job satisfaction:

> *"What they used to do felt thankless and redundant ... [coordinators] are the ones that the whole care team looks to, to be like, 'When are we getting this patient out?' So they feel a lot of responsibility ... [with] this system we have definitely seen some improvement in job satisfaction because it's actually effective at matching patients to [vacancies]."* (H3)

Throughout our study, participants discussed how the state provided data was out of date; in some cases home operators had either retired, closed, or were otherwise no longer operational. We received SMS responses from 28 homes that provided information such as *"Sorry, just retired"* (SMS). Without our system, a coordinator would not know that a home had ceased operation and might needlessly try to call. Instead, coordinators we interviewed discussed how our system enables them to perform more human-specific tasks rather than fighting through repetitive failed attempts at contacting non-operational homes:

> *"[This system] allows the coordinator to perform at the top of their skill set and have meaningful conversations with interested caregivers .. it takes away the inefficiencies and the time suck that is calling people to only get an out of service phone number."* (H2)

Although we did not quantitatively measure burnout for discharge coordinators using our system, this feedback suggests the system may help to reduce burnout by reducing redundant work, even if it may be for similar measured output with respect to a patient's post-acute hospital stay.

5.2 Why Care Homes Use Our System

We now analyze system usage among care homes and factors that contribute to continued usage, concluding that *community alignment matters*: it is essential to design an information gathering system that integrates with existing workflows and stakeholder priorities.

5.2.1 Care homes use our system. We saw sustained, rapid engagement throughout our deployment, with 99% of messageable homes responding at least once during the deployment: of the 1,047 homes messaged, 1,042 responded at least once. On average, about 40% of homes responded to any single survey. Figure 5 shows that engagement is sustained over the deployment. This level of voluntary engagement is higher than standard, modern phone polling rates. As a point of comparison, modern

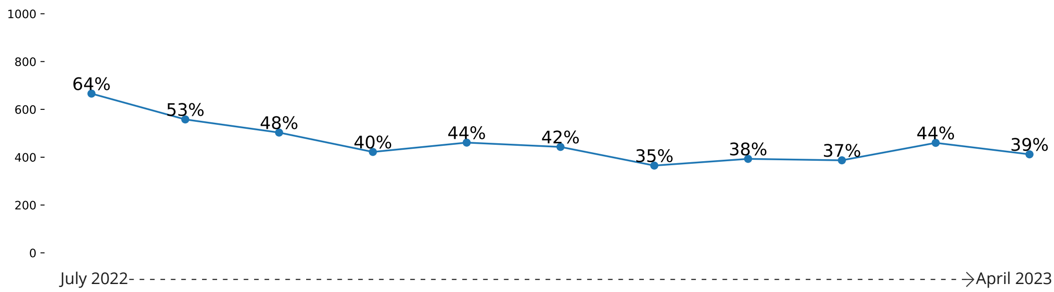

Fig. 5. A lineplot with 11 markers representing each survey, and corresponding percent of received responses across the participating 1,047 homes, between July 2022 and April 2023. We show an initial plateau from 64% response steadying out between 35% and 44%.

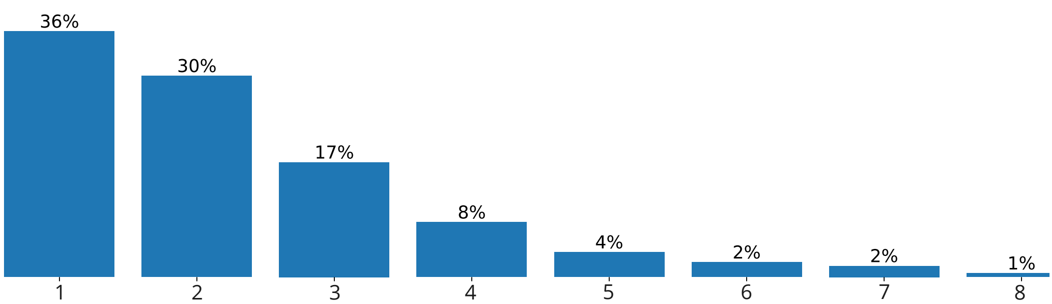

Fig. 6. Bar chart showing the proportion of responses received within ‘N’ hours, where by the end of the first hour, labeled ‘1’, 36% of responses are received. Response rates drop dramatically after the third hour; 83% of responses are received within the first three hours, following a long tail ending at hour ‘8’.

political polling by phone has less than a 5% response rate [6]. We also found that homes responded quickly: 83% of responses were received within three hours after a survey is sent (Figure 6).

We note that homes do not only respond when they have a vacancy; many respond to confirm a lack of a vacancy. Among homes that responded after not responding for two or more survey periods (about 6 weeks), only 10.8% stated they had a vacancy; the other 89.2% confirmed that they were full. Home operators also shared largely positive sentiment, with 22% of all received messages containing the word "thank".

5.2.2 Gaining Information. One reason that homes use the system is that they see how coordinators use the data they provide. Throughout our deployment, coordinators shared their use of vacancy and preference data to arrange phone calls, follow-up calls and visitations; subsequently, we saw homes, over text, indicate that they had *"received your call, will call you back"* (SMS). Interview participants confirmed this observation, for example: *"Yes, I received a call from [a coordinator], I remember ... I told her, do you want me to assess the patient?"* (CH4). These findings provide encouraging signs that home operators see their updates are being used by coordinators to make decisions.

Our system also provides more *personalized* information to homes, including about potential new candidate patients they might otherwise not have learned about. As discussed in Section 3.2, after a home confirms a vacancy and has an eligible match, they may be provided personalized information that includes de-identified information about specific potential patients. Coordinators saw our system leading to care homes that *“are highly motivated to use this system because they*

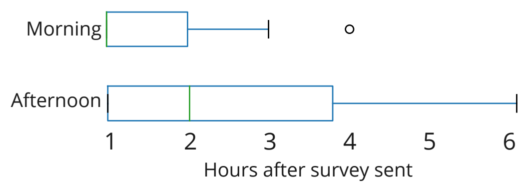

Fig. 7. Boxplots of when survey responses were received after a survey was sent, showing better engagement in the morning. We observe faster responses and a net increase in responses when survey waves are sent in the morning: on average, we receive responses from 401 homes when surveys are sent in the morning versus 336 homes when surveys are sent in the afternoon. Note that we cannot make causal claims regarding response rates, as survey timing was not randomized and we instead prioritized morning communication.

know that they do get way more information than they do from [coordinators]. They get accurate information. For those who have had successful placements, they know that it works" (H2).

