QUALITY MANAGEMENT JOURNAL
2023, VOL. 30, NO. 3, 168–186
https://doi.org/10.1080/10686967.2023.2211287

Uncertainty, emergence and adaptation: A complex adaptive systems approach to quality improvement

William Wilson[a], Scott McLachlan[b,d], Kudakwashe Dube[c,d,*], Kathleen Potter[e] and Nihal Jayamaha[a]

[a]School of Food and Advanced Technology, Massey University, Palmerston North, New Zealand; [b]School of Nursing, Midwifery and Palliative Care, King's College London, London, UK; [c]School of Fundamental Sciences, Massey University, Palmerston North, New Zealand; [d]Health informatics and Knowledge Engineering Research (HiKER) Group, London, UK; [e]Ryman Healthcare, Christchurch, New Zealand

ABSTRACT

The healthcare quality improvement (QI) literature is replete with examples stating that continued failure to regard healthcare as a complex adaptive system (CAS) reduces the effectiveness of quality improvement initiatives. Recommendations and strategies for managing change within CAS exist, but the specific mechanisms that bring about successful change within CAS and the implications for quality practitioners are under-explored. This article presents a generalizable model for QI within CAS and provides a specifically CAS explanation for incremental change. We develop a conceptual model from foundational CAS principles that is then operationalized as an agent-based simulation model. Our model captures critical complex system behavior in a generic manner easily applied to different improvement contexts. We tested the simulation model using a recognizably complex healthcare improvement case: reducing antipsychotic prescribing levels in aged residential care. Nonlinear phase transitions were observed in the agent network, conditioned on the network's ability to learn solution options and simultaneously maintain cooperation. We believe that a CAS explanation of change can assist practitioners navigating complex QI activities.

ARTICLE HISTORY
Received 28 November 2022
Accepted 15 February 2023

KEYWORDS
agent-based model; complex adaptive systems; healthcare improvement; improvement networks; problem-solving strategies

Introduction

Quality improvement (QI) is a fundamental and widely used approach to improving healthcare delivery. Despite significant investment, research, and effort over decades, sustainable, transformational change of healthcare systems is elusive (Braithwaite, Glasziou, and Westbrook 2020; Dixon-Woods and Martin 2016). One reason often cited for the limited effectiveness of QI initiatives in healthcare organizations is the failure to regard healthcare settings as *complex adaptive systems* (CAS) (Braithwaite et al. 2018; Institute of Medicine 2001; Khan et al. 2018). The implication from advocates for CAS thinking is that successful QI methods, developed and refined in business and manufacturing contexts, have inbuilt assumptions of *information availability*, *linearity*, and *system equilibrium* that might not apply to CAS (Braithwaite 2018; Khan et al. 2018, 3–6; Young et al. 2021). This is problematic if our QI methods assume information availability and stability, for example, when using a control chart. A further difficulty is the degree of interconnection between system elements within a CAS. Improvement approaches relying on isolating the individual parts of a system are considered unlikely to be successful within a densely interconnected complex system (Preiser et al. 2018). Taken together, the lack of certainty and difficulty locating precise intervention points are serious obstacles to QI.

Examining these claims critically raises two fundamental questions. *First*: is it valid to consider healthcare QI as a form of CAS? *Second*: if the characterization of healthcare QI as a CAS is valid, what do QI methods built on CAS principles look like? Regarding the first question, a large body of literature supporting the idea that healthcare QI is a form of CAS already exists, for example: Braithwaite (2018), Covvey (2018), Khan et al.

CONTACT: William Wilson bill.wilson@cdhb.health.nz School of Food and Advanced Technology, Massey University, Palmerston North, New Zealand.
*The current affiliation for Kudakwashe Dube is King's College London, UK.

(2018), Rouse (2008), and Thompson et al. (2016). A critical examination of the rationale and evidence base for this assertion is beyond the scope of this article. We take the position that the application of CAS principles to complex sociotechnical systems is valid and largely uncontroversial within the relevant discourse. Importantly, we also agree with John Paley's (2010) assertion that *specifically CAS explanations of healthcare QI* are still required. CAS explanations will help to answer the second general question: what do QI methods built on CAS principles look like? Specifically, the research question underpinning this study asks: *can complex QI be explained exclusively in terms of CAS phenomena?*

Focusing on a limited set of agent rules in accordance with CAS theory seems inconsistent with what healthcare QI practitioners perceive as their personal and professional agency. Paley requests explanations for outcomes that account for system "design" and describe agent behavior beyond preprogrammed responses to external stimuli (Paley 2010; Paley and Eva 2011). To move from natural ecosystems to human organizations requires plausible *mechanisms* equivalent to the biological adaptation processes in natural CAS, in short, an explanation that fits with our sense that human agency is more than preprogrammed rules and that our "systems" usually have a designed purpose. At the same time, a CAS explanation of QI needs to show that desired macro-level system outcomes can really be achieved via a small set of internalized agent behavioral rules (Paley and Eva 2011; Holland 2006, 2012).

This study presents a conceptual model for explaining QI within CAS using John Holland's foundational CAS theories as the starting point (Holland 2006, 2012). To test our model, it is first operationalized as an agent-based simulation. We then investigate face validity via a case study: *a multidisciplinary team seeking to reduce the overprescribing of antipsychotic medication in aged residential care (ARC) facilities*. This scenario has a strong prima facie claim to be genuinely complex in terms of both the problem being worked on and the characterization as a CAS. The case study results are used to explore applicability of the model and its potential as both a practical tool for improvement teams working in complex settings and as a novel approach for quality researchers investigating CAS phenomena.

Consistent with the CAS literature, we believe our simulation model will reveal system leverage points that could be influenced to achieve change in healthcare systems involving human agents (Holland 2012; Khan et al. 2018). We also anticipate that phase transitions (self-organized criticality) within the system may be observed, indicating that complex improvement projects may under certain conditions (determined by their unique attributes) transition from a "high probability of failure" state to a "high probability of success" state. We propose that development of a working CAS explanation for QI, together with the ability to identify and observe phase transitions within complex QI projects, will be a useful contribution to QI strategies in CAS environments.

The remainder of the paper proceeds as follows: The background contains a brief literature review covering the characterization of healthcare as a CAS, including the research gap between CAS and system design. The theoretical foundations for the conceptual model are then set out. The methods section describes how the conceptual model is developed and operationalized as an agent-based simulation model. Having fully described the agent-based model, the case study used for assessing face validity is introduced. Following presentation of the results, we discuss the findings and the implications for quality management, practice, and future research.

Background

Healthcare as a CAS

Complexity science has multiple originating intellectual traditions (Castellani and Gerrits 2021; Mitchell 2009). From within this broad range of complexity thinking, CAS ideas and language have resonated strongly with healthcare researchers (Thompson et al. 2016). A system can be classified as a CAS when some combination of the following criteria is present: (i) a large network of autonomous agents; (ii) dynamic, nonlinear interactions; (iii) a high rate of change; (iv) emergent higher-order effects; (v) the ability to adapt and learn; (vi) self-organization; (vii) coevolution; (viii) temporality, and (ix) system history (Chandler et al. 2016; Dooley 1997; Holland 2006, 2012). Conceptualizing healthcare as a complex system is often interpreted as *a rejection of reductionist solutions* that ignore interconnected system relationships and the use of deterministic, "machine" metaphors for the system (Young et al. 2021). The implication is that something very different is taking place within CAS.

