Title:

A simulation framework for mapping risks in clinical processes: the case of in-patient transfers

Authors:

[1]Adam G. Dunn PhD, [1]Mei-Sing Ong, [2]Johanna I Westbrook PhD, [1]Farah Magrabi PhD, [1]Enrico Coiera MBBS PhD, [3]Wayne Wobcke PhD

Affiliations:

[1]Centre for Health Informatics, Australian Institute of Health Innovation, University of New South Wales, Australia

[2]Centre for Health Systems and Safety Research, Australian Institute of Health Innovation, University of New South Wales, Australia

[3]School of Computer Science and Engineering, University of New South Wales, Australia

Funding:

Australian Research Council Linkage Grant (LP0775532) and National Health and Medical Research Council (NHMRC) Program Grant 568612.

Corresponding author:

Dr Adam Dunn, Centre for Health Informatics, University of New South Wales, Sydney, NSW 2052 Australia. Ph: +61 29385 8699 Fax: +61 29385 869 Email: a.dunn@unsw.edu.au

Manuscript information:

Word count estimate (excluding abstract, references and captions):
Abstract word count: 244
Number of Figures: 4
Appendices: 2
Pages: 19

I, as corresponding author, promise that I and all persons listed as coauthors on this submitted work have read and understand the "Originality of Manuscripts statement regarding submissions to JAMIA (available at http://jamia.bmj.com/site/about/originalityofmanuscripts.xhtml and confirm that this submission is a new, original work that has not been previously submitted, published in whole or in part, or simultaneously submitted for publication in another journal. Also, in accordance with the aforementioned policy, we have included as part of the submission any previously published materials that overlap in content with this new original manuscript.

The Corresponding Author has the right to grant on behalf of all authors and does grant on behalf of all authors, an exclusive license (or non-exclusive for government employees) on a worldwide basis to BOTH The American Medical Informatics Association and its publisher for JAMIA, the BMJ Publishing Group Ltd and its Licensees to permit this article (if accepted) to be published in Journal of the American Medical Informatics Association and any other BMJPGL products to exploit all subsidiary rights, as set out in our license (http://group.bmj.com/products/journals/instructions-for-authors/licence-forms).

Abstract

Objective: To model how individual violations in routine clinical processes cumulatively contribute to the risk of adverse events in hospital using an agent-based simulation framework.

Design: An agent-based simulation was designed to model the cascade of common violations that contribute to the risk of adverse events in routine clinical processes. Clinicians and the information systems that support them were represented as a group of interacting agents using data from direct observations. The model was calibrated the model using data from 101 patient transfers observed in a hospital and validate the results for one of two scenarios (a misidentification scenario and an infection control scenario). Repeated simulations using the calibrated model were undertaken to create a distribution of possible process outcomes. The likelihood of end-of-chain risk is the main outcome measure, reported for each of the two scenarios.

Results: The simulations demonstrate end-of-chain risk of 8% and 24% for the misidentification and infection control scenarios, respectively. Over 95% of the simulations in both scenarios are unique, indicating that the in-patient transfer process diverges from prescribed work practices in a variety of ways.

Conclusions: The simulation allowed us to model the risk of adverse events in a clinical process, by generating the variety of possible work subject to violations, a novel prospective risk analysis method. The in-patient transfer process has a high proportion of unique trajectories, implying that risk mitigation may benefit from focusing on reducing complexity rather than augmenting the process with further rule-based protocols.

I. Background

Adverse events occur in five to ten percent of all hospital admissions and result in significant social and economic costs (1-3). Hospitals are complex systems (4, 5), where care is delivered via a web of organisational (6, 7) and social interactions (8) supported by information systems (4, 9, 10). Violations of prescribed work practices, which include deliberate and accidental deviations from prescribed work practices (including communication errors), contribute to a large proportion of adverse events (11-13). Individual violations do not always lead to adverse events but rather, adverse events are often caused by a chain of violations (14-16).

Given the complexity of the socio-technical interactions in healthcare delivery, and the regularity with which work activity deviates from prescribed practice, existing methods of risk analysis that borrow heavily from risk analysis in engineered systems may not always be appropriate. Risk analysis involves identifying where risk exists, as well as identifying the causes and impacts of adverse events (17). Common methods of prospective risk analysis (as opposed to retrospective analyses such as root cause analysis, which is performed after an adverse event has occurred) include Failure Modes and Effects Analysis (FMEA) (18-22), and Hazard Analysis and Critical Control Points (HACCP) (23-25). More quantitatively-focused analyses include Probabilistic Risk Assessment, involving either Fault Tree Analysis (FTA) or Event Tree Analysis (ETA) (17, 26, 27).