Care home operators noted that our system may especially benefit homes that are more open to accepting patients and that have fewer constraints *"because they're willing to accept the challenge ... they don't look at so many things"* (CH2). Home operators also discussed how, with some patient descriptions, they can easily know *"if I look at those texts, I already have a feeling like oh yeah this [is not] appropriate for me"* (CH2). This suggests that the patient data received via SMS or through calls that result from information exchange over SMS is sometimes sufficient for homes.

Importantly, we note that although our system automates some personalized information transfer, there is nevertheless a substantial human component. Automated information provides the basis for a phone call between a coordinator and home, and a potential in-person assessment. The system thus combines automated information transfer and human skilled labor, focusing human effort on the type of information and work that cannot be automated.

5.2.3 Factors to Sustained Use. Beyond using our system to gain information, we note several factors that contributed to sustained use over time. A central design choice in our system is the exchange of information over SMS. As one home operator explained, many care home operators are older and would rather not have to learn an app *"or website or anything, they don't like that"* (CH2). Furthermore, software that may be normally used by coordinators may not be appropriate for care home operators: *"using [healthcare software] is a beast. It's not something that you can quickly get a hang of like SMS"* (H3).

Motivated by the need to align the system with existing workflows (which itself motivated the selection of SMS as a communication channel), we analyze whether survey response variance across waves is due to message timing, a factor discussed in prior work [7, 37, 61, 73]. Figure 7 shows that we observed more and faster response rates for messages sent in the morning, with responses decreasing and slowing down as the day progresses. This finding is consistent with the work day of a home operator: once the day has begun, patient care work becomes more pressing than engaging in an SMS survey, and is consistent with prior work showing that generally SMS participants in similar contexts will either respond promptly or not at all [37].

We also note that we sent personalized status update requests rather than general messages. We asked directed questions such as "Our last update from you was the following", as opposed to: "Do you have any vacancies?" This signaled to users that their messages were being stored and they had a status they could update, meaning their answers were not being ignored. When their status from the previous survey was communicated back to them, home operators often corrected it, for example, with *"actually no longer medicaid."* (SMS). We further enabled homes to respond by simply saying “confirm” to confirm their status, finding this response received high usage: of homes reporting an unchanged status, 40% used the “confirm” option.

5.3 Uncovering Fast-Changing Information and Preferences

We now analyze our system data to uncover (1) that vacancy information is fast-changing, with only about 20% of vacancies being reconfirmed in the following survey, and (2) that homes express preferences about patients, which may change over time. These findings underscore the importance of faster-updating, informative communication channels. Finally, we note the ambiguous effect of more information provided to care homes, and the challenges inherent in communicating subtle preferences via an automated system. This analysis was also used to inform our simulation of runaway labor (Figure 4).

5.3.1 How many vacancies are there at any given time and how long do they last? At each wave, care homes primarily responded to communicate if they have a vacancy. Figure 8 shows each response for each home over survey waves: whether homes responded, and how many vacancies they had. Following individual strands from left to right, representing single homes, one can see homes who, for example, were completely unresponsive, through to homes who consistently confirmed no vacancy, to homes who had communicated sporadic vacancies. On average, about 55 homes indicate a vacancy during each response. In our system, we impute non-response to a survey asking about vacancy as no vacancy; homes were both informed of this via messages and shown this status in the subsequent survey. This choice enabled coordinators to prioritize calling homes that were most likely to still have a vacancy. It also enabled homes to simply not respond if they don’t have a vacancy and thus may not see a benefit in responding.

In our system, vacancies are filled quickly. Of the average 55 homes that indicate a vacancy during a survey, only about 11 also confirm a vacancy in the following survey; a further 20 homes respond that they no longer have a vacancy. Figure 9 further shows the *conditional autocorrelation* for home vacancies across surveys: given that a home has a stated vacancy at a survey, what is the probability they state a vacancy τ surveys later? Only about 20% of vacancies last to the next survey and about 10% last 2 surveys. We note that this quick churn (as surveys are typically about 21 days apart, compared to the state data update frequency of 105) suggests that a data update rate of 105 days is insufficient to guide calling decisions: far before the data is updated, listed home capacities would no longer be informative.

5.3.2 Do homes have preferences about patients and are they consistently communicated? Beyond vacancies, home operators also express and change their capabilities and other preferences, such as the medical needs and patient sex. These preferences are important and substantially influence matches. Coordinators reiterated that *“if everyone could do everything, then it wouldn’t be so hard to match”* (H3). Here, we analyze preferences for patient sex or weight, which we saw were subject to change over time. For example, one home operator changed a preference to not take on heavier clients *"due to recent pregnancy"* (SMS). Preferences are thus important to communicate to coordinators attempting to place patients. In our data, 289 homes stated explicit sex preferences and 6% of these changed their preference at least once. Meanwhile, 340 homes stated weight preferences, with 10% of these changing their preference at least once during the deployment.

The consideration of what constitutes a preference is also subject to interpretation. For example, homes may be willing to forgo their preferences if it means receiving a patient. Coordinators pointed out that homes may *"blame [inability to take a patient] on a capability, when really it’s a preference"* (H3). This distinction, between a capability and a preference, we interpret as the difference between a hard and soft constraint. A capability, or hard constraint in this context, means the home is trained to handle a specific condition or has the requisite infrastructure (e.g., being

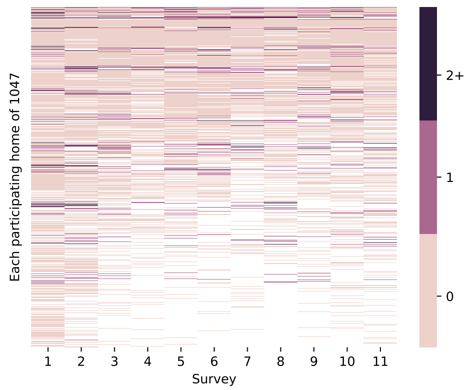

Fig. 8. A heatmap showing each participating home on the y-axis, and their vacancy status over 11 survey periods. Four increasing shades of color represent heat, the lightest, white, representing no response, makes up a quarter of the heatmap, the next darkest, beige, representing a confirmation of '0' vacancies, and makes up a large majority of the heatmap, with only a few speckled purple and black marks representing '1' and '2+' vacancies respectively.