In a CAS model of healthcare QI, system boundaries are no longer rigid and interactions between agents in the system are entangled and unpredictable. Agents apply internalized rules and self-organize their collective behavior. Successful change may require coevolution with connected agent networks. Outcomes of interest within a healthcare QI CAS are emergent, difficult to predict, and nonlinear. Notwithstanding these difficulties, meaningful patterns of change may be observable, even if they do not yield to precise analysis. Plsek and Greenhalgh (2001, 627) go so far as stating that: "*Ultimately, the only way to know exactly what a complex system will do is to observe it.*"

References to complexity thinking in healthcare research are now commonplace (Braithwaite et al. 2018; Chandler et al. 2016; Khan et al. 2018; Lorden et al. 2014; Turner and Baker 2019). Quality practitioners can readily access frameworks, strategies, and methods for better managing change within complexity. Table 1 lists six well-known and mature examples, noting the high-level emphasis of each framework.

Table 1 is not an exhaustive summary of all available frameworks; inclusion is based on maturity and active use supporting contemporary QI research and practice (Anderson et al. 2015; Augustsson, Churruca, and Braithwaite 2020; Bäcklander 2019; Kringos et al. 2015; Perla, Provost, and Parry 2013; Snowden and Rancati 2021). There are thematic commonalities to the list in Table 1, notably the emphasis on situational awareness and specific leadership/management behaviors. QI practitioners interested in working with complexity principles seem to have several options for guidance, and more variations along these same themes would be difficult to justify.

Despite the availability of complexity-informed tools, concerns remain over the effectiveness of healthcare QI under complex conditions (Dixon-Woods and Martin 2016). We are not suggesting that the current strategies or frameworks are ineffective, but we do believe that foundational CAS thinking still has more to offer. Given the continued pressure on healthcare systems worldwide and the potential impact of effective QI, it seems important to explore all options thoroughly. Carmichael and Hadžikadić (2019) suggest that systems thinking is vulnerable to hierarchical and complexity biases, that is, we tend to introduce hierarchy and complexity into our systems thinking whether or not this is justified. In contrast, explanations of natural CAS behavior are usually associated with a small set of agent rules and avoid hierarchical agent structure (Holland 2006). The frameworks summarized in Table 1 access CAS principles *indirectly*

Table 1. Frameworks for complexity-informed quality improvement.

Framework	Short description	Thematic emphasis
Soft Systems Methodology (Checkland 2000)	An approach to analyzing and responding to complex problems. *Soft* systems assume the presence of complex interactions and multiple valid perspectives of problems and solutions, in contrast to clear or stable (*hard*) systems. A cyclic learning "system" is used to propose, explore and test potential solutions from multiple perspectives.	Learning cycles Multiple stakeholder perspectives
System of Profound Knowledge (Deming 2000)	Contains four core elements: • Appreciation for a system (knowledge of interdependent parts and their interactions). • Theory of knowledge (what we know and how we can acquire new knowledge). • Psychology (understanding behavior). • Understanding variation (common and special cause variation).	Learning cycles Interdependence Human behavior
Cynefin (Kurtz and Snowden 2003)	A sense-making, situational awareness and decision support framework. Cynefin proposes that systems dynamically move between ordered and unordered domains, further broken down into *clear, complicated, complex*, and *chaotic* system states (or problem scenarios). Leadership, decision-making, and intervention actions differ based on the current system state and the level of knowledge, i.e., optimal management responses differ across different system states.	Situational awareness Sense-making Leadership skills and behavior Decision making
Complexity Leadership Theory (Uhl-Bien, Marion, and McKelvey 2007)	Leadership skills suited to knowledge-oriented activity rather than production focused. *Adaptive* outcomes: learning, innovation, adaptability. *Entangled* leadership roles: adaptive, administrative, and enabling leadership functions. Emphasizes the dynamic tension between administrative, bureaucratic functions and emergent, informal interactions.	Situational awareness Leadership skills and behavior Decision making
Adaptive Leadership (Heifetz et al. 2009)	Proposes leadership requirements based on situational complexity. Adaptive leadership fosters skills required to navigate *technical* and *adaptive* scenarios as appropriate. Simple problems require technical responses, complicated problems require technical and adaptive responses, and complex problems require adaptive responses. Balance between top-down, standardized procedure responses to clearly defined issues and bottom-up, staff-led learning to understand complex issues and build solutions.	Situational awareness Leadership skills and behavior Decision making
Model for Understanding Success in Quality (MUSIQ) (Kaplan et al. 2012)	Contextual factors model with recursive macro, meso, and micro levels. An assessment tool to examine enabling competencies for change. Encompasses leadership, quality management teams, and clinical microsystem teams that support and implement change.	Context and enabling factors Leadership hierarchy

through the filters of our sociotechnical norms such as leadership and decision-making. *Directly* harnessing CAS phenomena, free from hierarchical and complexity biases, is a valid and still under-explored alternative strategy (Carmichael and Hadžikadić 2019).

Skepticism toward the applicability of CAS principles within human sociotechnical contexts is rare in the literature but does exist and raises valid concerns. Yawson (2013) describes the wide gap that exists between theoretical papers and practice within human resource development research. A prerequisite of teaching the background systems theory is missing and there are few attempts to purposefully apply complexity theories as an alternative to linear frameworks (Yawson 2013). Within the healthcare as CAS discourse, questions over the role of human cognition and system design within CAS have been posited (Paley 2010; Paley and Eva 2011). We regard answering these questions of mechanism as an important research gap to be addressed.

Research gap

In the preceding sections, we have noted the general concerns within the literature as to the effectiveness of healthcare QI. We have also suggested that despite the availability of various complexity-informed solutions, under-explored avenues of research remain. Given the potential benefits of improved healthcare QI, we are motivated to explore in detail the notion of a CAS mechanism for change in systems where human agency is also present. In sociotechnical settings, the challenge of constructing an integrated framework for change is complicated by human agency and cognition. Paley and Eva (2011) contend that cognition in human agents, and the role of cognition in self-organization, results in a completely different process than for system agents who are not self-aware. For Paley and Eva, complex goal-oriented activity cannot literally be a CAS because CAS by definition lack a system-level design and agent intentionality (Paley and Eva 2011). The point is somewhat technical, relating to the validity of explanations for how CAS function, and Paley and Eva do not dismiss the usefulness of CAS as an analytical lens. Articulation of a CAS explanation (mechanism) for healthcare QI is still to be achieved. Khan et al. (2018) confirm that this gap remains and emphasize the need for ongoing research. Consistent with the general argument for better theoretical explanations of CAS, Greenhalgh and Papoutsi (2018) support the call for new complexity-informed research paradigms.

Deriving theory for a model of QI within CAS from reference literature

Recognizing both the intuitive appeal of interpreting complex social organizations as CAS and the valid question of system-level design or purpose, we propose that the QI within CAS scenario is best interpreted as a special case of CAS. At given moments in time, CAS agents involved in healthcare QI do coordinate their actions toward a goal, and that goal can be defined and understood at a level greater than one individual agent. In these circumstances, the "CAS" is not the nominal project team, organization, or population group that first comes to mind. The CAS is a *specific collection of agents within these entities, cooperating on a specific goal.* This definition fits with the CAS category distinction recently introduced by Wilson and Kirman (2016). In CAS_1, adaptation takes place at the system level, for example, the human immune system. In CAS_2, complex networks of agents, for example, multi-species ecosystems, respond adaptively to change and system-level change emerges from this agent-level activity (Wilson and Kirman 2016). Using this conceptualization of healthcare QI, the CAS in question becomes simply one of many interconnected agent networks within a larger sociotechnical system such as a hospital. The function of adjacent and interlinked networks also needs to be adequately accounted for in any explanation of a given network's behavior.