As quantitative methods, FTA and ETA may be used to prospectively analyse the likelihood of systemic failure given the individual likelihoods of specific faults and events. While the methods address multiplicative (combinations of events that must occur together to create risk) and additive (any one event can trigger an adverse event risk) likelihoods, the methods have general limitations. Firstly, the methods rely on a linear dependency of faults to create the structure of static trees, whereas healthcare environments are more dynamic than this – actors modify their behaviour based on what they know about the environment as it changes in time – they have non-linear dependencies. Secondly, workarounds (28, 29) contribute to risk in the same manner as other violations but evolve as a consequence of optimising the multiple objectives of healthcare delivery – quality, safety and efficiency. Existing methods of prospective risk analysis consider faults as binary entities (26) and do not consider the manner in which workarounds may still achieve the intended goal after diverging from the prescribed work practice.

Our aim was to provide an alternative approach to this problem of risk assessment by explicitly simulating the trajectories of routine processes as they deviate from prescribed work practices, to determine which trajectories may lead to adverse events, and estimate how often these risky trajectories occur. We define a trajectory as a single instance of a routine work process, which is characterised by sequential and parallel actions taken by both human and information system agents. An end-of-chain risk is defined by counting the proportion of trajectories in which a specified adverse event is possible – either via misidentification of a patient or compromised infection control in the two scenarios we examine in detail.

II. Methods

In what follows, we describe the general methods for constructing an agent-based model for prospective risk analysis and then illustrate the method using an in-patient transfer case study. We additionally describe the simulation approach and validation of the model. For the purpose of this research, we use the term *violation* to describe any intended deviation from a prescribed work process (such as a workaround), unintended deviation (such as a slip) or a communication error (in which an information exchange does not occur or is erroneous). Fig. 1 is a schematic representation of arbitrary workflow process illustrating how violations create divergence away from the prescribed work practice, leading to risk for a proportion of the trajectories. Each fork in the trajectory diagram represents a possible violation.

Figure 1 approximately here

To simulate work processes we used agent-based modelling to specify and simulate the behaviour of actors in a clinical process, as well as the information systems the actors use. An *agent* in an agent-based model may be described as an entity that is situated in, perceives and affects an environment in order to achieve one or more goals (30, 31). In multi-agent systems, the agents interact with each other and their environment to achieve individual goals that may or may not be in conflict with the goals of others, and may coordinate their behaviour to achieve shared goals.

In an agent-based model, agents know how to carry out their own roles within the work process, including communicative actions that involve other agents or the environment. Agents do not need to have knowledge of the entire work process but rather, only knowledge about the prescribed work practices they require to achieve their own goals. For example, in the case study presented below, the Porter agent has knowledge of what activities to perform

if the patient has an infectious status but does not know how to perform the patient ID check usually performed by the Ward Nurse (see Appendix).

Agents hold beliefs about the current state of the environment and these beliefs may be partial or incorrect. Agents update their beliefs as they participate in the work process and when they perceive changes to the state of their environment, including objects and other agents. For example, the Porter agent updates his belief about the infectious status of the patient after interacting with the Transfer Form – if the information is incorrect on the Transfer Form, then the Porter may believe the patient is not infectious (see Appendix).

An action is an atomic element in the work process performed by one agent, or concurrently by two or more agents. An agent may undertake an action within their repertoire when triggered by a change in the environment, required by the process, or requested by another agent. When there are multiple actions that may be performed an agent has the autonomy to choose an action. If there is more than one way to perform an abstract action, agents have the autonomy to choose how to perform that action, or to omit the action entirely, depending on their beliefs. For example, the Porter must use adequate infection control for an infectious patient but may choose to perform the required actions in different orders or at different times during the work process (see Appendix)

Brahms (Agent iSolutions, NASA Ames Research Center, California, USA) is an example of an agent-based modelling environment (32, 33) and it is appropriate for modelling clinical work processes because it captures the fine-scale decision-making and behaviour of the human agents and their interactions with information systems, other agents and the environment, without onerous requirements for modelling the cognition of actors. Brahms is developed to be used to model environments at a more abstract level than models of individual cognition (34-38), and a less abstract level than discrete event systems used to simulate patient flow over much longer time periods (39-43).