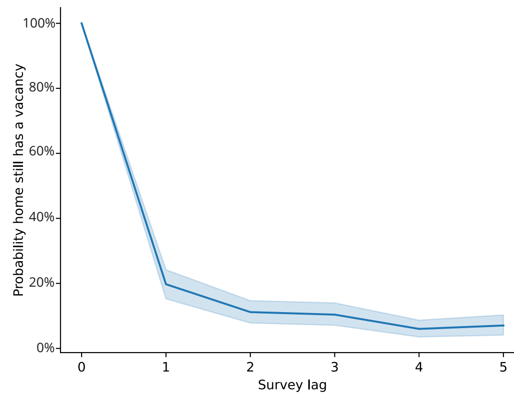

Fig. 9. A plot that shows: if a home states a vacancy in a survey, the probability that the same home also states a vacancy τ surveys (21 days) later, i.e., the probability a vacancy lasts over τ surveys. The plot shows that only 20% of vacancies are reconfirmed 1 survey after they are initially shared, with confidence intervals ±5%. This suggests that the state data, which updates on average every 105 days, would largely be stale and show vacancies that no longer exist.

able to accommodate dialysis). Such capabilities differ from soft constraints, or preferences, such as preferring that patients are ambulatory or not, or preferring patients between specific weight brackets, which may be more subject to change. One home operator shared, *"I thought [I would only take] private [pay] initially, but just as a preference"* (CH4). Four messages in our dataset also shared, for example, that homes had not yet received a call from a coordinator, with one stating *"update update update, never get anything"* (SMS). However, when we analyzed their responses, we found they had provided constrained preferences, such as a low weight limit that few patients meet and would likely not yield a match. This unintended consequence of a preference yielding no possible matches occurred for a small number of homes. Future work is needed on how to communicate potential consequences of such preferences and to enable homes to express more flexible preferences versus hard constraints (e.g., that they would prefer patients below a certain weight but may be willing to accept someone above that).

5.3.3 What level of preference communication is appropriate? Additional complicating factors to the placement workflow include accurately communicating patient information and home preferences, and also accounting for different incentives of care homes and coordinators when information is exchanged. Care homes must cope with the challenges associated with accepting a patient and want to avoid taking patients that might be too challenging. Meanwhile, coordinators may be working to place patients in a home to reduce their length of stay and may be wary of homes that refuse patients. We now study the role of such incentives during the placement process.

When asked to consider if their preferences were indeed necessary, care home participants elaborated that, especially with the complex patients our system primarily serves, they perceive that coordinators want to get them discharged and, as a result, not all information may get conveyed about a patient that is seemingly feasible for them to foster:

> *"Their focus is 'Get this patient out'. Like their acute stay is up, 'Let's get them out, find a home for them.' But when I asked them, 'Hey, so does that patient have a guardian? Who takes care of the finances? Do you know how much the patient makes a month?' [Coordinators] are like, 'Oh, I don't know' ... their level of care, sometimes they don't send it with the referral."* (CH1)

Gathering accurate information is a priority for home operators when receiving or looking for patients; they know that others might be incentivized to place a patient, but they will be the ones to bear the burden if finances are not in place or the patient is more challenging than communicated.

Coordinators raised concerns around the volume of patient information that could be shared with a home over SMS. They also noted that *"[homes] can be picky about who they place into their homes"* (H6). Coordinators further stated that although homes might make more informed decisions given more information, *"they might get super nitpicky. They see something in [the patient's information] that's from like 10 years ago and it may not even be relevant to the patient anymore"* (H6). They further expressed that communicating complex information might mislead homes: *"[homes] don't understand, so even if they got the [patient's] chart they would be seeing outdated data"* (H1). Similarly, home operators discussed how, ultimately, automated information exchange alone will be insufficient for the final task of placing a patient:

> *"I think it's very important that actual assessment should be done. Because it's impossible. I could have received all those texts... I would still want to do my assessment."* (CH2)

These challenges underscore a limit to communicating preferences through automated channels like SMS and stakeholders' concerns around accurately conveying preferences; these include coordinators' concerns that care homes may strategically alter their behavior to 'game the system' as a result of *"too much transparency"* [39]. Coordinators also believed that patient matching should be less subject to preferences to promote equitable patient treatment: *"If we're truly dealing with patients from an equitable standpoint... [matching] should be pretty objective based on the patient's care needs"* (H3), and to this end furthered a need for stricter certification requirements due to the difficulty in matching:

> *"[State licensing] should make … you get certified just if you can [cover] all those capabilities ... [and if not] can [state licensing] mandate that they use [your system] and mandate that they do these check-ins?"* (H3)

However, our state participant clarified that their work is focused not on deciding what capabilities care homes can cover, and instead handling *"reports of agencies placing [into] unlicensed care homes"* (S1), further clarifying the purpose for their data updates, and suggesting potential limitations to policies addressing placement challenges:

> *"The problem is ... care home operators will intentionally misrepresent that they have vacancies to get more business, or at least [referrals], so we seldom get updates that they don't have any vacancies … unfortunately there's no requirements in the administrative rules that they provide us with accurate data as far as [vacancies are] concerned."* (S1)

5.4 Opportunities and Limitations of Automating Placement Processes

Lastly, we discuss considerations associated with improving scheduling and assessment outcomes and the role agencies play in patient placement.

5.4.1 Opportunities to improve placement beyond initial calls and conversations. We model placement as shown in Figure 1: getting a coordinator to make a call and have a conversation that leads to interest in a patient, then a scheduled visitation, and finally a placement. Our system currently seeks to improve the initial parts of this process: improvements to calls which led to conversations and interest in patients that accounted for preferences. Beyond these initial stages, participants envisioned future improvements to our system that might help later stages in the placement process, namely scheduling and visitations.

Care home operators saw our system as potentially enabling them to conduct in-person patient assessments on their own terms:

> *"You [could] have standing appointments that are shared over SMS. Open assessment or date of assessment should be applicable to all the caregivers, like the date and time. Because you will have more caregivers that would respond to that."* (CH2)

Suggesting that, by allowing homes to schedule into open assessment slots, they would be more proactive in taking on patients. On the other hand, coordinators cautioned *"even if there were open appointments, we need to wake up the patient and prepare them"* (H5). They agreed that *"we should be making it so that we're available for the [home] ... we should be able to accommodate that"* (H3). Currently, assessments are scheduled in low volumes one at a time. Participants saw the potential for our system to scale up assessments, potentially increasing efficiency, though they cautioned about the challenges managing multiple caregivers assessing one patient: *"Oh, it's gonna be chaotic . . . all the caregivers at the same time doing assessment for one patient"* (CH2). Coordinators shared that while they do schedule single assessments at a time, they rarely have a single assessment result in a placement, meaning scheduling assessments also requires overbooking interest in a patient such that multiple homes are given a chance to decide on taking on a patient.

With respect to visitations, coordinators empathized that the challenge in assessing a patient reaches beyond scheduling, and that care homes carry out difficult work simply to visit patients in-person: *"They have to find a ride, they have to find a replacement to replace them [at the home]. It's not like the case management agency is doing all this for them"* (H2). A coordinator also noted:

> *"I think for a caregiver who's taking care of [activities of daily living], you shouldn't need to do an in-person assessment. It could all be telehealth. Because you could show them the wound. You could show them everything. The head to toe. You could get to know the person, lay eyes on them, start to match the condition with the person. So I think it would lead to a much quicker placement if they didn't have to visit [in person]."* (H3)

This participant went on to add that the speed with which SMS messages are received means *"you could do like, 'Hey, I have you [texting] right now. Let's switch to a [virtual assessment]'"* (H3).