Theories of CAS exist; our research applies one of these theories and demonstrates that appropriate CAS thinking can be meaningfully applied to self-aware human agents. Explanations that demonstrate CAS applicability in human sociotechnical contexts address Paley and Eva's concerns over mechanism. More importantly, CAS principles have significant implications for QI practice and management. CAS theory predicts that system-level change emerges from the localized agent interactions where the agent behavior conforms to a small number of internalized rules (Holland 2012; Kauffman 2000: Mitchell 2009). Alongside any "conventional" improvement actions taking place, understanding the agent adaptation process in a QI setting thus becomes a necessary condition for successful change to occur.

Perhaps the clearest formal articulation of CAS behavior has been provided by John Holland (1992, 2006, 2012). Holland defines four essential CAS attributes: (i) *parallelism*, where large numbers of agents interact by sending and receiving signals; (ii) *conditional action*, where agent actions are conditional on signals and rules available to the agent; (iii) *modularity*, where activities can be combined to form larger "subroutines," and (iv) *adaptation and evolution*, where the agents change over time, via adaptations that improve a situation (rather than random change). Holland (2012) proposes that adaptation requires *credit assignment* (weighting and rewarding individual signals and rules within the parallel activity) and *rule discovery* (when the agent's existing rules are incomplete or inadequate). Working in the field of computational biology, Holland specified a formal algorithmic process, a *classifier* system, to simulate this detailed agent behavior of signal detection, signal evaluation, rule evaluation, rule generation, and rule selection (Holland 2006, 2012; Urbanowicz and Moore 2009). Satisfactorily expanding or replacing this classifier "adaptive" kernel presents the biggest challenge in proposing that CAS principles are appropriate for nuanced, sophisticated human behavior.

Conceptualizing improvement as a problem-solving process provides a bridge between rule-based agent behavior within a CAS and the goal-oriented activity of human agents. Adaptation directed toward improvement implies problem-solving to some degree, which may range from trivial to extremely difficult. Formal problem-solving strategies, such as *means-end analysis* or the *goals-assumptions-elements-operators* model (Anderson 2005; Cronin and Weingart 2007), offer a practical way of introducing purpose, awareness, and human agency into the agent adaptation process.

Plausible agent mechanisms are required to replace the biological processes that occur within natural (non-self-aware) CAS, that is, inherited and learned behavioral rules, variation, genetic mutation, and natural selection (Holland 2006, 2012; Kauffman 2000). Problem-solving processes can deliver this learning and adaptation function (Anderson 2005; Doolittle 2014; Newell and Simon 1972). Cooperative problem-solving establishes, tests, and evolves the mental models or *schema* of participating agents (Cronin and Weingart 2007; Dooley 1997; Mohammed, Ferzandi, and Hamilton 2010). Problem-solving requires knowledge which, if not initially present in the agent network, must then be acquired, that is, learned. A learning gain from each learning activity is therefore posited as critical to acquiring sufficient knowledge within the agent network (Anderson 2005; Deming 1986; Doolittle 2014; Reed and Card 2016).

Completing the generic QI as CAS model requires two further elements. The problem-solving actions of agents are attempts at system regulation and control. To regulate a system, agents need feedback signals and paths from their environment to understand and respond to the continually changing system status (Ashby 1961; Umpleby 2008). Finally, CAS are open systems without fixed boundaries (Holland 2006; Nair and Reed-Tsochas 2019; Umpleby 2008) and some form of network boundary delineation is required. The network will evolve as agents join and leave and implementing change may also require other agent networks to adapt if change is to occur (coevolution with other agent networks). Where the CAS stops and the external environment begins is not fixed. Nair and Reed-Tsochas (2019, 86) describe this permeable boundary as the *enacted and interpreted boundary* within a CAS, dynamically linked to the external environment.

The minimum requirements for a model explaining QI within CAS can now be specified. A system capable of carrying out deliberate improvement actions requires (i) the circulation of system information (signals) among a network of agents; (ii) some level of shared interpretation of issues, goals, and assumptions; (iii) problem-solving capability and resources adequately matched to the problem and wider context; and (iv) an ability to implement changes within the defined system as well as closely related systems (agent networks) that may be impacted (Cronin and Weingart 2007; Holland 2012; Newell and Simon 1972). Having outlined the theoretical foundations for our model, we now describe the *complex quality improvement network* (CQIN) model of QI within CAS.

Approach and method

Bringing together the foundations from the preceding sections, the theory represented by the CQIN model is as follows: *within CAS, intentional system-level change, in the form of quality improvement, is achieved through coordinated changes to agent schema.* This theory, expressed as a testable hypothesis for this initial study, becomes: *(H1) Updating and aligning agent schema takes place fully within the constraints of the agent behavioral rules.* Evidence to support this hypothesis would confirm that agent behavioral rules for cooperation, learning, and updating internal schema, rather than being interpreted as simplifying or diminishing human agent activity, can instead be treated as the CAS leverage points (Holland 2012; Wilson and Kirman 2016). In testing the theory, this study follows the approach recommended by Davis, Eisenhardt, and Bingham (2007) for *theory building via simulation methods*.

Constructs and definitions for the CQIN conceptual model

The CQIN conceptual model of a generic problem-solving network is presented in Figure 1. A network of agents (A) cooperates in deliberate problem-solving activity (G-B), receiving and interpreting system feedback (B-E-F and C-D-F). The agent interactions are moderated by the complexity of the setting (H) and exogenous factors from the external environment (I). The model elements shown in Figure 1 are conceptual constructs, that is, abstract variables and relationships we are not able to directly observe (Meredith 1993). Definitions for the CQIN constructs are contained in Table 2.

CQIN contains only those constructs and relationships required for problem-solving across an agent network.

Assessing the CQIN conceptual model as theory

Satisfactory theory requires (i) definitions of the theory elements and constructs; (ii) identified domain boundaries for applicability; (iii) propositions for the relationships/associations between constructs, and (iv) the ability to make predictions or provide explanations (Wacker 2008; Weber 2012). Definitions for the CQIN model constructs and their interactions have been introduced. The boundaries of the model and the applicable domain for the theory (QI within CAS) have also been specified. In the following sections, we present the detailed associations between the variables and offer explanations for QI within CAS system behavior based on a simulation case study.

Agent-based simulation model

Empirical measurement of CAS is challenging. Simulation modeling offers a feasible starting point for data generation and system observation when data collection is difficult or impossible. *Agent-based modeling* (ABM) is particularly suited to complex systems involving autonomous agents (Holland 2006; Preiser et al. 2018). The agent substrate is a critical requirement for studying CAS because it allows the complexity to be approached from the granular level of the agent activity, rather than attempting to faithfully define the full scope and detail of very large systems "from the outside in." Research consciously applying ABM to quality improvement is not yet common but is increasing. Ornstein et al. (2020) used the theory of "rugged landscapes" from evolutionary biology to model how different types of public health implementation efforts can be suboptimal when not matched to the applicable complexity scenario. Seid et al. (2021) developed an ABM model of a collaborative learning health system and modeled how patients and clinicians influence engagement levels and the overall level of knowledge available within the system. Both these studies found that simulation was particularly

