A systems approach to healthcare:

Agent-based modeling, community mental health, and population well-being

by

Barry G. Silverman[a], Nancy Hanrahan[b], Gnana Bharathy[a], Kim Gordon[a], Dan Johnson[a]

[a] - Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104

[b] - School of Nursing, University of Pennsylvania, Philadelphia, PA 19104

Contact: basil@seas.upenn.edu

Abstract

Purpose: Explore whether agent-based modeling and simulation can help healthcare administrators discover interventions that increase population wellness and quality of care while, simultaneously, decreasing costs. Since important dynamics often lie in the social determinants outside the health facilities that provide services, this study thus models the problem at three levels (individuals, organizations, and society).

Methods: The study explores the utility of translating an existing (prize winning) software for modeling complex societal systems and agent's daily life activities (like a Sim City style of software), into a desired decision support system. A case study tests if the 3 levels of system modeling approach is feasible, valid, and useful. The case study involves an urban population with serious mental health and Philadelphia's Medicaid population (n=527,056), in particular.

Results: Section 3 explains the models using data from the case study and thereby establishes feasibility of the approach for modeling a real system. The models were trained and tuned using national epidemiologic datasets and various domain expert inputs. To avoid co-mingling of training and testing data, the simulations were then run and compared (Sect. 4.1) to an analysis of 250,000 Philadelphia patient hospital admissions for the year 2010 in terms of re-hospitalization rate, number of doctor visits, and days in hospital. Based on the Student T-test, deviations between simulated vs. real world outcomes are not statistically significant. Validity is thus established for the 2008 to 2010 timeframe. We computed models of various types of interventions that were ineffective as well as 4 categories of interventions (eg, reduced per-nurse caseload, increased check-ins and stays, etc.) that result in improvement in well-being and cost.

Conclusions: The 3 level approach appears to be useful to help health administrators sort through system complexities to find effective interventions at lower costs.

1. Introduction: Community health promotion and quality of life

Social and human ecologic determinants of health and health differentials have recently risen in the political agenda of the US and other societies leading to renewed interest in disease prevention, in health promotion, and in the systems approach [1] to improving quality of life (QoL) in order to reduce reliance on sick care. We view the management and promotion of health as a complex systems problem. It is a system of systems since it has numerous layers from the biologic to the cultural and many of its layers are themselves networked systems (eg., human physiology, a family, a community). Further it is a complex system since many of its parts are purposeful, have their own (often hidden) motivations, and behave probabilistically often leading to unexpected, emergent patterns.

One of the best tools for attempting to understand and better manage complex systems is modeling and simulation. Our goal is to place a decision support tool into the hands of health administrators. Ideally, this decision support system (DSS) includes an agent-based model of various stakeholders' motivations and micro-decision making -- like a Sim City for health – that allows users to easily simulate and visualize all of the interacting parts and thereby confront costs, outcome improvements, and benefits of alternative health promotion projects in different regions. By using systematic modeling that incorporates the most salient aspects of social determinants of health and systems science, the DSS will help administrators standardize assessment and decisions about programs and expenditures and thereby manage scarce resources more efficiently and effectively and achieve desired outcomes.

But what exactly should be modeled? How do we know what to model and what to leave out? After all modeling is expensive and simulation is time consuming (though not modeling can be far more costly). Fortunately, systems science suggests that to understand any system, it is important to represent its three levels -- how its parts work (humanization problem), how the whole works (control problem), and how it interacts with its containing system (environmentalization). If we model a community, the parts might be the individual residents and the various practices that serve the community; the whole would be all the practices and health system services available to community residents, and the container would be the community including its population groups and members, organizations they belong to, and society at large.

We hypothesize that a DSS is useful for studying how to improve operations at all three levels. First, since the traditional fee-for-service business model of healthcare is unsustainable over the long haul, it is important to understand how policy changes impact the redesign of a given practice. A DSS allows decision makers to use systems thinking. Systems' thinking is foreign to most US-trained doctors and the medical system with its traditional focus on pathogens, risk factors, treating disease, episodic care, and on

2

running clinics as a transactional fee-for-service business. One needs a systematic way to evaluate the plethora of disruptive innovations proposed for fixing the situation: e.g., medical homes, retail medicine, nurse practitioners providing primary care, reimbursable self-run support groups, health club activities, etc.: e.g., see Christiansen et al [2]. Which alternative interventions yield the greatest savings and the greatest improvement in health status? What are the tradeoffs? How sensitive are overall results to changes in particular features that might be implemented? How long will it take for the interventions to begin producing net savings? How might screening programs that identify these illnesses at an earlier stage improve outcomes?

Making decisions about resource utilization became a high priority with the 2010 Affordable Patient Care Act that mandates accountability for quality of care while reducing costs [3]. Multiple chronic illnesses and mental health conditions place the highest burden of cost on health care budgets and these costs are mostly shouldered by public payment systems. DSS are needed to assist health care administrators improve quality at reduced costs. However, few studies have used agent-based modeling despite advantages of accounting for complex social determinants in the analysis.

In short, this paper reports on results to date of our investigation of the usefulness of three levels of models of a community – the overall healthcare "system", the various practices and services that comprise its "parts", and the key stakeholders (organizations, groups, individuals) in the overall "containing society". Over the past 15 years, we have successfully applied this 3-level systems modeling approach for the US Department of Defense and State to analyze societal instabilities overseas and how to influence them for the better [4-6]. The current paper examines a prototype and its validity and usefulness in addressing instabilities in a large urban healthcare system. We named the model *Simulating Urban Mental Health Operations* (SUMHO) and we model health systems of the city of Philadelphia.

1.1. Case study: community mental health in Philadelphia

The population with serious mental illness (SMI) with comorbid medical problems represents one of the most complex resource allocation problems for public health administrators.[7] People with chronic mental and physical illness have great need and are associated with some of the highest costs of health care and the worst outcomes. For the SMI, rehospitalization rates and medication errors elevate the cost of care by $44 billion dollars a year [8]. Certain conditions in the social construct of the overall health care delivery system make individuals with SMI extremely vulnerable. The delivery systems for mental health, substance and physical health care operate independently, communicate with one another inefficiently and

often have different financing arrangements and policies [8]. These systems are virtual silos. Patients negotiating these fragmented health care systems find them not only burdensome but perilous and result in exacerbation of symptoms and rehospitalization.   For example, these challenges ignite problems following discharge from a hospitalization for a psychiatric condition and result in striking statistics: 38% of clients with serious mental illness relapse within 30 days of discharge and must be rehospitalized; 48% show high utilization of emergency services [7]; and people with SMI die 25 years earlier than those in the general population [9]. DSS could improve the effective and efficient management of this population.