5.4.2 Considerations around agency involvement. While participants imagined several future improvements enabled by our system, our system exists within a complex multi-stakeholder ecosystem. As such, non-hospital community partners can play important roles assisting in patient placement, which our system does not replace. One such entity is a Case Management Agency (CMA) who, in some cases, can serve as an alternative match-maker between discharge coordinators and homes. Their role was self-described during our interviews as *"make sure the [patient] is placed in a good place... we see them once a month... we make sure the caregivers are doing their job"* (CM1). Our CMA participant emphasized the importance of their human work in this process, namely in clarifying misconceptions about a client, in a way that simple information exchange may not be able to accomplish: *"I see the [coordinators] are like, 'Help me place this patient.' I said, 'No, there's no guardian. Don't bother me with that one. I don't want to stress out'"* (CM1). They further view their role as protecting care homes from difficult patients or identifying patients whose needs differ from their medical charts.

Since our system was developed in collaboration with and for use by hospital staff, CMAs are not its intended users. Some coordinators raised concerns about the system bypassing CMAs and *"going behind CMA backs"* (H4), because coordinators *"already send referrals to CMAs"* (H4). Other coordinators argued that our system acted to help coordinators *"assist CMAs with finding interested caregivers, so if we find an interested caregiver, we refer them to the CMA"* (H1).

At the same time, coordinators and care homes saw significant benefits to direct communication with each other. Indeed, they discussed how CMAs often ask care home to call coordinators anyway and that, by proactively communicating with care homes, coordinators share information that homes would not have received from a CMA about a patient: *"the [CMAs] usual response to patients that are less than easy is 'no interested caregivers'"* (H1). Instead, the calls facilitated by our system help coordinators voice their perspective on the patient directly to potential caregivers.

Simultaneously, we saw that care homes wanted to have more direct communication with coordinators, perceiving that, by exchanging their preferences over SMS, they would receive more frequent calls directly from coordinators. They were also dissatisfied with CMAs: *"CMA's are charging like crazy for placement. [If] I share updates over SMS, I want a call from you [the coordinator] not the CMA"* (SMS). We see clear benefits to directly connecting coordinators and care homes via our system. Such communication may both help and harm CMAs, who work as middlemen in the patient placement pipeline. Future work might consider if or how it may be appropriate to expand data access with CMAs.

6 Discussion

This paper presents an information sharing and decision support system that helps coordinators place older adult patients into adult foster care homes. As a working system in a high-stakes domain, it must continuously engage with stakeholders, integrate with existing workflows, and enable coordination among stakeholders. Here, we discuss challenges and opportunities associated with the long-term viability of such a system: (1) the importance of centering improving workflows as opposed to just end metrics; (2) opportunities to support coordination and flexibility in patient placement processes; and (3) sustainably and ethically scaling across contexts.

6.1 Centering improvement of workflows as opposed to end metrics

Our system tackles what can be understood as a matching problem: how to best match patients ready for discharge with homes who have vacancies. As such, natural end success metrics may include the number of matches our system finds, the quality of those matches, and how long post-acute patients must stay at the hospital waiting for a placement. Related high-stakes dynamic matching and algorithmic decision support literature has followed this approach of primarily measuring against such end metrics, such as organ exchange [5] with respect to maximizing organ availability while minimizing donor swaps; risk assessment tools in criminal justice [32] that aim to recommend whether someone should be released; resource allocation in homelessness settings to decide allocations that minimize simulated recidivism to homelessness [64, 80]; and matching in refugee resettlement [3] to maximize placement success, such as employment likelihood.

However, although we are able to quantify that at least a third of the patient placements we focus on can be traced back to our matching efforts (Section 5.1.1), we cannot yet compare against the counterfactual placement rate for these patients without our system, i.e., if they would have been placed regardless of our system. Instead, our contribution is to show that, by provisioning accurate, timely information to support decisions along a pathway to placement—even absent sophisticated matching algorithms optimizing for an end-metric—we create improved stakeholder workflows. We demonstrate a foundation towards dynamic matching that also prioritizes maintenance of up-to-date information (as opposed to a static dataset) and enablement of matching given the data.

Towards building on this foundation, one example for how our system might optimize toward an end-metric might be to gather data about prior patient placements, and generate a forecast for placement difficulty in order to help workers in prioritizing the placement of new patients as they come in. Such an approach would be similar to the use of algorithmic decision support for child welfare [22] which helps decide how to move forward while balancing potential new allegations of child maltreatment with risk assessment scores. However, as several recent studies point out [39, 69], such an approach may lead to misalignments between algorithmic predictions and worker decision-making processes. While the decision support tools studied in this literature and our system both support decision making at the earliest stages, such as when a new case or patient is available, our work instead strives to center the immediate daily needs of the workers we seek to support. This means working to surface fresh data on a regular basis about the potential available resources for *both* care homes and coordinators: patients and vacancies, and thereby facilitating timely exchange of up-to-date information.

6.2 Supporting coordination in patient placement processes

Key to our focus on supporting daily workflows is aiding the coordination work required for patient placement. As discussed in Section 2.2, our system acts as a distributed coordination mechanism between a heterogeneous array of external, individually-operated care homes, and an internal team of hospital discharge coordinators. As detailed in Figure 1, discharging a patient is a multiple step process that requires substantial coordination, starting with discharge coordinators making calls to potential homes and ending with placement after a visit by a care home operator. In collaboration with stakeholders, our design focused on the initial stages of this system: the rate at which calls lead to conversations with compatible homes with vacancies.

One reason to focus on these initial stages is the relative ease with which technology is able to automate the collection of vacancy and some capability and preference information from homes, compared to the complexity of orchestrating an assessment. This information can be relatively standardized, does not require patient-specific details, and reduces the likelihood that coordinators repeatedly call homes without vacancies. By collecting this information via SMS, a widely accessible communication channel, we see how our system successfully establishes a *"shared platform of communication"* [11], which Bossen and Grönvall discuss as a key requirement to enabling coordination among diverse stakeholders.

At the same time, our design heeds calls for coordination mechanisms that are *"flexible enough to enable standard and ad hoc communication and coordination"* [11, 15]. In our case, although there are benefits to maximally standardizing messaging and exchange of preferences and patient needs, there are also benefits to providing flexibility in how care home preferences are set and communicated (Section 5.3.2). One way our system affords this flexibility is by enabling home operators to respond via natural language, rather than forcing homes' responses to conform to pre-specified structures.