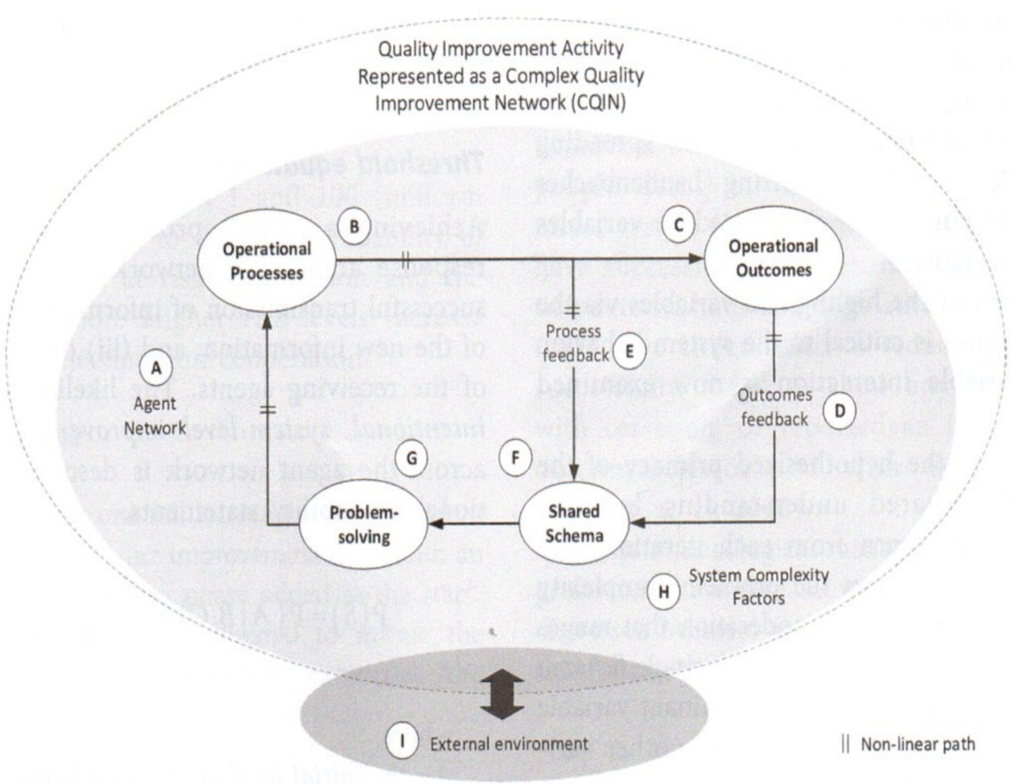

Figure 1. Generic complex quality improvement network.

Table 2. CQIN conceptual constructs.

CQIN construct	Definition
Agent network (A)	The agent network is made up of all the participants and affected parties of the QI activity. For healthcare QI activity, the agent network encompasses clinicians, patients, QI practitioners, and supporting functions, within the interpreted environmental boundaries.
Operational processes (B)	Operational processes and activity as relevant to the improvement context, for example, processes and applicable clinical protocols in the healthcare context.
Operational outcomes (C)	Defined, measurable process and system outcomes.
System signals (D), (E)	Feedback on the health outcomes (e.g., changes in a specified performance measure) as well as wider system feedback (e.g., changes in the environment, resources, process timeliness, or costs).
Shared schema (F)	The level of coherently aligned goals, assumptions, priorities, and improvement options across the agent network.
Problem-solving (G)	The domain knowledge, problem solving strategies, and resources (time, funding, cognitive bandwidth) available to the network.
System complexity factors (H)	The elements of complexity present. Using Ladyman and Wiesner's (2020) conditions for complexity: numerosity of elements; numerosity of interactions; diversity of elements; non-equilibrium; and feedback. Products (outcomes) of complexity: non-linearity; self-organization; adaptive behavior; and memory (Ladyman and Wiesner 2020).
External environment (I)	The external environment is that part of the contextual setting that falls outside the *interpreted and enacted boundary* (Nair and Reed-Tsochas 2019). As this boundary is dynamic and may vary between agents, it acts as a moderating factor, in the form of *coevolution constraints* (impact from shared and adjacent networks).
Learning gain	*The iterative nature of CQIN is difficult to convey via a static diagram. CQIN functions as a series of learning cycles. A learning gain from each iteration is posited as critical to acquiring sufficient knowledge within the agent network (Deming 1986; Doolittle 2014; Reed and Card 2016).

useful for exploring non-observable interaction effects between variables.

The CQIN simulation model was built using the NetLogo modeling software (Wilensky 1999). Building on a technique used for modeling virus spread in a computer network (Stonedahl and Wilensky 2008), the CQIN model simulates the uncertainty associated with agent interactions through a stochastic procedure. Introducing a configurable probability to each agent interaction overcomes the observation challenge of knowing which interactions will succeed and which will fail by randomizing this probability across the agent network activity.

Table 3 details the elements of the CQIN conceptual model now mapped to specific, configurable variables within the simulation model. To define and run the simulation, additional types of variable are required to provide construct definition, spreading processes, limits, and system reporting (Laubenbacher et al. 2007). The core threshold procedure variables are identified in Table 3.

The interaction of the highlighted variables via the threshold procedures is critical to the system behavior. The posited variable interaction is now examined more closely.

Figure 2 reflects the hypothesized primacy of the schema sharing (shared understanding between agents), an ability to learn from each iteration, and the moderating impact from the problem complexity and external constraints (i.e., a moderation that ranges from very little influence to fully inhibiting). *Schema share effectiveness* is regarded as the dominant variable as it is a necessary precondition of all the other variables as well as contributing in its own right (Cronin and Weingart 2019; Dooley 1997; Holland 2006, 2012). Relating Figure 2 back to the CQIN conceptual model in Figure 1, the path diagram shown in Figure 2 posits how the conceptual constructs of shared schema (F) and problem-solving (G) interact within an agent network carrying out improvement activity.

Within CAS, a small set of internalized agent rules are sufficient to produce complex agent activity (Holland 2012; Kauffman 2000; Mitchell 2009). Agent behavior in the CQIN model is governed by five rules only: (i) request to share information; (ii) request for cooperation; (iii) reject a request; (iv) accept a request, and (v) withdraw from the network. The threshold equations for these agent behaviors follow and a state transition diagram of the model's operational phases is provided in the Appendix.

Threshold equations

Achieving a shared problem understanding and response across the network is conditional on (i) successful transmission of information; (ii) acceptance of the new information; and (iii) ongoing cooperation of the receiving agents. The likelihood of achieving *intentional, system-level improvement outcomes* (S) across the agent network is described using conditional probability statements.

$$P(S) = P(A \mid B, C) = \frac{P(B \cap C \mid A)P(A)}{P(B \cap C)} \quad [1]$$

where:

A = *the initial level of shared schema between agents*
B = *successful sharing of relevant information*

C = *accepting and acting on the relevant information*

To establish successful information circulation across the network (B):

$$P(B)=\left(\sum_{i=1}^{n}a_i P\left(x<DE^F\left(\frac{GE^F}{H}\right)\left(\frac{(100-I)}{100}\right)\right)\right) \quad [2]$$

where:

a = *the agent attempting to share information;*

D = *the schema share effectiveness;*

E = *the learning gain per improvement iteration;*

F = *the improvement iteration count;*

G = *the problem solving effectiveness;*

H = *the problem space complexity; and*

I = *the coevolution constraints*

x = *random* $x \in \mathbb{N}$, $x \in [1,100]$

For each agent attempting to share with a connected neighbor, the total value of the agent sharing activity is compared to a random number between 1 and 100 (uniform distribution). The stronger the information-sharing and cooperation is, the more likely they are to be successfully transmitted among agents.