Urban behavioral health systems have long been a safety net for people with serious mental illness, caring for some of the sickest and poorest individuals. Just in the city of Philadelphia alone, community mental health care is an $850 million dollar annual industry that supports up to 500,000 Medicaid beneficiaries; of these beneficiaries, approximately 100,000 are individuals with multi-complex mental and medical health problems and consume nearly 75% of the resources [7]. These systems rely on federal and state funding and have experienced deep budget cuts as states close their budget deficits. Now more than ever, public health administrators must spend resources efficiently, reducing costs while still providing essential services. Because of the complexity of the system, however, there is great need for tools to guide administrators' resource utilization decisions.

1.2. Why Agent-based modeling (ABM) & modeling challenges

ABM is a new way of understanding social systems that grew out of complexity science and artificial intelligence. When equations and principles cannot adequately describe and predict a complex system's overall macro-behavior, analysts can use a mixture of quantitative and qualitative methods to construct agent-based models of key stakeholders and their micro-decision making. Initially, these start out as simple rule-based models, which grow more sophisticated as they are validated and details are added. When these models are run, one can observe synergies that emerge, study new equilibria that arise in response to policy variables, and explore the causes of micro and macro-behavior patterns that have been observed in the real world.

A system is a whole that cannot be divided into its interdependent and interacting parts without losing the essence (synergy). Further, the micro-decision making of each of its parts has an effect on the macro-behavior of the whole (emergence). A unique challenge of social systems is that there are many sub-systems that are themselves purposeful systems -- many levels of self-determining functionality, from the depths of the cognitive up to the heights of the economic institutions and political strategies -- and one

must find ways to encapsulate them in hierarchies or networks, so that different levels may be meaningfully studied. This network of teleologic sub-systems, this independent purposefulness of the parts, is the distinguishing trait of social from other types of systems. It means that social systems are the hardest ones, the most complex – what are referred to as ill-structured and having "wicked" dilemmas.

ICSU [1] also points out that "Systems analysis involves the use of mathematical modelling to represent processes and relationships and how desirable end points may be achieved." To pursue a systems approach to inform city (Health and quality of life) planning, we thus must be capable of modeling complex social systems (synergy, purposefulness, emergence) at multiple levels: the individual, social groups, institutions, and so on. We chose to do this via agent-based modeling (ABM). There is no universally accepted definition of the agent approach, but for our purposes here, we will state that *agents are software entities that have mental states and can sense, think, and act with some degree of autonomy to carry out goals of their own choice.* ABM helps us learn about a complex system and its dynamics since we model the micro-decision making of the system's parts (e.g., diverse agents' choice functions) and use that to study the emergent macro-behavior of the collective.

2. Survey of agent methods

Realizing that ABM is a useful paradigm for studying urban health promotion and well-being is not, by itself, sufficient insight to guide us as to what kind of agent modeling we should attempt. The ABM field has been around for some time, and there are now many types of ABMs to select from. Specifically, the literature can be roughly segmented into 3 categories of agents: broad-shallow, narrow-deep, and mixed. We address each of these in the following three sub-sections. We posit that a complex social system like urban health and well-being will best be addressed by mixed modeling agents, a concept we define more fully in this section. But briefly, this means that the agent modeling environment allows us to model the agents as purposeful (i.e., cognitive entities) yet existing within higher level social constructs that influence them and that they can influence as well. In short, we need to make use of many of the approaches to agent (and other types of) modeling.

2.1. Broad-shallow category of agents

This is the most widely implemented category of agents. It is sometimes also referred to as Multi-Agent Systems (MAS), though that term can be a bit ambiguous as it can cover everything from a few agents interacting on a team up to billions of agents interacting on a landscape. Generally, by broad-

shallow agents, we posit that the agent modeling environment is reusable across many domains (broad), that agents are scalable (potentially large numbers of them), and generally their decision making involves only a few simple rules, as few as needed to make the agents believable and social (they are shallow). Thus there is very little cognition in these agents. A popular example is a grid where each agent's next state is updated to reflect the states of a majority of agents in nearby grid cells. This cell agent approach has found utility for modeling the spread of ideas, disease, and products: [10]

Alternatively, MAS agents are often simple Finite State Machines and the rules are hand-scripted behaviors. Examples are flocking, collision avoidance, rules of war (moving, shooting and fighting – in videogames, movie characters, etc.), and simple decision making such as moving away from threats (predators, enemies, diseased people, adverse states of nature, etc.) and toward safer locales and/or more acceptable experiences. Also, this occurs with agents in game theory [11].

A popular focus for discrete event simulation is the modeling of workloads and workflows in settings within hospitals and clinics: e.g., see Brailsford et al [12]. One of the few agent-based approaches to this problem is HealthSim of Christiansen [2]. Each agent has physiology models that represent the dynamics of various diseases. HealthSim also models workflows which can be recursively divided into component steps such that each step represents an action or series of actions that is carried out by participants filling various roles (patient, healthcare provider, etc) and with the help of specific resources and facilities for a given time interval. HealthSim thus illustrates a number of useful workflow modeling ideas that, in theory, should be extendable beyond its current pathogen-based, disease treatment focus, an effort we attempt here.

2.2. Narrow-deep category of agents

While broad-shallow agents are widely applicable and scale well, their shallowness implies that there are many domains in which they will fail to satisfy. There are times when one needs to have aspects of agent behavior be guided in a deep way by biologic, sociological, psychological, economic, political, etc. fields of concern. To address this concern, researchers fielded a growing number of cases of narrow-deep ABMs. These typically add depth in one or two dimensions and often rely on one or a few narrowly focused models drawn from a given discipline such as social networks, cognitive models, physiologic models, and so on.