However, this flexibility yielded new challenges. For example, we saw how some homes unknowingly expressed stringent preferences that precluded them from matching with any patient (Section 5.3.2). Such homes may then become frustrated that using the system does not lead to calls from coordinators about patient matches. Moreover, since a key benefit that homes derive from our system is gaining information about potential patients (Section 5.2.2), homes may strategically alter their survey responses, for example by indicating an openness to accept any patient, when in reality they have constraints. While this may lead to the home receiving more calls from coordinators, these calls may increase coordinator workload and be unlikely to actually lead to a match.

These challenges suggest rich opportunities for future work, towards a system that effectively solicits granular care home preferences and constraints in a way that *"enables a comprehensive overview of required information and establishes shared understandings of information among collaborators"* [11]. For example, the system might alert homes whose resulting number of patient matches is either very low or very high, and provide opportunities for them to consider adjustments to their expressed preferences and constraints.

Beyond care home preferences, we see further design opportunities to increase flexibility in ways that aid the work of discharge coordinators. In particular, our current system design sends surveys to homes at fixed intervals every 21 days, a period that was decided in collaboration with hospital staff. Such a design may push discharge coordinators to schedule their work according to when new vacancy data has freshly arrived. However, a more flexible model might instead seek to instead provide fresh data to coordinators on demand, changing the system so that it conforms to coordinators' desired schedules. For example, a discharge coordinator planning to call homes about a patient could request fresh data; the system could then send a survey to a subset of homes that it determines are most likely to have a vacancy, and provide just-in-time data to aid the coordinator. Across these future efforts, a key system design consideration will remain the focus on how technology might automate work that can easily be automated, while supporting coordinators and home operators in their more specialized, patient-specific workflows.

6.3 Sustainability and scaling present additional challenges

When deploying computational systems for high-stakes contexts such as ours, research has highlighted the need to plan early on for "sustainability, maintenance, and repair" [57]. To ensure long-term sustainability, our system was built in close collaboration with our hospital partner and carefully attends to their and other stakeholders' priorities and workflows. As a result, the system has been continually operating since February 2022 and, importantly, our hospital partner has assumed financial costs of running the system, a clear indication that it provides value.

At the same time, we saw challenges to ensuring equitable access to the system. As detailed in Section 5.2.1, some care homes did not have listed phone numbers and could not be reached via SMS, while some do not regularly respond to our surveys. Especially as hospital discharge coordinators increasingly rely on our system, more work will be needed to ensure these homes are not excluded from patient placements.

In addition, only discharge coordinators in our collaborating hospital can currently access our system and data. However, potentially expanding the system to include other hospitals in Hawai'i will require caution. For example, complications may arise if discharge coordinators at different hospitals "compete" to place patients with the same set of homes. Strategies will need to be developed to ensure that both homes and hospitals equitably benefit from the system and contribute to updating the system's data after calls with homes.

Finally, future work is needed to understand if or how our system might work in other regions or domains. Most states in the US offer some form of residential long-term care, and specifically small scale adult foster care [31]. However, prior work [44, 48] has highlighted the significant challenges in generalizability and cautioned against assuming expansion of deployments, and discussed the important role that domain specific expertise can play in the success of any such expansion or deployment [21]. For example, our system currently requires some manual labeling of home responses before the database is populated; this labeling is not possible without an understanding of the kinds of preferences homes have, as well as various stakeholder incentives (Section 5.3).

7 Conclusion

We design, build, deploy, and evaluate a conversational SMS system in Hawai'i, to aid in the discharge of older adult patients from a hospital to adult foster care homes. We provide qualitative and quantitative evidence that our system improves participants' workflows by providing more timely, accurate information to inform matches, aiding coordination between care homes and hospital discharge coordinators. Taken together, our findings advance the CSCW community's knowledge of how to design and deploy multi-stakeholder information collection and sharing systems in high-stakes domains.

8 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 Millenium Scholar Program, NG is supported by NSF CAREER IIS-2339427 and an Amazon research award.

References

[1] Joanna Abraham and Madhu C Reddy. 2010. Challenges to inter-departmental coordination of patient transfers: a workflow perspective. *International Journal of medical informatics* 79, 2 (2010), 112–122.
[2] Mark S. Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing Knowledge and Expertise: The CSCW View of Knowledge Management. *Computer Supported Cooperative Work (CSCW)* 22 (2013), 531–573. https://api.semanticscholar.org/CorpusID:8183637
[3] Narges Ahani, Paul Gölz, Ariel D Procaccia, Alexander Teytelboym, and Andrew C Trapp. 2021. Dynamic placement in refugee resettlement. *arXiv preprint arXiv:2105.14388* (2021).
[4] Asbjørn Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2021. Street-Level Algorithms and AI in Bureaucratic Decision-Making: A Caseworker Perspective. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 40 (apr 2021), 23 pages. doi:10.1145/3449114
[5] Itai Ashlagi and Alvin E. Roth. 2021. Kidney Exchange: An Operations Perspective. *Manag. Sci.* 67 (2021), 5455–5478.
[6] Michael Bailey. 2023. A New Paradigm for Polling. *Issue 5.3, Summer 2023* (2023). https://api.semanticscholar.org/CorpusID:260274230
[7] George Balabanis, Vince Mitchell, and Sarah Heinonen-Mavrovouniotis. 2007. SMS-based surveys: strategies to improve participation. *International Journal of Advertising* 26 (2007), 369 – 385.
[8] Karthik S. Bhat, Amanda K. Hall, Tiffany Kuo, and Neha Kumar. 2023. "We Are Half-Doctors": Family Caregivers as Boundary Actors in Chronic Disease Management. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 111 (apr 2023), 29 pages. doi:10.1145/3579545
[9] Pernille Bjørn and Morten Hertzum. 2011. Artefactual Multiplicity: A Study of Emergency-Department Whiteboards. *Computer Supported Cooperative Work (CSCW)* 20 (2011), 93–121. https://api.semanticscholar.org/CorpusID:1890691
[10] Claus Bossen, Lars Rune Christensen, Erik Grönvall, and Lasse Steenbock Vestergaard. 2013. CareCoor: Augmenting the coordination of cooperative home care work. *International journal of medical informatics* 82, 5 (2013), e189–e199.
[11] Claus Bossen and Erik Grönvall. 2015. Collaboration in-between: The care hotel and designing for flexible use. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 1289–1301.
[12] Tone Bratteteig and Ina Wagner. 2013. Moving healthcare to the home: The work to make homecare work. In *ECSCW 2013: Proceedings of the 13th European Conference on Computer Supported Cooperative Work, 21-25 September 2013, Paphos, Cyprus*. Springer, 143–162.
[13] V Braun and V Clarke. 2006. Using thematic analysis in psychology. *Qual. Res. Psychol.* (2006).
[14] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019).
[15] Federico Cabitza and Carla Simone. 2013. Computational Coordination Mechanisms: A tale of a struggle for flexibility. *Computer Supported Cooperative Work (CSCW)* 22 (2013), 475–529.
[16] Christine Caffrey. 2012. *Residents living in residential care facilities: United States, 2010.* Number 91. US Department of Health and Human Services, Centers for Disease Control.
[17] Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 104 (nov 2019), 24 pages. doi:10.1145/3359206
[18] Clara Caldeira, Matthew Bietz, Marisol Vidauri, and Yunan Chen. 2017. Senior Care for Aging in Place: Balancing Assistance and Independence *(CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1605–1617. doi:10.1145/2998181.2998206
[19] Shiye Cao and Chien-Ming Huang. 2022. Understanding User Reliance on AI in Assisted Decision-Making. 6, CSCW2, Article 471 (nov 2022), 23 pages. doi:10.1145/3555572