To accept new information and continue cooperation (C):

$$P(C)=\left(\sum_{i=1}^{n}b_i\frac{P((x<J)\cap P(x<K))}{P(x<K)}\right) \quad [3]$$

where:

b = *the individual agent receiving/acting on the new information;*

J = *the range of network interactions at risk*

K = *the rejection probability*

x = *random* $x \in \mathbb{N}$, $x \in [1,100]$

A random number between 1 and 100 (uniform distribution) is compared to the joint probability of the number of agents at risk of rejection and the probability of rejection. Higher risk levels increase the likelihood of rejection/non-cooperation.

Model assumptions:

- One iteration is one full learning cycle, equivalent to a *plan-do-check-act* improvement cycle with an initial signal detection phase added at the start.
- The equations can be adjusted to mimic the learning cycle phases, for example, the coevolution constraint can be set to be active (or more strongly influential) in the "do" and "act" phases, which require more interaction at the network boundaries.
- "Sharing information" and "cooperating" are entirely abstract constructs and treated as equivalent activities. They gain meaning based on which learning cycle phase applies at the time, for example, *notice this signal* (review), *is this signal important* (check), *what should we do* (plan), *can you assist us with this action* (do).
- Similarly, rejection or noncooperation are also abstract and may mean any of unwilling or unable to receive, absorb, agree or participate.
- Rejection/noncooperation can be treated as permanent or able to be reset after each iteration.

Having introduced both CQIN models in detail, we conclude the methods summary with a description of the real-world case study used to configure the simulation model.

Verification case study—antipsychotic deprescribing in aged residential care (ARC)

A significant proportion of people living in residential aged care are exposed to antipsychotic medication, with approximately one-third of residents receiving a medicine in this class at some point during their admission. Antipsychotic use is associated with significant adverse effects and an increased risk of premature death, especially for people with dementia. Studies of antipsychotic deprescribing in people living in residential care have produced mixed results (Van Leeuwen et al. 2018). Some trials have successfully withdrawn antipsychotics in participants without apparent adverse effect, but other trials have either failed to reduce antipsychotic use successfully, failed to show any benefit associated with cessation, or reported an increase in adverse outcomes associated with deprescribing (Van Leeuwen et al. 2018).

Nonpharmacological management strategies are the preferred alternative to antipsychotics for managing distressed behavior in people with dementia, but while these strategies reduce the need for pharmacological

Table 3. CQIN elements mapped to modeling parameters.

CQIN constructs	Configurable parameters in ABM
Agent network (A)	• Number of networked agents (nodes)**
	• Network type (e.g., preferential attachment or random)
	• Minimum node degree
	• Network centrality measures
	• Proportion of network cooperating
Operational processes and Outcomes (B), (C)	• Outcome measure at start
	• Targeted outcome measure value
	• Elapsed time (time "units")
System signals (D), (E)	• The range of signal detection amongst agents
	• Signal amplitude
	• Signal validity (relevant to objectives)
Shared schema (F)	• Schema share effectiveness**
Problem-solving (G)	• Problem-solving effectiveness**
	• Problem-solving resources
System complexity factors (H)	• Problem space complexity**
	• Network rejection scope**
	• Network rejection probability**
	• Minimum cooperation threshold
External environment (I)	• External (coevolution) constraints**
	• QI phase**
Learning	• Learning gain per cycle**
	• Improvement increments
	• Number of planned iterations
	• Number of iterations completed (count)**

**Core variables used in the threshold procedures

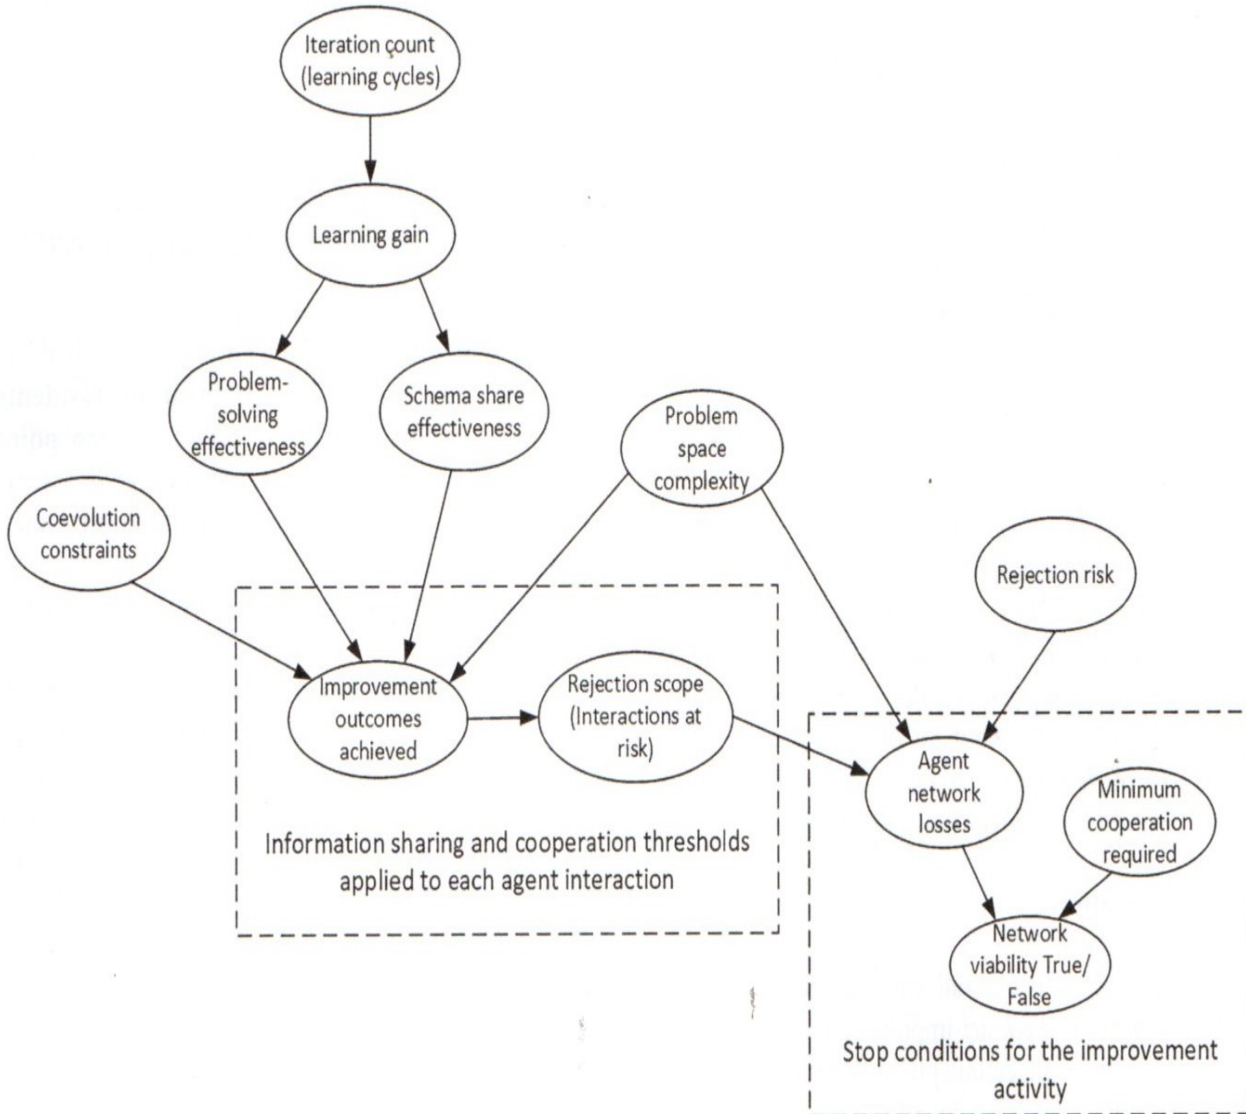

Figure 2. Core ABM variable interaction.

and physical restraints, they are not a panacea (Birkenhäger-Gillesse et al. 2018). Inadequately controlled aggressive or agitated behavior creates physical risk for care staff, other residents, and the affected individual. Other behavioral symptoms of dementia can cause stress and burnout in carers, reducing their capacity to provide effective person-centered and compassionate care.