A. Formulation process

A user of the simulation framework formulates a new model by firstly translating a set of work practices and associated violations into the design of an agent-based model. Prescribed work practices are context-specific implementations of one or more guidelines, policies or protocols. There may be multiple ways of completing a task that are in accordance with a policy, so we describe a set of prescribed work practices that lead to a safe outcome. The

requirement for modelling in an agent-based model is that the prescribed work practices describe the process in enough detail for actions to be represented at the same level of abstraction as the observed violations.

Having implemented a formal description of the agents, and their beliefs and actions, the violations enacted by agents are assigned probability functions, based on the empirical data. When an agent encounters a system state in which a violation may be enacted, the agent modifies its beliefs or environment in different ways depending on the outcome of the probability function. This implies that two simulations instantiated in exactly the same way may have different trajectories and thus, different outcomes in regards to risk.

Each simulation is run individually within the Brahms environment and exported as time-stamped logs of activities and changes in beliefs. A log of activities and changes in beliefs describes a trajectory. The process is implemented such that if all decisions are taken as expected (no agents ever enact a violation of the prescribed work practice), then the agents behave according to the prescribed work practice. When the likelihood of violations is non-zero, the behaviour of the agents may diverge from the prescribed work practices.

By aggregating the trajectories of repeated simulations that are instantiated consistently according to a specified scenario, it is possible to quantify and examine the distribution of trajectories and distinguish those in which a risk of an adverse event is present. The number of simulations required to accurately map a set of trajectories is described in Table 1. While relatively few simulations are required to detect a single end-of-chain risk, rare trajectories that contribute to end-of-chain risk require many more. Analysis involves measuring the proportion of simulations that contribute to or pass on risk during the main stages of the process.

Table 1 approximately here

B. Case study – in-patient transfers

The process of inter-ward patient transfer we chose to simulate is an example of a process that is characterised by *ad hoc* communication, inconsistency and poor team coordination (44). Actors involved in the in-patient transfer process are expected to conform to infection control precautions and patient identification checks defined by policy (45, 46). The aim of the process modelled here is to safely transfer the correct patient from a ward to the radiology department within the hospital. The process begins when it is scheduled by the coordinator in

radiology in response to a request from a ward. The coordinator instructs a hospital porter to transfer a patient from the ward. The porter is given a transfer form, which contains information about the patient including multiple-redundant identification information (including a unique identifier) and any specific requirements for infection control associated with this patient. The coordinator may additionally communicate these requirements verbally to the porter. On arrival at the ward, the porter hands over the transfer form to a ward nurse for a sign-off – this includes verification of the patient's identification and the correctness of the information on the transfer form. The ward nurse may additionally communicate transfer requirements to the porter verbally. If a nursing escort from the ward is required to accompany the porter, a clinical handover will take place between the ward and radiology nurses.

A set of violations associated with the inter-ward patient transfer process were recorded in a separate study involving 101 transfers (47). In that study, four violations were observed in each transfer, on average. In that study, the results suggest that a failure to perform patient identification checks was a significant issue (occurs in 42% of transfers) and the use of adequate infection control precautions was used as an example for redundancy analysis. Poor compliance rates were suggested as the main cause of low system reliability. The data from this observational study were used to create the prescribed work practice in the model (the actions that conform to the policies (45, 46)) and to populate the model with the likelihoods of each violation. Ethics approval for the protocol was received from the New South Wales Department of Health Ethics Committee (EC00290) with application number 08/019 and site specific application 08/043.

In our model, we represent the in-patient transfer process using 4 human agents, 6 objects and 186 activities, of which 31 are driven by an empirically-defined likelihood. Besides the human agents, we also define the non-human information systems as agents, namely the transfer form and the patient record, since they participate in communication acts and hold possibly erroneous information. A simplified version of the process is represented in the results section below. Violations most associated with misidentification and infection control are labelled where they occur in the process. The flow of activities is broadly represented from top to bottom for the significant agents in the process (including the coordinator, the radiology nurse, the porter, the patient's record and identification, the transfer form and the ward nurse), following the solid arrows. Communications between each of the agents are illustrated as vertically-oriented communication channels. Although represented

simplistically in comparison to the full specification of the behaviour, the schematic captures the main interactions that may create, pass along or ameliorate risk.