An example of the narrow-deep category of agents may be found in the program designed by Schlessinger & Eddy [13]. It appears that Archimedes does what it does exceedingly well because it is

focused on only two primary topics: physiologic modeling and care treatment impact modeling. Archimedes is first and foremost a set of models all within one person, one model for each of multiple diseases. It is largely pathogen-based, disease oriented, and risk factor focused. For example, while it looks at multiple diseases and preventing risk factors before (or treating diseases once they) occur, it does not consider the overall community health promotion.

This brings us to another example of narrow-deep agents - that of cognitive modeling. Agents with cognitive models are capable of independent and autonomous perception, thinking/deciding, and reasoning about action execution. A cognitive architecture is a broadly-scoped, domain-generic computational cognitive model, capturing essential structures and processes of the mind, to be used for a broad, multiple-level, multiple-domain analysis of cognition and behavior [14]. Some well-studied cognitive architectures are ACT-R [15], SOAR [16], and Cognet [17], among others [18]. Instead of a superficial model specifically designed to study emergence (often in an ad hoc way as with shallow rules), using a cognitive architecture forces modelers to think in terms of the mechanisms and processes centered on primitives of cognition as envisioned in the cognitive architecture, and therefore such explanations are deeper explanations.

Despite these theoretical advantages, when it comes to implementation, many of the existing cognitive models suffer from the fact that they model a single, unembodied person doing only a few tasks. And while they may do those tasks as well as an actual person, few of them model the range of tasks that a person does during the course of a day. Thus for example, we have cognitive models that solve algebra problems OR Lisp program problems, but nothing else. None of the mainstream cognitive models have a body or physiology, so they do not experience fatigue, stress, or diseases. Worse still, most cognitive models are incapable of modeling a community, a social group, or even relationships between people.

2.3. Mixed modeling category of agents

Mixed modeling operates at the level between the macro and micro of the two former categories. They are not generally able to scale as well as the macro or shallow-broad category since they have more depth and detail. And, they are not as deep and fine-grained as the micro level models. However, they tend to integrate many of the capabilities of both categories – so they combine cognitive functioning with physiologies, affective computing with relationships, social networks and cultural norms, and action choices with workflows. Mixed modeling agents are still bottom up – they include emergence of macro-behavior from complex collectives of micro-decision making agents – and as such are the best fit for the type of well-being motivation and reasoned action type of modeling that we are interested in here.

According to Zacharias et al [19], there are very few agent systems that occupy this space, though they do cite the PMFserv socio-cognitive agent framework utilized here. In that vein, for the past 15 years, we have been conducting research on scientific agent-based models that could simulate a wide array of agent behaviors (cognitive, affective, relationships, daily activities), economic/religious/political organizations, and societal systems. This research has been sponsored by the US military and various intel agencies in order to model various cultures, leaders, followers, and communities around the world. It is currently matured and, in March 2011, it won 1st place for best Artificial Intelligence/Pattern of Life in the Federal Virtual World Challenge: see Village Video and Silverman [20]. It is now capable of being ported over to civilian uses. Hence, this article on our porting efforts to date.

3. Toolbox for mixed modeling of social systems

In terms of the history of this toolbox, we started by modeling individuals from the social psychological and mind-body perspective (deep-narrow) and then added more models that broadened the capability in the economic-political-sociologic dimensions. Specifically, the components included here are as follows and as documented, respectively, in Silverman et al. [6, 20]. In 2011/12 we applied these tools to model the community mental health of Philadelphia. In this section, we will illustrate the models with examples from this domain. Section 4 will then show results from running these models.

3.1. The PMFserv architecture

The basic building block is a model of an agent's bodily needs (physiology, fatigue, hunger, injuries, belonging, etc) and mind. The agent mind is based on Ortony et al [21] where agents apply their moralistic value system, form relationships, appraise the world, and autonomously carry out courses of action. A number of recognized instruments are used to profile the traits, cognitions, relations, alignments, and reasoning of individual leaders, followers, and others. A performance moderator function (PMF) is a micro-model covering how human performance (e.g. perception, memory, or decision-making) might vary as a function of a single factor (e.g. hunger, need for sleep, injury, event stress, time pressure, grievance, and so on). PMFserv culls from the literature dozens of best available PMFs and synthesizes them within a unifying framework and thereby offers a family of models where micro-decisions lead to the emergence of macro-behaviors within an individual. PMFserv includes a plugin architecture that facilitates turning on and off different models and trying new ones. An intel agency, for instance, sponsored us (2000-2005) to

implement part of their leader profiling methodology inside of PMFserv. There are too many models to review them all here, but we will sample a few of them.

As an example of just the cognitive appraisal loop, let us examine some agents representing different populations within the city of Philadelphia. The primary components of this framework are Motives, State, and Actions as summarized in Figure 1 and explained in what follows.

Figure 1: Framework of agent behavior in PMFserv

Agent motives – In cognitive appraisal, an agent's motives arise based on its value system. In PMFserv, Goals, Standards, and Preferences (GSP) Trees specify a value system that is based on multiattribute utility theory and Bayesian probability mathematics. In this context, utility refers to the non-monetary measure of satisfaction one derives from various outcomes or situations. These utilities are multi-attribute, meaning that there are several attributes or values (hence, the use of a tree) that are considered when evaluating the utility of a certain outcome. Finally, the use of Bayesian mathematics allows a frequency distribution of the past choices of an agent to define the relative important of each branch of the value system to the agent. An example of a portion of such a value system for a cognitive agent is presented in Table 1. There we see that Goals are the near-term events that an agent wants to happen, standards are the socially and morally-imposed guidelines for behavior, and preferences are the ideals about the long-term and the desired state of the world. In order to drive its own behavior, an agent evaluates these three dimensions to assess the utility of any potential action it could take as well as those taken by other agents. For instance, an agent with the bad patient archetype value system described by Table 1 is apathetic and short term gratification-oriented. It also believes its choices have little impact on outcomes (perceived control) and it views itself as generally ineffective (low efficacy). In the case of SUMHO, we plan to research the branches more fully in terms of what defines a relevant range of patient behaviors, and then based on actual patient histories; we will try to derive important agent archetypes or personas in terms of specific values we assign to each node in a GSP Tree for that personality.