[20] Centers for Medicare & Medicaid Services. 2023. Institutional Long-Term Care. https://www.medicaid.gov/medicaid/long-term-services-supports/institutional-long-term-care/index.html. Accessed: 05/02/2023.

[21] Parmit K. Chilana, Amy J. Ko, and Jacob O. Wobbrock. 2015. From User-Centered to Adoption-Centered Design: A Case Study of an HCI Research Innovation Becoming a Product. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015). https://api.semanticscholar.org/CorpusID:16966854

[22] Allegheny County. 2023. The Allegheny Family Screening Tool. https://www.alleghenycounty.us/Human-Services/News-Events/Accomplishments/Allegheny-Family-Screening-Tool.aspx Accessed: 09/13/2023.

[23] Judd N. L. Cramer and Alan B. Krueger. 2016. Disruptive Change in the Taxi Business: The Case of Uber. *Entrepreneurship & Law eJournal* (2016). https://api.semanticscholar.org/CorpusID:155799194

[24] CTIA. 2019. CTIA - Messaging Principles & Best Practices. https://www.ctia.org/the-wireless-industry/industry-commitments/messaging-interoperability-sms-mms. (Accessed on 05/10/2023).

[25] Andrea Cuadra, Jessica Bethune, Rony Krell, Alexa Lempel, Katrin Hänsel, Armin Shahrokni, Deborah Estrin, and Nicola Dell. 2023. Designing Voice-First Ambient Interfaces to Support Aging in Place. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 2189–2205. doi:10.1145/3563657.3596104

[26] Brian DeRenzi, Nicola Dell, Jeremy Wacksman, Scott Lee, and Neal Lesh. 2017. Supporting Community Health Workers in India through Voice- and Web-Based Feedback. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) *(CHI '17)*. Association for Computing Machinery, New York, NY, USA, 2770–2781. doi:10.1145/3025453.3025514

[27] Brian DeRenzi, Leah Findlater, Jonathan Payne, Benjamin Birnbaum, Joachim Mangilima, Tapan Parikh, Gaetano Borriello, and Neal Lesh. 2012. Improving Community Health Worker Performance through Automated SMS *(ICTD '12)*. Association for Computing Machinery, New York, NY, USA, 25–34. doi:10.1145/2160673.2160677

[28] 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

[29] Economic Policy Institute. 2022. Residential Long-Term Care Workers. https://www.epi.org/publication/residential-long-term-care-workers/.

[30] Administration for Community Living. 2021. Profile of Older Americans. https://acl.gov/aging-and-disability-in-america/data-and-research/profile-older-americans. Accessed: 05/01/2023.

[31] Paying for Senior Care. 2023. Medicaid's Adult Foster Care and Adult Homes State by State. https://www.payingforseniorcare.com/medicaid-waivers/adult-foster-care Accessed on: 09/10/2023.

[32] Sharad Goel, Ravi Shroff, Jennifer L Skeem, and Christopher Slobogin. 2021. The accuracy, equity, and jurisprudence of criminal risk assessment. *Research Handbook on Big Data Law* (2021). https://api.semanticscholar.org/CorpusID:238153081

[33] Ben Green and Yiling Chen. 2021. Algorithmic Risk Assessments Can Alter Human Decision-Making Processes in High-Stakes Government Contexts. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 418 (oct 2021), 33 pages. doi:10.1145/3479562

[34] Tom Gross. 2013. Supporting effortless coordination: 25 years of awareness research. *Computer Supported Cooperative Work (CSCW)* 22 (2013), 425–474.

[35] Azra Ismail, Divy Thakkar, Neha Madhiwalla, and Neha Kumar. 2023. Public Health Calls for/with AI: An Ethnographic Perspective. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 354 (oct 2023), 26 pages. doi:10.1145/3610203

[36] Azra Ismail, Deepika Yadav, Meghna Gupta, Kirti Dabas, Pushpendra Singh, and Neha Kumar. 2022. Imagining Caring Futures for Frontline Health Work. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 480 (nov 2022), 30 pages. doi:10.1145/3555581

[37] Douglas Johnson. 2016. Collecting Data from mHealth Users via SMS Surveys: A Case Study in Kenya. *Survey practice* 9 (2016), 2824.

[38] Rosalie A. Kane, Robert L. Kane, L H Illston, John A. Nyman, and Michael D. Finch. 1991. Adult foster care for the elderly in Oregon: a mainstream alternative to nursing homes? *American journal of public health* 81 9 (1991), 1113–20.

[39] Anna Kawakami, Venkatesh Sivaraman, Hao Fei Cheng, Logan Stapleton, Yanghuidi Cheng, Diana Qing, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022). https://api.semanticscholar.org/CorpusID:247958405

[40] Jon M. Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2017. Human Decisions and Machine Predictions. *Economics of Networks eJournal* (2017). https://api.semanticscholar.org/CorpusID:6050316

[41] Neha Kumar and Nicola Dell. 2018. Towards Informed Practice in HCI for Development. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 99 (nov 2018), 20 pages. doi:10.1145/3274368

[42] Vivian Lai and Chenhao Tan. 2018. On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection. *Proceedings of the Conference on Fairness, Accountability, and Transparency* (2018). https://api.semanticscholar.org/CorpusID:53774958

[43] Amanda Lazar, Hilaire J. Thompson, Shih-Yin Lin, and George Demiris. 2018. Negotiating Relation Work with Telehealth Home Care Companionship Technologies That Support Aging in Place. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 103 (nov 2018), 19 pages. doi:10.1145/3274372

[44] Allen S. Lee and Richard L. Baskerville. 2003. Generalizing Generalizability in Information Systems Research. *Inf. Syst. Res.* 14 (2003), 221–243. https://api.semanticscholar.org/CorpusID:6087319

[45] Charlotte P Lee and Drew Paine. 2015. From The matrix to a model of coordinated action (MoCA) A conceptual framework of and for CSCW. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 179–194.