When initiating an antipsychotic for people with behavioral and psychological symptoms of dementia, prescribers must consider the potential benefits and harms of *using* antipsychotics along with the potential benefits and harms of *not using* them. When evaluating deprescribing of antipsychotics, similar consideration is given to both arms (*use* versus *not use*) of this harms and benefits dichotomy. There is no simple guideline or one-size-fits-all solution to this complex issue. Antipsychotic deprescribing is an ideal initial verification case for CQIN simply because it is so complex. Seeking to minimize potentially clinically unnecessary prescribing involves a large heterogenous network of actors across several organizations and professional disciplines

The case study organization is a large private provider of ARC services in New Zealand and Australia. This organization maintains an interdisciplinary medications advisory committee that oversees performance and improvement of medication-related activities. The committee covers clinical and operational functions and includes general practice representatives, a pharmacist, a geriatrician, a specialist researcher, nurses, and improvement support staff. The interventional approach adopted is premised on providing timely, targeted prescribing data to prescribers and encouraging discussion and consideration of alternatives. The underlying hypothesis is that ensuring the appropriate information is readily available to those making patient-level decisions, combined with raised awareness of the alternatives, will allow frontline prescribers to identify early where any reduction in antipsychotics may be appropriate. Following an initial approach to gauge interest in participation, low-risk ethics approval was obtained from our primary research institution.

We began by using expert elicitation to understand the complex interactions and establish plausible starting values for the simulation model variables. This was not a trivial task given the novelty and abstract nature of CQIN. Clinicians and improvement facilitators are thoroughly versed in clinical and change management details, but CQIN focuses on shared understanding, network relationships, and the overall effectiveness of actions. We worked with small subsets of the medications advisory committee to first introduce and illustrate the concepts and then consider potential values of the variables in the context of the deprescribing project. Obtaining valid expert assessments of probability is recognized as difficult (Fenton and Neil 2019; Morgan 2014). An elicitation protocol was developed and the information gathering process was highly iterative, updating as new information or understanding became available. Monitoring of the project is ongoing; however, after several rounds of elicitation activity a credible set of starting values for simulation was established.

Results

Simulation experiments involved networks of cooperating agents running through sequences of learning (improvement) cycles. We measured the resulting progress of each network toward a prescribed improvement target. Understanding the stop conditions is important to interpreting the results. The sequence stops under two conditions: either the programmed number of cycles are completed or the number of cooperating agents has fallen below the threshold. Where the network stopped prior to reaching the improvement target, we classified this as a failed sequence. Each complete learning sequence was repeated a minimum of 1,000 times to observe the stochastic variation among the iterations. For this study we conducted three separate experimental tests:

- Test 1: Assess the probability of the deprescribing CAS achieving its improvement objectives, based on the elicited values from the case study.
- Test 2: Measure the influence of the key *schema share effectiveness* variable.
- Test 3: Conduct a sensitivity analysis using the elicited values from the deprescribing case as the control.

Simulation of deprescribing project

For Test 1, the simulations predicted that the desired 10% reduction in antipsychotic prescribing would be achieved with a probability of 69%, based on the system variables as configured. This result is based on 69% of the 1,000 simulated improvement sequences being successful and 31% failing. Test 1 also revealed a substantial nonlinear transition zone in the progression toward project success. This transition is shown in Figure 3. The y-axis of Figure 3 shows progress toward the goal state of a 10% reduction in

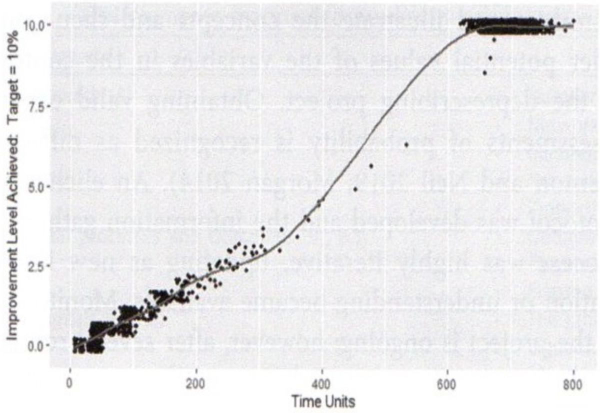

Figure 3. Improvement level reached.

inappropriate prescribing. The x-axis displays "time units"—an abstract representation of elapsed time using the program clock cycles. Elapsed time is a proxy indicator for the amount of learning achieved. Any learning acquired within an improvement cycle improves the starting point for the subsequent cycle; thus, problem-solving increases over time as more cycles occur.

Each dot in Figure 3 is one complete repetition of a full learning cycle sequence. The lower left-hand quadrant reveals that those sequences that "fail" predominantly do so very early in the process. As time (and learning) progresses, reaching the target of 10% becomes increasingly likely, and failure becomes rare after about one-third of the planned 20 QI cycles have taken place.

Impact of schema share effectiveness

Test 2 examined the impact of varying the critical *schema share effectiveness* variable across five ordinal levels as shown in Figure 4. The levels of *schema share effectiveness* were set from 10 (very low) to 90 (very high), with 2,000 repetitions taken at each level. Meeting the improvement target of 10% is indicated as target achieved = true or false.

Figure 4 reveals again the distinct threshold between two system states, with one state being a low probability of success and the second state being a high probability of success. This result indicates that *schema share effectiveness* variable must meet a minimum threshold for success to be probable, and beyond this threshold failure becomes unlikely. A second, smaller transition zone is also observed within the successful repetition subset. This effect is caused by the time delay introduced when an initial schema share request is unsuccessful, but sufficient agents nevertheless remain active in the network to allow subsequent attempts. These repetitions are therefore ultimately successful but take longer to reach the improvement target.

Sensitivity analysis

Test 3 reveals that successfully complete a learning sequence is highly dependent on all the initial system values. Table 4 shows the impact of stepwise changes in key CQIN variables. Continuing with variation in *schema share effectiveness* as an illustration, the final row in Table 4 shows the percentage of successful repetitions ranging dramatically from the initial value of 69% down to less than 10% and then up to over 90% (test configurations 1–3 in Table 4).

The results shown in Table 4 reflect the nonlinear and holistic nature of CAS, that is, the impact of even small interventions on individual elements is unpredictable *and* contingent on the settings of other critical elements. Some interventions have nonlinear effects (disproportionate to input changes), based on the threshold settings for the variable, for example, varying the *rejection chance* between 5% and 15% moves the likelihood of success from near-certain to impossible (test configurations 6 and 7). Individual change levers can be adjusted, but *all* the variables must be on the right side of their respective thresholds for successful problem solving.

Test 3: Sensitivity test details:

- Each test configuration consisted of 1,000 repetitions.
- Configuration 1 utilized the initial estimated parameters determined in conjunction with project members.
- Configurations 2–17 then varied 1 parameter in stepwise increase or decrease from the initial value.
- The percentage of repetitions fully meeting the desired improvement objective (10% reduction) is shown in the final row.