Table 1: Portion of a value system of a cognitive agent (bad patient archetype)

Agent state properties - In the current SUMHO prototype, the well-being of an agent is described along three dimensions: their physiological, mental, and socioeconomic state. Each of these three dimensions contains numerous metrics, called state properties, each of which describes a different aspect of the state

9

of an agent. Table 2 shows the state properties that an agent representing a mental health patient in the current version. These states are variously computed by dynamical models in the agent (eg, hunger, stress, mood), models and/or rule sets external to an agent (eg, employment opportunities, injuries, and social support), and some are set as initial variables that might be changed stochastically over time (eg, income, living conditions, psychotic intensity).

<div align="center">Table 2 – Categories of state properties in an agent</div>

3.2. FactionSim

This level of PMFserv adds in a model of the social and organizational roles that exist in a community of interest (e.g., multi-state, state, sub-state, or region/neighborhood) and that may be played by the PMFserv agents. Developed under a 3 year grant from AFOSR, FactionSim implements a number of recognized scientific theories from sociology (mobilization, leader-follower theories, motivational congruence, social norms, etc) and political-economy (developmental economics). To apply it to a region, one profiles the capabilities, norms, allegiances, enmities, etc. of the formal and informal organizations and networks as well as the key stakeholder roles.

In our Philadelphia community mental health case study, we organized the groups according to the widely used four quadrant clinical integration model [22]. The groups categorize those who have low mental and physical health problems (Quandrant I), high mental health problems and low physical health problems (Quandrant II); low mental health and high physical health problems (Quandrant III); and, high mental and physical health problems (Quandrant IV). Thus, individuals in Quadrants I& II function well with few medical and mental health problems; in contrast to those who have greater functional limitations from physical health problems with high mental health comorbidities—Quadrant III or greater functional limitations from mental health problems with high medical comorbidities—Quadrant IV. Individuals categorized in Quadrant III & IV often have major social needs due to poverty, unemployment, unstable housing. These people are often insured by Medicaid or Medicare or both. The population with serious mental illness (SMI) and comorbid medical problems represents one of the most complex social problems for public health administrators. Druss et al [23] show that two thirds of people with SMI do not access timely physical or mental health care and do not receive adequate treatment.

Causes and consequences of wellbeing are woven in a complex web of social–cultural–technological conditions and associated human decisions. For instance, belonging to a group (quadrant),

sets up social norms such as what such a member is expected to be capable of. But this is tempered by a number of other factors as well. For example, do they have family supports and live in a safe area? Or are they living in dangerous neighborhoods where they might be preyed upon resulting in loss of their payment checks and medications, or exposed to drug dealers and become dependent on substances. The agent state properties are also used to capture and represent further variability such as, among other factors: acuity and risks of disease levels in each Quadrant, age, income, gender, race, migrant status, social support system, and social class.

3.3. StateSim

This toolset introduces a model of a state (or cross state or sub-state) region including all the important organizations, the relevant portions of government and its institutional services, economic and security conditions, political processes, domestic practices and external influences. StateSim adds plugins and models atop FactionSim including a population model, economic services models, and the actual institutional agencies that allocate public goods and services (or not) to the factions in that region of interest. StateSim was originally built for three DARPA programs. To date, it has been applied in Afghanistan, Iraq, Palestine, Africa, Sri Lanka, Bangladesh, Thailand, and the Koreas (as well as for UK soccer hooligans and USA crowds) to model (forecast) emergence of state instabilities (insurgency, rebellion, domestic political violence, repression, etc.) which are macro-behaviors that emerge from the micro-decisions of all the participants, including the diverse leaders and followers. In tests run during 2010, DARPA indicated StateSim was better than 80% accurate in over 240 backcast trials they put it through. One can use this model to experiment on and study operations that might influence a region's instabilities and to assess the primary, secondary, and tertiary effects of different courses of action on the stakeholder groups and actors. In an extension for the Army in 2011-12, we scaled up this capability so that StateSim now also includes a module that permits many 10,000s of less-cognitive, follower agents. These exist in networks that carry out the workflows of their daily lives. They also execute the cognitive leader agents' courses of action decisions in a spatial and temporally realistic fashion.

Our goal for SUMHO was to utilize StateSim to represent the behavioral health care services of a community such as Philadelphia, though it also minimally portrays other governmental and non-governmental services (eg., police, courts and justice system, private hospitals, social service agencies, places of employment, and so on) that might impact the lives of the individuals being modeled. Briefly, in terms of Medicaid healthcare services, *Pre-Admission* represents the periods when individuals do not

interact with hospitals or psychiatric resources. Depending on the status of income level and mental stability, individuals will attend regular checkups or take prescribed medicine. *Admission* to a hospital requires physician approval and physicians are accessed in many different ways. Health providers attending patients, policemen, and family or significant others may present patients to a hospital for a physician evaluation. Individuals can also present themselves for a hospital admission. In some cases, individuals seek hospital admission for shelter and meals for the night, or "three hots and a cot," an inefficiency for hospitals. This leads to *Assessment*, or the different ways patients are redirected to receive the care they need. Inpatient *Treatment* or services can be divided into four general categories:

1. ER/Intensive Care – immediate care is required
2. Acute Treatment – urgent, short-term treatment required
3. Sub-acute Treatment –in between acute and chronic treatment; slightly urgent, medium-term treatment required
4. 24-Hour Bed –patient is provided with food, bed and is released following day

The *outpatient* stage is crucial in assessing relapse rates, which are, as noted above, very high in Philadelphia.

Figure 2 – Workflow example

An important capability of this tool is that it captures, represents, and runs the workflow of individual agents in all of the groups and organizations one is interested in modeling. In the current SUMHO prototype, we authored 43 workflows for the residents' daily life activities (e.g., sleep/wake up, get dressed, take medication, do errands, etc.) and for workers in such as those in Hospitals, Homeless Shelters, and Jails. Each of these workflows involve steps that the agent must take, participant roles that must be filled and carried out for the workflow to succeed, various inputs/resources that get consumed (e.g., facilities, space, medicines, supplies, equipment, and so on), and specific outputs or outcome possibilities. As but one example, Figure 2 shows an agent executing the "Take Medicine" workflow. Once an agent decides to Take Medicine (i.e., as its highest utility action choice in earlier Sect 3.1), it must first evaluate whether it has the medicine supply necessary to successfully execute the workflow. If the agent does not have sufficient supply (or believes that it does not), the workflow ends in failure. However, if there is sufficient supply, then the agent takes the medicine, which increases its health and decrements its medicine supply, thus resulting in successfully completing the workflow.