[46] Jiachen Li, Bingrui Zong, Tingyu Cheng, Yunzhi Li, Elizabeth D. Mynatt, and Ashutosh Dhekne. 2023. Privacy vs. Awareness: Relieving the Tension between Older Adults and Adult Children When Sharing In-Home Activity Data. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2, Article 353 (oct 2023), 30 pages. doi:10.1145/3610202

[47] Qi Ma, Alan Hoi Shou Chan, and Pei-Lee Teh. 2021. Insights into Older Adults' Technology Acceptance through Meta-Analysis. *International Journal of Human–Computer Interaction* 37 (2021), 1049 – 1062. https://api.semanticscholar.org/CorpusID:234182018

[48] Lena Mamykina, Arlene M. Smaldone, Suzanne R. Bakken, Noemie Elhadad, Elliot G. Mitchell, Pooja M. Desai, Matthew E. Levine, Jonathan N. Tobin, Andrea Cassells, Patricia G. Davidson, David J. Albers, and George Hripcsak. 2021. Scaling Up HCI Research: From Clinical Trials to Deployment in the Wild.. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI EA '21)*. Association for Computing Machinery, New York, NY, USA, Article 50, 6 pages. doi:10.1145/3411763.3443437

[49] Angela Mastrianni, Lynn Almengor, and Aleksandra Sarcevic. 2022. Alerts as Coordination Mechanisms: Implications for Designing Alerts for Multidisciplinary and Shared Decision Making. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 9 (jan 2022), 14 pages. doi:10.1145/3492828

[50] Joy Ming, Elizabeth Kuo, Katie Go, Emily Tseng, John Kallas, Aditya Vashistha, Madeline Sterling, and Nicola Dell. 2023. "I Go Beyond and Beyond" Examining the Invisible Work of Home Health Aides. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 59 (apr 2023), 21 pages. doi:10.1145/3579492

[51] 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.

[52] Terry D. Moore, Thomas P. McDonald, and Kari Cronbaugh-Auld. 2016. Assessing Risk of Placement Instability to Aid Foster Care Placement Decision Making. *Journal of Public Child Welfare* 10, 2 (2016), 117–131. doi:10.1080/15548732.2016.1140697 arXiv:https://doi.org/10.1080/15548732.2016.1140697

[53] Leslie A. Morgan, J Kevin Eckert, Ann L. Gruber-Baldini, and Sheryl Itkin Zimmerman. 2004. Policy and Research Issues for Small Assisted Living Facilities. *Journal of Aging & Social Policy* 16 (2004), 1 – 16. https://api.semanticscholar.org/CorpusID:5729600

[54] Nasrin Navabi, Fatemeh Ghaffari, and Zahra Jannat-Alipoor. 2016. Older adults' attitudes and barriers toward the use of mobile phones. *Clinical Interventions in Aging* 11 (2016), 1371 – 1378. https://api.semanticscholar.org/CorpusID:20267943

[55] Fidèle Ngabo, Judith Nguimfack, Friday Nwaigwe, Cathy Mugeni, Denis Muhoza, David R Wilson, John Kalach, Richard Gakuba, Corrine Karema, and Agnes Binagwaho. 2012. Designing and Implementing an Innovative SMS-based alert system (RapidSMS-MCH) to monitor pregnancy and reduce maternal and child deaths in Rwanda. *The Pan African Medical Journal* 13 (2012).

[56] American College of Emergency Physicians. 2022. *Emergency Department Boarding Crisis Sign-on Letter.* https://www.acep.org/siteassets/new-pdfs/advocacy/emergency-department-boarding-crisis-sign-on-letter-11.07.22.pdf Accessed: 05/02/2023.

[57] Chinasa T. Okolo, Srujana Kamath, Nicola Dell, and Aditya Vashistha. 2021. "It cannot do all of my work": Community Health Worker Perceptions of AI-Enabled Mobile Health Applications in Rural India. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (2021). https://api.semanticscholar.org/CorpusID:233987440

[58] Trevor Perrier, Nicola Dell, Brian DeRenzi, Richard Anderson, John Kinuthia, Jennifer Unger, and Grace John-Stewart. 2015. Engaging Pregnant Women in Kenya with a Hybrid Computer-Human SMS Communication System *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1429–1438. doi:10.1145/2702123.2702124

[59] Anthony Poon, Sarah Giroux, Parfait Eloundou-Enyegue, François Guimbretiere, and Nicola Dell. 2020. Baccalauréat Practice Tests in Cameroon: The Impact of SMS-Based Exam Preparation. In *Information and Communication Technologies and Development (ICTD '20)*. 13. doi:10.1145/3392561.3394646

[60] Anthony Poon, Matthew Luebke, Julia Loughman, Ann Lee, Lourdes Guerrero, Madeline Sterling, and Nicola Dell. 2023. Computer-Mediated Sharing Circles for Intersectional Peer Support with Home Care Workers. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 39 (apr 2023), 35 pages. doi:10.1145/3579472

[61] Stephen R. Porter, Michael E. Whitcomb, and William H. Weitzer. 2004. Multiple Surveys of Students and Survey Fatigue. *New Directions for Institutional Research* 2004 (2004), 63–73.

[62] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Wortman Vaughan, and Hanna Wallach. 2021. Manipulating and Measuring Model Interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 237, 52 pages. doi:10.1145/3411764.3445315

[63] Noshad Rahimi and Rachanida Koosawangsri. 2013. Selection of smart phone plans. *2013 Proceedings of PICMET '13: Technology Management in the IT-Driven Services (PICMET)* (2013), 426–448. https://api.semanticscholar.org/CorpusID:44254097

[64] Aida Rahmattalabi, Phebe Vayanos, Kathryn Dullerud, and Eric Rice. 2022. Learning Resource Allocation Policies from Observational Data with an Application to Homeless Services Delivery. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (2022). https://api.semanticscholar.org/CorpusID:246275792

[65] James R. Reinardy and Rosalie A. Kane. 1999. Choosing an adult foster home or a nursing home: residents' perceptions about decision making and control. *Social work* 44 6 (1999), 571–85.

[66] James R. Reinardy and Rosalie A. Kane. 2003. Anatomy of a Choice: Deciding on Assisted Living or Nursing Home Care in Oregon. *Journal of Applied Gerontology* 22 (2003), 152 – 174.

[67] Sheria Robinson-Lane, Erica Solway, Dianne Singer, Matthias Kirch, Jeffrey Kullgren, and Preeti Malani. 2022. Older adults' preparedness to age in place: findings from the national poll on healthy aging. *Innovation in Aging* 6, Supplement_1 (2022).

[68] Devansh Saxena, Karla A. Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the U.S. Child Welfare System. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020).