Model validity and verification

Content validity requires that all necessary elements are included in the model and that unnecessary elements are excluded. In computer modeling, face validity refers to the faithful representation of an underlying conceptual model (Korb, Geard, and Dorin 2013). In other words, do the model outputs accurately represent the concepts as interpreted? By deriving the simulation model explicitly from each conceptual model element, we ensured that the two models are highly consistent. The simulation model includes the

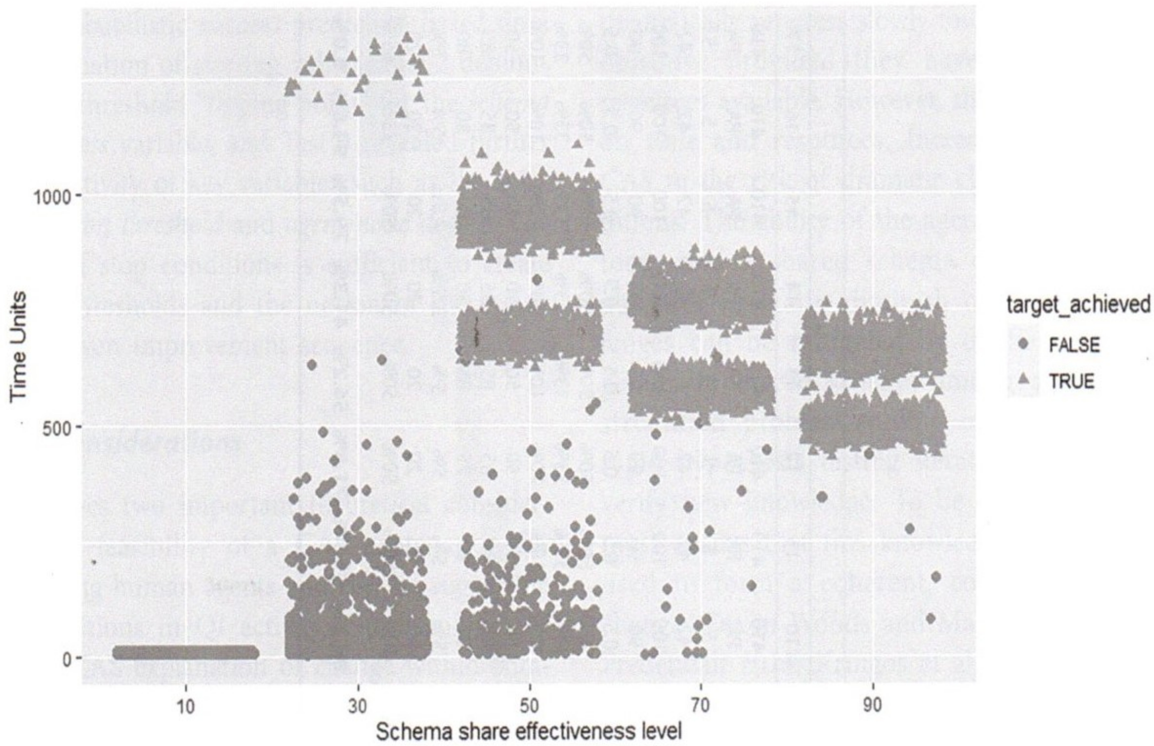

Figure 4. Impact of varying the levels of schema share effectiveness.

elements necessary for agent network information processing, cooperation, and problem solving, in a manner consistent with the initial theoretical assumptions. Sensitivity analysis tested this performance over a wide range of input conditions.

Verification of a computer simulation model asks whether the model, accurately recreating the underlying concepts internally, is an appropriate representation of the real-world phenomena being modeled (Korb, Geard, and Dorin 2013). Comprehensive empirical assessment of CQIN is part of our ongoing research. The results presented here support preliminary verification, based on the model's ability to emulate a complex healthcare QI example involving a large interdisciplinary team.

Our approach to validation and verification of the agent-based model followed Bayesian principles as recommended by Korb, Geard, and Dorin (2013). Expert opinion provided a subjective Bayesian prior understanding: in our example, expertise includes the original CAS theory underpinning the model and the frequent specialist reviews within the case study to confirm face and content validity of the resulting models. Sensitivity analysis then provided opportunity to review model outputs and assess internal validity of the model variables (Korb, Geard, and Dorin 2013). The simulation model produced predictive outputs and these data provide a likelihood in Bayesian terms. The results do not represent a single target prediction, but a range of probabilities given the specific input conditions.

Discussion

Critical assessment

The simulation experiments tested the feasibility and underlying mechanisms for a purely CAS explanation of QI. Values for key CAS variables were taken from a real-world example and tested under scenarios that would be infeasible to observe empirically. Simulation also provided the ability to "repeat" an improvement project. In addition to supporting model validation, repetition allowed a considered assessment of the role of uncertainty and chance within the CAS interactions.

The hypothesis under test is: *(H1) Updating and aligning agent schema takes place fully within the constraints of the agent behavioral rules.* The process of responding to a signal for change, acquiring sufficient knowledge, and cooperating to achieve change was achieved exclusively through the programmed agent sharing and cooperation rules. No other mechanism for aligning schema existed within the simulation. Therefore, the ability of the simulated agent network to successfully reach the improvement target directly supports H1. System learning was managed at the agent behavior level of granularity, without recourse to agent hierarchy or more complex agent rules (Carmichael and Hadžikadić 2019).

The results of all three experimental tests also support the general CAS principles of uncertainty, nonlinearity, dependence on starting values, and dependence on the values of other system variables. Test 1 produced

Table 4. Sensitivity analysis for the antipsychotic deprescribing case.

ABM variable	Test configurations																
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Number of networked agents	410	410	410	410	410	410	410	410	410	410	410	410	410	410	410	410	410
Network type: (preferential attachment 'PA')	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA	PA
Minimum node degree	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5	5
Initial signal detection	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42	42
Signal base	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100	100
Signal target	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90	90
Signal increment	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Schema share effectiveness	50%	40%	60%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
Rejection threshold	33%	33%	33%	25%	40%	33%	33%	33%	33%	33%	33%	33%	33%	33%	33%	33%	33%
Rejection (disengage) chance	10%	10%	10%	10%	10%	5%	15%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%
Problem-solving capability	50	50	50	50	50	50	50	40	60	50	50	50	50	50	50	50	50
Problem space complexity	85	85	85	85	85	85	85	85	85	80	90	85	85	85	85	85	85
Coevolution constraints	80	80	80	80	80	80	80	80	80	80	80	75	85	80	80	80	80
Learning gain	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	5%	2.5%	7.5%	5%	5%
Planned iterations	20	20	20	20	20	20	20	20	20	20	20	20	20	20	20	20	20
Required minimum proportion of network cooperating	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	33%	66%
Percentage of repetitions meeting full improvement target	**69.0%**	**9.0%**	**94.0%**	**98.3%**	**14.0%**	**99.6%**	**0.0%**	**53.9%**	**75.6%**	**88.3%**	**55.9%**	**71.8%**	**56.2%**	**41.3%**	**79.5%**	**94.6%**	**0.0%**

a plausible, probabilistic success prediction based on a realistic combination of starting values. Test 2 demonstrated a clear threshold "tipping point" for the *schema share effectiveness* variable, and Test 3 revealed further nonlinear sensitivity of key variables such as the *minimum cooperation threshold* and *agent node degree*. The introduction of stop conditions is sufficient to create these system thresholds and the notion of success or failure for a given improvement sequence.