In order to facilitate the interaction of community residents with the baseline healthcare system, we implemented a metric called the Global Assessment of Functioning (GAF) [24]. The GAF is a numerical score (from 0 to 100) used by mental health providers to subjectively rate the psychological, social, and occupational functioning of adults – in other words, how well that individual is able to cope with living. For the purpose of SUMHO, the Physiological and Mental State Properties of an agent are used to calculate an agent's GAF score. This GAF score changes over time as the State Properties of the agent change (both stochastically and in response to various Action Bindings), which will lead the agent to interact with the mental health system in various ways. Figure 5 provides a high-level overview of the simulation environment.

Figure 3 – High level view of patient daily life/health activities for the baseline SUMHO scenario

As Figure 3 suggests, once the GAF drops below a certain threshold, there is a chance an agent will be admitted to the healthcare system. There are numerous ways this could happen. For example, an agent may submit themselves to a clinic, or a police officer may forcibly submit an agent to a clinic if the officer encounters a patient that has become a danger to themselves and those around them. The first case would be highly unlikely for someone with extreme illness, as their paranoia might prevent them from seeking treatment, while the second case might be more likely in such extreme cases. Once this has occurred, an assessment from a physician as well as approval by the managed care team (both of which may be impacted by policies enacted by the model user) determine what kind of inpatient care a patient will receive: intensive, acute, sub-acute, or 24-hour bed. These treatments will persist until the agent's GAF has stabilized at a sufficiently high threshold, after which the agent is discharged from the health system with a particular mix of outpatient treatment. The user of the model is able to specify what types of outpatient care will be available and under what conditions (i.e. GAF scores) each will be applied to an agent being discharged.

4. Results: model validation and analysis of interventions

To develop a three level, mixed-model is a large undertaking. In our experience, this is best approached via a spiral software development where successively better versions are evolved on each spiral (as mentioned in Sect. 1.2). The goal of the first spiral (focus of this paper) should always be to try and encompass the interaction of all the parts of a system, admittedly at a high level. This is meant to be a

first approximation that sacrifices depth, details and precision in exchange for breadth and roughly correct interactions of parts since a system cannot be studied in pieces. It is best to approach this spiral with light MAS agents (as in Sect. 2.1) so one can focus on the component functioning and agent roles in all three system levels. Also, it is useful to begin research on more cognitive agents as we presented in Sect. 3.1 -- the progress on developing those agents was adequate as assessed by domain experts. The assessment of the light agents for spiral 1 is the focus of this section. Specifically, this section begins with model calibration and validity assessment of spiral version 1, a version built for us by Michan et al [25], though calibration and assessment results were obtained by the fourth author of this paper. Once that is presented, SUMHO is then applied to a set of sample policy analyses to illustrate how it can be used to impact population mental health, patient quality of experience, cost effectiveness, and facility usage. The conclusions will return to the lessons learned and questions for subsequent spirals.

4.1. Calibration and model validation

To begin, the SUMHO prototype must be calibrated to examine if key model inputs and outputs for patient behavior and mental health policy match current data of the mental health system in Philadelphia. The calibration process consisted of a three-step process. First, current healthcare practices/workflows, clustering of data into agent archetypes and various distributions of patient state properties were derived via (a) consultation with mental healthcare subject matter experts (SMEs) from our School of Nursing and Medicine (b) from our main model training dataset -- the Collaborative Psychiatric Epidemiology Survey (CPES) [26]. Initial data analysis provided insights into the basic trends on mental health patients and initial distributions for percent of agents flowing over time around the pathways of Figure 5 as a function of health status (GAF score) and social attributes. We also verified parameters using research on the SMI population such as Folsom et al [27], Brown et al [28], Hanrahan et al [29, 30] . Using all this data, we analyzed the overall sample population and then focused on populations with similar disease severities (earlier Quadrants I – IV). It was also helpful to sort agents by disease or type of condition– e.g., manic disorders, depressive disorders, and substance abuse. Decomposing the population into these similar attribute groups allowed the formation of clusters of agents, and thus produced a more heterogeneous population of patients that parallels real data. The inputs also initialize the population sizes and correlated variables such as, among others, how many agents of each type are employed, adhere to lifestyle guidelines (smoking, non-smoking, take drugs), live with family support, are homeless, etc.. These inputs are all added to the patient state properties.

Once all these agent state and input parameters were derived and the daily life workflow model was built, the simulation was then ready to run as a second calibration step. As the simulation runs, each agent makes daily activity choices. In our first version, these are not cognitive agents, but 10,000s of light MAS agents carrying out basic daily activities and flowing around the pathways of Figure 5. Their health related choices are guided by their health status, and the likely distributions of behavior as derived in our research. (ie, probability distributions around each choice), subject to initial conditions and constraints imposed by the inputs and state properties of a given agent. For instance, a low adhering, living alone, depressed agent has lower probability of taking a meds dose than does a high adhering, married, schizophrenic. After an initial run, we corrected for any significant deviations between the model results and the CPES dataset by making incremental adjustments to policy parameters and key patient inputs to make the results from SUMHO as consistent with the training data as possible.

While the same agent will tend to make the same choice under identical conditions, a number of events and their outcomes are random draws in the simulation (e.g., if you try to take your meds, they may have been stolen if you live in a high crime district). Thus the output is stochastic. To account for this non-determinism, SUMHO was next run for 1,000 simulated days in each of 10 trials. The results obtained with each trial of the simulation were slightly different each time. However, when using an average of 10 trials, the variability (standard deviation as % of mean) of each aggregate output parameter we used to judge the validity of SUMHO was below 5%. Only the means of these runs are shown in this article.

As a third calibration step, the averages of these outcomes were compared to the average results of statistics from a second dataset (a separate test dataset) specific to Philadelphia as we will now review. This dataset is from a macro-level study of all mental health incidents and hospitalizations for over 250,000 patients from 2008 to 2010. Specifically, this dataset was used to calculate expected relapse rates, annual hospitalization days per patient, and annual number of admissions per patient. Given the focus of mental health policy on reducing the rate of hospitalization relapse among individuals with severe mental illnesses, important metrics to calibrate would be annual and 100-day relapse rates. As can be seen in Figures 4a - b, the SUMHO prototype is reasonably accurate at replicating the true relapse rates, particularly the annual relapse rate. Student t-tests on these summary metrics were attempted to examine whether deviations from the actual data were significant enough to warrant further model calibration. According to these tests, none of the deviations between SUMHO's results and the actual data were significant at the 5%, 10%, or 15% levels.