C. Structural and behavioural validation

In order to demonstrate the validity of the approach, we follow the formal process detailed by Barlas (48), which includes both structural and behavioural validation. Structural testing was done here via the analysis of boundary conditions, since the structure necessarily conforms to the prescribed work practices. Therefore, the model is checked to ensure that agents always perform the prescribed work practice in the absence of any violations. For the scenarios we examine, this means that a misidentified patient will be recognised at the first identification check, and adequate infection control is always used when a patient is recorded as being infectious. The resulting trajectories should be the same for every simulation and always match the prescribed work practice. This was confirmed by repeated simulation, in which the likelihood of all violations was set to zero. Behavioural testing involves comparing the behaviour of the model with the behaviour observed in the real world. This validation is performed for the infection control scenario but is not performed for the misidentification scenario because misidentified patients were never observed in the associated observation study. Consequently, we did not perform tests to determine the predictive validity of the model and the behaviour of the model is validated by demonstrating that the model reproduces the aggregate behaviour under boundary and a sample of realistic conditions. In the latter case, the model is calibrated by the same data against which it is tested.

D. Simulation experiments

The purpose of running repeated simulations is to determine the range of potential trajectories that evolve as a consequence of combining individual violations within a framework that captures the technical and social constraints associated with interacting agents in a hospital environment. Each simulation is instantiated by defining values for the agents and objects (creating the scenario) and pre-defined likelihoods for each individual violation (creating the potential for divergence from the prescribed practice). While the violations are defined by empirical results, the properties of the agents and objects are instantiated to reflect a specific scenario – in the case study this is either the misidentification scenario or the infection control scenario. In order to produce coherent measures of risk, the resulting trajectories (one for each simulation) are categorised as being incomplete, completed without risk, and completed subject to risk, where the risk is defined by the context of the scenario.

In the misidentification scenario, a patient requires transfer from the ward to radiology but is incorrectly identified at the outset (the ID band does not match the patient's record), indicating an initial risk of 100%, which is reduced during multiple identification checks. In this case, the risk is ameliorated during the routine process. The purpose of the scenario is to see if the process catches or passes along the risk of misidentification. A completed transfer with a misidentified patient may lead to a wrong patient/location/test adverse event.

Alternatively, in the infection control scenario, new risks evolve as a consequence of inadequate infection control. In this case, an infectious patient (correctly identified as infectious on the patient record) is transferred to Radiology. There is no initial risk but the effects of individual violations combine to generate the potential for inadequate infection control and an end-of-chain risk of contagion. In this case, the purpose of the scenario is to see if the process generates new risks in relation to infection control. If inadequate infection control is used during a patient transfer, the process may lead to hospital-acquired infections, which are considered adverse events.

III. Results

In the misidentification scenario the process begins with 100% risk of misidentification because the simulation is instantiated such that the patient's record information does not match the patient. Through each simulation, identification checks involving the patient, patient ID band, patient record and transfer form are used to confirm the identity of the patient and any of these checks can ameliorate the risk. The reduction in risk is illustrated by the reduction in thickness of the aggregated trajectories represented in Fig. 2. In Fig 2, the boxes are a simplification of the policy, the vertical arrows indicate chronology for the six agents and the horizontal channels indicate communication between agents. The risk of misidentification is ameliorated partially by the Ward Nurse and to a greater degree by the Radiology nurse during identification checks performed with the patient (Fig. 2). Of the 500 simulations, 323 featured either the capture of a misidentification or a complete transfer where the misidentification error is never caught. The remaining simulations were halted as a consequence of other factors including patient unavailability and equipment failure. Of the 323 that met the criteria, 26 (8.05%) were complete transfers in which the misidentification was not caught. In 297 (87.3%) simulations, the patient reached Radiology before the misidentification error was identified. Since there were no misidentified patients observed in the 101 observed patient transfers, we are unable to validate the likelihoods.