[69] Devansh Saxena, Karla A. Badillo-Urquiola, Pamela J. Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proceedings of the ACM on Human-Computer Interaction* 5 (2021), 1 – 41. https://api.semanticscholar.org/CorpusID:235765287

[70] Kjeld Schmidt and Carla Simone. 1996. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Computer Supported Cooperative Work (CSCW)* 5 (1996), 155–200. https://api.semanticscholar.org/CorpusID:1122632

[71] Peter G Scupelli, Yan Xiao, Susan R Fussell, Sara Kiesler, and Mark D Gross. 2010. Supporting coordination in surgical suites: Physical aspects of common information spaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1777–1786.

[72] Phoebe Sengers, John McCarthy, and Paul Dourish. 2006. Reflective HCI: Articulating an Agenda for Critical Practice. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (Montréal, Québec, Canada) *(CHI EA '06)*. Association for Computing Machinery, New York, NY, USA, 1683–1686. doi:10.1145/1125451.1125762

[73] Jennifer Sinibaldi and Anton Örn Karlsson. 2016. The Effect of Rest Period on Response Likelihood. *Journal of Survey Statistics and Methodology* 5, 1 (11 2016), 70–83. doi:10.1093/jssam/smw027 arXiv:https://academic.oup.com/jssam/article-pdf/5/1/70/11127076/smw027.pdf

[74] Brenda C Spillman and Kirsten J Black. 2021. *Size of the long-term care population in residential care: a review of estimates and methodology*. https://aspe.hhs.gov/reports/size-long-term-care-population-residential-care-review-estimates-methodology-1 Accessed: 05/02/2023.

[75] Allan Stisen, Nervo Verdezoto, Henrik Blunck, Mikkel Baun Kjærgaard, and Kaj Grønbæk. 2016. Accounting for the invisible work of hospital orderlies: Designing for local and global coordination. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 980–992.

[76] Yanfang Su, Changzheng Yuan, Zhongliang Zhou, Jesse Heitner, and Benjamin Campbell. 2016. Impact of an SMS advice programme on maternal and newborn health in rural China: study protocol for a quasi-randomised controlled trial. *BMJ Open* 6 (2016).

[77] Yuling Sun, Xiaojuan Ma, Silvia Margot Lindtner, and Liang He. 2023. Data Work of Frontline Care Workers: Practices, Problems, and Opportunities in the Context of Data-Driven Long-Term Care. *Proceedings of the ACM on Human-Computer Interaction* 7 (2023), 1 – 28.

[78] Charlotte Tang, Yunan Chen, Karen Cheng, Victor Ngo, and John E Mattison. 2018. Awareness and handoffs in home care: coordination among informal caregivers. *Behaviour & Information Technology* 37, 1 (2018), 66–86.

[79] Nervo Verdezoto, Naveen Bagalkot, Syeda Zainab Akbar, Swati Sharma, Nicola Mackintosh, Deirdre Harrington, and Paula Griffiths. 2021. The Invisible Work of Maintenance in Community Health: Challenges and Opportunities for Digital Health to Support Frontline Health Workers in Karnataka, South India. *Proc. ACM Hum.-Comput. Interact.* 5,

CSCW1, Article 91 (apr 2021), 31 pages. doi:10.1145/3449165

[80] Hau yin. Chan, Eric Rice, Phebe Vayanos, Milind Tambe, and Matthew H. Morton. 2018. From Empirical Analysis to Public Policy: Evaluating Housing Systems for Homeless Youth. In *ECML/PKDD*.

[81] Kirti Zeijlmans, Mónica López, Hans Grietens, and Erik J. Knorth. 2017. Matching children with foster carers: A literature review. *Children and Youth Services Review* 73 (2017), 257–265. doi:10.1016/j.childyouth.2016.12.017

[82] Petros Zerfos, Xiaoqiao Meng, Starsky H. Y. Wong, Vidyut Samanta, and Songwu Lu. 2006. A study of the short message service of a nationwide cellular network. In *ACM/SIGCOMM Internet Measurement Conference*. https://api.semanticscholar.org/CorpusID:7455983

[83] Zhan Zhang and Aleksandra Sarcevic. 2018. Coordination Mechanisms for Self-Organized Work in an Emergency Communication Center. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018), 1 – 21. https://api.semanticscholar.org/CorpusID:53249047

9 Appendix.

9.1 Modeling Patient Matching and Efficient Labor

Algorithm 1 Patient Placement Simulation

```
Data: Ground Truth Vacancies and Capabilities Y, Knowledge of Ground Truth Ŷ, List of
Patients (θ0, θ1... θN)
function PlacementSimulation(patients, facilities, data refresh τ)
    Sort Ŷ by descending vacancies (e.g. 1, 1 ... 0, 0)
    for patient in patients do
        Set matched = False, update ground truth, Y
        if patient = increment of τ then
            Perform knowledge refresh: set Ŷ = Y; re-sort by descending vacancies
        while not matched do
            Call next facility from Ŷ where the facility-patient pair is compatible
            if Facility has a vacancy then
                Assign facility to patient and update Y, Ŷ
            Update facility state and knowledge for called facility
    Return list of calls performed per patient
```

Simulation description. Our simulation is detailed as Algorithm 1. The simulation proceeds over time, with N facilities, each with at most one vacancy (guaranteeing at least one facility that has a vacancy). Each time step (day), the true status of each facility (both vacancy status and patient preference) may update. Whether they have a vacancy is updated with probability p (otherwise, their status remains the same as the previous time step). Whenever their status is updated, they have a vacancy with probability q, and no vacancy otherwise. Preference updates are analogous. However, the information system used by discharge coordinators does not have perfect data: rather, the information system receives a knowledge refresh every τ time steps. At each time step, a discharge coordinator is attempting to place a single patient. They do so using the latest data that they have available, calling only facilities that they believe have a vacancy and have preferences that match the patient.

We calibrate the parameters to the data from our system. We set $N = 1,047$ facilities, with probability of a vacancy at $q = 0.05$ (as, at any given time, about $\frac{55}{1,047}$ facilities respond that they have a vacancy). Vacancy update probability p is set to be consistent with Figure 9, in which about 20% of facilities that have a vacancy at day t still have a vacancy about 21 days later at day $t + 21$:

i.e., p is set to be the solution of $(1-p)^{21} + (1-(1-p)^{21})q = 0.2$, implying $p \approx 0.08$.[3] For simplicity, we just model sex preferences, and model these preferences as static: 60% of patients are male, 40% are female; 80% of facilities have no sex preference, while 13% prefer females. Finally, we simulate two data refresh rates: every 21 days (matching our system), and every 105 days (matching the mean update cadence for the state data).

Received January 2024; revised July 2024; accepted October 2024

[3]Either this facility never had their vacancy status updated, or had their status updated at least once and the last update led them to having a vacancy.