Figure 4: Actual versus model-predicted results

Other metrics that were considered included the average number of doctor visits made by patients per year and the average number of days patients stay in the hospital per year (Figure 5). T-tests on both of these metrics produced quite large p-values, which suggest that the deviation between SUMHO's output and the actual data are not statistically significant.

Going forward, the model should be further tested as new data becomes available for 2011 and 2012. Specifically, the results from new policies enacted in recent years (e.g., the Affordable Care Act) will lead to numerous changes. It would be interesting to update and use SUMHO to predict their outcome and then compare that to actual results. However, the impact of these policies will not be reflected in the data of mental health incidents for at least another year. Consequently, the calibration of SUMHO presented here matches "current" mental health policy consistent with policies in place from 2008 through 2010, the time frame of the macro-level dataset.

Figure 5: Statistical t-tests for annual number of doctor visits and days in hospital

4.2. Testing interventions

After successfully assessing the validity of SUMHO's parameters and agent specifications, we experimented with various interventions in order to see their impact on healthcare costs and treatment effectiveness. Surprisingly, we were able to craft a basket of interventions that together resulted in a significant improvement in patient well-being and a dramatic reduction in healthcare costs. Interventions included are discussed in 4.2.1 Summary of Effective Interventions. There were many interventions, however, that were determined to be ineffective at simultaneously improving patient well-being while reducing costs. A few such policies are discussed in 4.2.2 Summary of Ineffective Interventions. The results of our proposed intervention basket are discussed in 4.2.3 Impact of Proposed Interventions.

4.2.1. Summary of effective interventions

Table 3 shows results from the analysis of four policy changes that showed effectiveness.

Reducing the number of patients assigned to an Advanced Practice Nurse (APN): A previous study of the effectiveness of an Advanced Practice Nurses (APNs) intervention providing home visits to

individuals with serious mental illness assigned caseloads of 30 patients [28]. We hypothesized that by reducing the size of these caseloads, nurses may be better able to address the unique needs of each patient and spend more time with them to ensure that they are taking adequate care of themselves, taking their prescribed medications, and attending appropriate therapy sessions. The metric in SUMHO associated with this policy is "Maximum Patients Assigned to APNs".

Assigning more check-ins with physicians following hospitalization: Patients are assigned check-ins with physicians following hospital discharge. Increasing the number of follow-up appointments per month within 30 days of hospital discharge may promote greater stability among patients living on their own without the support of a family unit. The metric in SUMHO associated with policy is "Frequency of Outpatient Check-Ins (per Month)".

Requiring that patients remain in the hospital for addition days once they reach stability: A standard managed care practice to control utilization of hospital days is to suspend funding for inpatient care once the patient has been assessed to have reached sufficient stability and no longer a danger to themselves or others. This assessment is conducted through the use of the GAF score. We hypothesized that the current policy may results in more relapses because temporary emotional stability could be mistaken for true stability. Shorter lengths of stay in hospitals yield patients who remain fragile at discharge and vulnerable to daily stresses of life. Altering this policy by requiring that patients remain hospitalized an additional two days after they reach stability (as determined by a GAF assessment) may reduce the likelihood of relapse and hospital readmission. The metric in SUMHO associated with this policy is "Required Hospitalization Days Above Min GAF".

4.2.2. Summary of ineffective interventions

In stating a compulsory volunteer program for patients with repeated hospitalizations: One hypothesis about why many mental health patients relapse repeatedly is that they lack meaningful ways to contribute to society. To that end, engagement in regular community service could provide a sense of purpose and order. After experimenting with such a policy, it was found to have a minimal impact on patients with moderate condition severity while leading to substantially more stress for patients with severe conditions or comorbidities. Metrics in SUMHO associated with this policy are "Instate Compulsory Volunteer Program?" and "Min Annual Hospitalizations for Volunteer Program".

Providing cellular phones to patients without access to a home phone number: Across the United States, nearly 7 % suffer from a diagnosed severe mental illness and a prevailing problem is difficulty

maintaining communication with these individuals. Case managers and APNs must keep track of a patient's whereabouts and visit them directly for a check-up. To ensure communication, it has been suggested that such patients be provided with a cellular phone so that healthcare practitioners can call patients. After implementing this intervention in SUMHO, it was found to be ineffective and costly for numerous reasons. Firstly, homeless patients without a source of stable income except welfare payments often decided to sell their phones or trade them. Secondly, even among those who opted to keep their phone, patient groups at the greatest risk for breakdown are also groups that tend to cooperate with mental health professionals the least – they often avoid answering the door for visiting APNs or cooperating with prescribed therapy schedules. Providing these individuals with a cell phone did not address this key behavioral characteristic. Metrics in SUMHO associated with this intervention are "Provision of Cellular Phones?" and "Min Annual Hospitalizations for Provision of Cell Phone".

   Setting a maximum number of hospital admissions per year: A criticism about the current institutional focus of the mental healthcare system is that it builds a system of dependency – patients become unable to deal with the stresses of daily life on their own because they know that they can rely on hospitalization and institutionalization to provide them with food, therapy, and a place to sleep. Further, many homeless individuals have been found to fake mental illness and other health problems in order to be admitted to a hospital in order to receive a "cot" and a "hot meal". By setting a maximum number of times a patient can be permitted to stay overnight in the hospital, costs within the health system were drastically reduced within SUMHO as individuals were not able to develop a dependency on the health system for the provision of their basic necessities for life; however, this was accomplished at the cost of patient well-being. Specifically, the average GAF score for patients with severe schizophrenia and complex comorbidities dropped dramatically from an average of about 40 to less than 30. Therefore, we do not recommend this as a valid approach. The metric in SUMHO associated with this policy is "Max Admissions per Year".

4.2.3. Impact of proposed policies

Table 3 highlights the policy changes we propose based on the results of SUMHO. In order to assess the impact of our proposed policies, we looked at many key metrics for patient well-being and healthcare system costs. As a summary, below are figures generated by SUMHO for several of these metrics. For patient well-being, we considered the average GAF score for mentally ill agents and the 100 day relapse rate. For cost effectiveness, we considered the average daily cost per agent and the availability of hospital beds for mentally ill patients. As the results indicate, the above basket of policies improves patient well-

being while reducing healthcare system costs and preserving capacity within the healthcare system (Figures 6a-d).