Theoretical considerations

This study raises two important theoretical considerations: (i) the feasibility of a CAS explanation for change involving human agents and (ii) the suggestion of phase transitions in QI activity. Our results show that a purely CAS explanation of change within complex sociotechnical contexts is feasible. We used *schema share effectiveness* to represent how the goals, assumptions, and problem-solving options form the foundation of shared problem representation among agents (Anderson 2005; Cronin and Weingart 2007, 2019; Holland 2012; Mohammed, Ferzandi, and Hamilton 2010). Problem-solving provides the cognitive framework necessary to replace the biological processes taking place in natural CAS. However, the share/cooperate rules on their own do not guarantee success. Sufficient knowledge must be accessed or generated, selected, *and* applied in real-world tests of fitness. At the agent level, adaptation is the updating of the internal schema. If system conditions allow, system-level adaptation takes place when the level of aligned schema across the agent network is sufficient to enact change and successfully reach a goal state (Carmichael and Hadžikadić 2019). Our micro-level agency ultimately has consequences that effect the system, and these effects manifest at the macro (system) level (Carmichael and Hadžikadić 2019; Wilson and Kirman 2016).

The adaptation of a complex adaptive system involves learning (Doolittle 2014; Holland 2006, 2012). If a chosen action is working, any progress is understood via feedback signals and incorporated into the next actions. If the chosen action is not effective, either the implementation or the underlying change hypothesis needs to be revisited. Such "double loop" learning is long established in the literature (Ashby 1961; Argyris 1976; Umpleby 2008) but hard to achieve in QI practice (Reed and Card 2016). In this iterative manner, an agent network progresses through the available problem space until an effective solution is found (Anderson 2005; Newell and Simon 1972). In contexts that we might describe as low-complexity or benign, a low-skilled agent network could theoretically progress slowly toward their improvement objective provided they have sufficient time and resources available. However, there are very real limits on time and resources. Increased time exposes the CAS to the risk of dramatic change in external conditions. The ability of the agent network to maintain focus and a shared schema on this particular QI objective may also diminish over time. These challenges can be mitigated by deliberately building a genuinely shared schema among agents and using structured problem-solving methods, for example, rapid hypothesis testing iterations, to generate and verify new knowledge. To be effective, practitioners must ensure that this knowledge is not isolated but used to form a coherent, compounding force for change (Dixon-Woods and Martin 2016; Kovach and Fredendall 2013; Kringos et al. 2015).

Variations in the agent network structure, together with the uncertainty associated with agent interactions, creates conditions where cascading failure is possible (Mitchell 2009). A large, heterogenous agent network is made up of multiple communities and clusters with different levels of connection (Barabási 2016). In these conditions, individual agents with high node degree who stop cooperating have a large impact and can trigger a cascade of network losses. A genuinely complex problem means that shared problem interpretation and physical cooperation are difficult to achieve, and this results in a higher probability of resistance. These network losses lead to the phase transitions observed in our simulations. The transition between states is a form of self-organized criticality, manifesting as a cascading failure triggered from the expected micro-level system activity (Bak and Chen 1991). For the deprescribing project simulations, we defined when the system transitions would occur *via* a *maximum level of noncooperation* variable (initial value set as 50% of the network). Accepting that the location of the resulting transitions is unlikely to be empirically precise, it is nevertheless axiomatic that *any* agent network has a minimum critical mass for cooperation to be sustained. System state transitions have significant implications for QI practice. Working with a CAS explanation of change, identifying the *minimum* values of the critical CAS variables and where a given improvement network sits in relation to the system transition points becomes essential.

Practical considerations

CQIN offers QI practitioners two practical contributions: (i) a complexity risk assessment tool and (ii) the imperative of maintaining agent network viability.

As the results in Table 4 demonstrate, the starting conditions for QI activity are critical and *all* elements must function at a minimum viable level. QI teams can use CQIN to assess and reduce their exposure to CAS-specific risk factors. Navigation from the abstract elements of the CQIN model back to tangible actions is straightforward, determined by which CQIN variable is the focus for action. For example, if available evidence suggests that the *schema share effectiveness* is low due to poor intergroup communication, improving this underlying cause becomes a practical intervention to prioritize. A lack of solution options within the network's current knowledge is another example. In this case, seeking external guidance from outside the network, or running trials to generate new options, would be the important tasks. Addressing underlying causes is no different than for conventional QI, but practitioners can use CQIN to identify CAS focus areas most likely to shift system outcomes.

Applying CAS framing to QI dramatically exposes the importance of maintaining sufficient energy and information flow within the agent network. Attention to the rate of network resistance or noncooperation is therefore just as important as direct actions toward improvement goals. Assuming the existence of a minimum effective cooperation level for agent network viability is sensible, even when precise measurement is impossible. Recognizing that a viability threshold exists allows QI practitioners to factor agent network topology and cooperation rates into their actions. In practice, this might mean increasing the communication paths between disparate communities of agents and increasing the node degree (the number of agent connections) of key individual agents (Barabási 2016; Bar-Yam 2004).

Limitations

We have described the relationships between the variables encoded in our probability equations, but an important caveat is that these equations make no attempt yet to weight the variables precisely. Weighting requires more empirical data and verification to confidently determine the values of any weighting or scaling factors. Elicitation of accurate probability values from subject experts is a challenging task in practice, and the process does not have a set endpoint. The objective for this first phase of our work has been to construct a plausible model and not to add information that can only be obtained empirically.

Additional and individually important features of any specific CAS might also be proposed for inclusion within CQIN, for example, modeling the influence of leadership, communication, and trust across the agent network (Lorden et al. 2014). Our agent network was modeled as non-directional, in that initiation of sharing or cooperation between two agents could take place in either direction. In real-world scenarios, some degree of hierarchy between agents and the directionality of system signals and schema sharing is assumed. CQIN focuses on the ultimate *impact* of these variables, that is, poor leadership or low trust will eventually result in low shared schema representation. The level of granularity chosen helps CQIN achieve its goal of high generalizability, while still retaining an ability to deconstruct the contributing factors.

Conclusion

This paper has presented the CQIN model for understanding QI in healthcare CAS environments. CQIN provides a working example of a *CAS explanation for QI activity*. The shared mental model or schema sufficient to achieve change evolves via a very small number of internalized agent rules (in our example, five). A CAS explanation of change directly addresses a research gap regarding design within CAS. The emergence of shared mental models is not an example of design imposed from outside the agents, nor is it a teleological system purpose. Two conceptual pivots help to apply CAS thinking to QI. *First*, the "complex adaptive system" should be understood to be a network of agents acting adaptively. *Second*, the recognition that generating the desired *system coordination* behavior from simple rules does not replace or challenge human agency. Countless nuanced actions and inputs contribute to the decisions to share information or cooperate with other agents. CAS rules are not in conflict with agency; However, our micro-level agency ultimately has consequences that effect the system, and these effects manifest at the macro (system) level. Complex problem-solving within a CAS requires that we navigate gaps in shared understanding, high levels of uncertainty, nonlinearity in process outcomes, and lack of timely feedback signals. An explicitly CAS explanation for QI is valuable because it provides QI practitioners directly relevant insights into managing these difficulties. Our CAS explanation of QI is still rudimentary, and there are many opportunities for further research. Immediate research priorities include modeling the impact of hierarchical network structures, further empirical measurement, and assessing methods for incorporating expert perceptions of agent network performance and risk.