Table 3: Preliminary policy recommendations

Figure 6: The new policy improves metrics for quality of life (a & b), cost (c), and capacity (d)

5. Discussion and lessons learned to date

This paper started out by exploring the usefulness of three levels of models of the health of a community – the individual population members, the healthcare "system" including the various practices and services that comprise its "parts", and the overall "containing society". The first and last are critical since the decision making of individuals in the population and societal determinants are the primary drivers of living well. We have successfully applied this 3-level systems modeling approach for the US Department of Defense and State to analyze societal instabilities overseas and how to influence them for the better. The current paper explored whether this could be successfully transported to modeling and simulating mental health of an urban community. A case study was attempted involving the Philadelphia Medicaid population seeking mental health services – roughly 100,000 members. Section 3 explained a first approximation prototype of the agent models where they are guided by their overall health status score and population distribution functions relevant for their respective state properties (disease/physiologic, mental, and socio-economic). This successfully demonstrated the feasibility of quantifying all the organizational, group, and individual agent parameters.

Despite the limits of the first SUMHO prototype, Section 4 demonstrated that it has substantial validity relative to population behavior and mental health policies and practices in Philadelphia in the 2008 to 2010 timeframe. The prototype was able to recreate several summary statistics (e.g., re-hospitalizations, 100 day relapse rate, number of doctor visits, and days in hospital) with no statistical difference from the averages of the population drawn from a sample of 250,000 patients. Finally, the SUMHO prototype was applied to identify policies that are effective as well as those that are ineffective. Four categories of effective policies and 3 categories of ineffective ones were found by the use of the model. The basket of effective policies were projected forward in time and the results indicate the policies improve patient well-being while reducing healthcare system costs and preserving capacity within the healthcare system. This is the overall result that one ideally wants from a systems modeling effort. We thus can posit that there is utility in the 3

levels of system modeling approach, at least for the case study behavioral healthcare system of Philadelphia.

The agent framework we are using readily supports fuller modeling of patient physiology, cognition, and socio-economic determinants, as well as the workflows and caregiving of providers. Adding depth into these dimensions will allow fuller assessment of interventions and their impacts. In our earlier military research, we started with simpler agent models and added model depth and causal richness as we proceeded to later spirals or versions. Further research is thus warranted to add flesh to the component models. As one example, the underlying PMFserv agent framework allows one to replace the physiologic state properties with a model of physiologic dynamics. Also, we hope to extend agent cognition to the point where they use individualized value trees (goals, standards, and preferences, or GSPs) as described in the PMFserv discussion of Sect 3.1. Indeed, this is consistent with the theory of reasoned action [31] which presumes a person's behavioral intention is based on his or her attitudes about that particular behavior and the current subjective norms based on societal influence. PMFserv supports quantifying of attitudes via the GSP approach, though we are interested in adopting more formal instruments for profiling these in the future. As to the social norms, the current SUMHO prototype clusters agents into heterogeneous categories based on disease type, severity quadrants, and social property levels (family support, neighborhood, lifestyle). The underlying agent models exist for dynamically modeling activities of daily life. Adding all these models to a future version of SUMHO would more fully inform the analysis and be more useful for assessing the impact of interventions.

It is worth closing by comparing our results found via agent-based modeling (ABM approach) to what could have been discovered with a classical regression approach. Social determinants are necessarily prominent features when configuring best practice and optimal health outcomes for individuals with chronic and complex illnesses. For example, people with serious mental illness (SMI) have complex mental and medical conditions that are deeply-rooted in a health, socioeconomic, and environmental context [3]. As our survey of the literature indicated, regression models are the prevailing approach to answering research questions. Data mining of patient record systems to find interesting patterns and form trend line predictions is often utilized. These predictions tend to rely on some form of regression equation. However, regression models lack the deep dive into social drivers of health outcomes that are intrinsic to ABM methods.

To illustrate, we examine results from a randomized control trial (RCT) that used a sample of individuals with serious mental illness [29]. The purpose of the RCT was to examine the effect of a nurse intervention on lowering readmission rates. Counter to hypothesis, the group that received the nurse

intervention had significantly higher readmission rates than the control group. At face value, the RCT finding shows the nurse intervention is ineffective. However, another plausible explanation is that the intervention group received attention from a skilled nurse who was available to them 24/7 and assessed the need for additional health services. Usual care of SMI does not include such attention.

In contrast, the ABM study reported in this paper used a comparable sample and found that lowering the number of patient cases followed by the nurse could lower patient readmissions and reduce costs. Compared with the RCT, the ABM results provide health administrators guidance about resources management and accountability for improving quality and reducing costs. Additionally, ABM methods account for the influence of psychological, social, economic, and environmental factors that are dynamic features of problems associated with the SMI population. In short, ABM provides a different perspective from regression models and different research questions can be posed and answered.

Specifically, one could certainly use regression to find effective and in-effective interventions if one had completely filled in patient records and socio-economic data sets. However, the MAS approach has illustrated the ability to generate missing patient data based on statistical inferences about archetypes and to use that data in monte carlo style experiments to project likely behaviors and impacts of interventions. Spiral 1 of our research has thus demonstrated capabilities not possible with classical regression alone, provided one accepts the calibration and validity analysis. Even more significantly, once a research group satisfies itself with the validity of this approach for modeling a given population and set of interventions, the MAS opens the door to a new paradigm with no parallel in the regression world. Specifically, since one can drill down into the agents mental processes, one can determine what caused their behaviors. One can look inside them to learn about their decision making and emergent behavior rationale. This is the reason for all the excitement about the MAS approach.

[1] Carpenter, S. R., Folke, C., Norström, A., Olsson, O., Schultz, L., Agarwal, B., et al. Program on ecosystem change and society: an international research strategy for integrated social–ecological systems. *Current Opinion in Environmental Sustainability*, 4, 1 (2// 2012), 134-138.

[2] Christiansen, J. H. and Campbell, A. P. *HealthSim: An agent-based model of simulating health care delivery*. City, 2003.

[3] Bauer, A. M., Thielke, S. M., Katon, W., Unutzer, J. and Arean, P. Aligning health information technologies with effective service delivery models to improve chronic disease care. *Preventive Medicine* (Jun 22 2014).

[4] Bharathy, G. K. and Silverman, B. *Applications of Social Systems Modeling to Political Risk Management*. Springer, City, 2012.

[5] Silverman, B., Bharathy, G. and Kim, G. The new frontier of agent-based modeling and simulation of social systems with country databases, newsfeeds, and expert surveys. *Agents, simulation and applications. Taylor and Francis*2009).

[6] Silverman, B. G. Systems social science: a design inquiry approach for stabilization and reconstruction of social systems. *Journal of Intelligent Decision Technologies*, 4, 1 2010), 51-74.

[7] Rothbard, A., Chhatre, S., Zubritsky, C., Fortuna, K., Dettwyler, S., Henry, R., et al. Effectiveness of a High End Users Program for Persons with Psychiatric Disorders. *Community Mental Health Journal*, 48, 5 (2012/10/01 2012), 598-603.

[8] Jencks, S. F. Defragmenting care. *Ann Intern Med*, 153, 11 (Dec 7 2010), 757-758.

[9] Colton, C. W. and Mandersheid, R. W. *Congruencies in increased mortality rates, years of potential life lost, and causes of death among public mental health clients in eight states*. Center for Disease Control, City, 2006.

[10] Axtell, R. and Epstein, J. M. Growing artificial societies: social science from the bottom up. *Brookings Inst. Press: Washington, DC [ua]*1996).

[11] Parsons, S. and Wooldridge, M. Game Theory and Decision Theory in Multi-Agent Systems. *Autonomous Agents and Multi-Agent Systems*, 5, 3 (2002/09/01 2002), 243-254.

[12] Brailsford, M., Clark, D. and Gissing, R. Simulation brought to life. *Nurs Stand*, 28, 1 (Sep 4-10 2013), 64.

[13] Schlessinger, L. and Eddy, D. M. Archimedes: a new model for simulating health care systems—the mathematical formulation. *Journal of biomedical informatics*, 35, 1 2002), 37-50.

[14] Newell, A. Unified theories of cognition. *Cambridge, MA: Harvard University*1990).

[15] Anderson, J. R., Matessa, M. and Lebiere, C. ACT-R: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction*, 12, 4 1997), 439-462.

[16] Cooper, R. and Shallice, T. Soar and the case for unified theories of cognition. *Cognition*, 55, 2 1995), 115-149.

[17] Manoj, B., Rao, R. R. and Zorzi, M. CogNet: a cognitive complete knowledge network system. *Wireless Communications, IEEE*, 15, 6 2008), 81-88.

[18] Gluck, K. A. and Pew, R. W. *Modeling human behavior with integrated cognitive architectures: Comparison, evaluation, and validation*. Psychology Press, 2005.

[19] Zacharias, G. L., MacMillan, J. and Van Hemel, S. B. *Behavioral Modeling and Simulation: From Individuals to Society*. Committee on Organizational Modeling from Individuals to Societies, National Research Council of the National Academies. Washington, DC: National Academies Press, City, 2008.

[20] Silverman, B. G., Pietrocola, D., Nye, B., Weyer, N., Osin, O., Johnson, D., et al. Rich socio-cognitive agents for immersive training environments: case of NonKin Village. *Autonomous Agents and Multi-Agent Systems*, 24, 2 2012), 312-343.

[21] Ortony, A. The cognitive structure of emotions. *Journal of Intelligent Decision Technologies*, 4, 1 1990), 51-74.

[22] Bartels, S. J. Caring for the Whole Person: Integrated Health Care for Older Adults with Severe Mental Illness and Medical Comorbidity. *Journal of the American Geriatrics Society*, 522004), S249-S257.

[23] Druss, B. G., Zhao, L., Cummings, J. R., Shim, R. S., Rust, G. S. and Marcus, S. C. Mental comorbidity and quality of diabetes care under Medicaid: a 50-state analysis. *Med Care*, 50, 5 (May 2012), 428-433.

[24] Jones, S. H., Thornicroft, G., Coffey, M. and Dunn, G. A brief mental health outcome scale-reliability and validity of the Global Assessment of Functioning (GAF). *The British Journal of Psychiatry*, 166, 5 1995), 654-659.

[25] Michan, Z., Warner, J. and Zhang, L. mhsSIM: Mental Health System Simulator. An Agent-Based approach to modeling mental health policies & interventions2012).

[26] Alegria, M., Jackson, J. S., Kessler, R. C. and Takeuchi, D. *Collaborative Psychiatric Epidemiology Surveys (CPES), 2001-2003 [United States]*. Inter-university Consortium for Political and Social Research (ICPSR) [distributor], City, 2009.

[27] Folsom, D. P., Hawthorne, W., Lindamer, L., Gilmer, T., Bailey, A., Golshan, S., et al. Prevalence and risk factors for homelessness and utilization of mental health services among 10,340 patients with serious mental illness in a large public mental health system. *Am J Psychiatry*, 162, 2 (Feb 2005), 370-376.

[28] Brown, R. S., Peikes, D., Peterson, G., Schore, J. and Razafindrakoto, C. M. Six Features Of Medicare Coordinated Care Demonstration Programs That Cut Hospital Admissions Of High-Risk Patients. *Health Affairs*, 31, 6 (June 1, 2012 2012), 1156-1166.

[29] Hanrahan, N. P., Wu, E., Marcotte, J., Kelly, D., Ma, C., Aiken, L., et al. Randomized clinical trial of the effectiveness of home-based advanced practice nursing on outcomes for individuals with serious mental illness and HIV. *Journal of Nursing Research*2011).

[30] Solomon, P., Hanrahan, N. P., Hurford, M., DeCesaris, M. and Josey, L. Lessons Learned from Implementing a Pilot RCT of Transitional Care Model for Individuals with Serious Mental Illness. *Archives of Psychiatric Nursing*2014).

[31] Blank, M. B., Hanrahan, N. P., Fishbein, M., Wu, E. S., Tennille, J. A., Ten Have, T. R., et al. A randomized trial of a nursing intervention for HIV disease management among persons with serious mental illness. *Psychiatr Serv*, 62, 11 (Nov 2011), 1318-1324.