Figure 2 approximately here

In the infection control scenario, the risk begins at 0% (an infectious patient is not an infection control risk if all actors adhere to policy) and accumulates as a consequence of missing information on the transfer form or in verbal exchanges, and the Porter's violations relating to adequate infection control precautions. The risk of inadequate infection control is generated as a consequence of missing or incorrect information on the transfer form, and from violations committed by Porters with access to the correct information (Fig. 3). Of the 500 simulations, 345 were complete (others were stopped for reasons such as patient unavailability and equipment failure during transport). For the 345 simulations that were complete, the number of times that the agents did not ensure adequate infection control precautions was 84 (24.3%), assuming post-completion tasks were always completed. As a consequence of having received misinformation, the lack of information exchange, or the lack of response to correct information, the Porter completes the process with an infection risk in 80 (23.2%) cases. The Ward Nurse completes the process with an infection risk in 37 (10.7%) cases. The results may be interpreted to mean that, in the presence of the workarounds observed, the process is effective in ensuring adequate infection control 76.8% of the time.

Figure 3 approximately here

In order to validate the results of the infection control scenario, we compare the observed behaviour of the agents in response to an infectious patient with the simulated behaviour of the agents in response to an infectious patient. In the observations, 27 transfers involved an infectious patient. Of those, 12 transfers were completed without the porter using infection control (which gives a 95% CI 24% to 65%). In comparison, 80 of the 345 simulated transfers were completed without the ward nurse or the porter ensuring adequate infection control (which gives a 95% CI 19% to 28%). The results suggest that the simulated results match the observed results ($p=0.0136$), however, we would be hesitant in using this to form a conclusion due to the small number of observations and because it is only one behaviour amongst the multitude represented in the simulations.

IV. Discussion

The value of using this simulation framework is that it allows a user to quantitatively examine how individual violations combine along trajectories of routine work practices to create risk. The results suggest two main conclusions. Firstly, the proportion of unique

trajectories is very high as a consequence of the number of redundant steps required in passing along information about the patient's identity and infection control requirements. Secondly, chains of percolated errors that lead to risk (or fail to ameliorate risk) are visible when the results of many repeated simulations are aggregated.

The simulation framework presented here complements existing prospective risk analysis methods. It requires more detailed observation of work practices and individual violations than with traditional prospective risk analysis such as FTA and ETA, and consequently has some benefits over existing methods. Firstly, the trajectories offer more precision because the simulation framework explicitly models the behaviour of the individuals in the scenario and the information systems with which they interact (rather than logically combining the probabilities of individual violations). A further benefit of this precision is that it allows a user to uncover a fuller range of possible work activity trajectories, of which many may not have been observed (this is indicated by the number of unique trajectories in the simulations). Finally, in traditional forms of prospective risk analysis, the user encodes probabilities for a pre-determined risk of their choice, whereas the simulation framework is scenario-driven. In this case, extra information is available about the likelihood of conforming to the prescribed practice, near misses, rare events and potentially unforeseen risks.

A. The in-patient transfer process

By simulating the in-patient transfer process with zero violations, we have tested the correctness of the policy as it is implemented in the prescribed work practices of the hospital. If no violations occur, the process of identification check will mitigate any incoming risk of transferring the wrong patient at the ward, and porters will always maintain adequate infection control. The presence of violations (however they may have evolved in the context of the patient transfer process) leads to the presence of risk and thus the policy is not perfectly reliable in that it is unable to adapt to violation. Like any work process, the in-patient transfer process is both procedurally correct and unreliable at the same time.

Regarding the misidentification scenario, the simulations suggest that the process will capture approximately 92% of wrong patient errors, leaving 8% in which there is the potential for the error to persist through the entire process. The ward nurse often defers the identification check to the later redundant identification check performed by the radiology nurse. A great deal of process efficiency is lost here since a misidentified patient is more likely to be transferred before the error is caught by the process. We might speculate that this may have

evolved as a consequence of diffusion of responsibility (26, 49), where individuals are unaware of whether another redundant step will or has been taken, in this case, assuming that the check will be performed. This issue highlights the complex nature of redundancy – it appears that the existence of redundant checks across disparate locations and roles may have decreased the levels of compliance. In line with the conclusions suggested by Ong et al. (47), we advocate increasing compliance in the existing identification checks before introducing new ones.

Regarding the infection control scenario, the risk of infection is generated by violations in information flow, namely omissions and errors in the transfer form or verbal exchanges between the coordinator, porter and ward nurse. The results suggest that the individuals will take enough steps to ensure adequate infection control in three-quarters of patient transfers involving an infectious patient. Furthermore, the results suggest that communication errors and procedural errors are approximately equal in their contribution to the risk (in 10.4% of cases, the Porter does not have the correct information, and in another 12.8%, the Porter does not use adequate infection control despite having correct information), so changes at either point in the process would be appropriate. Since a substantial proportion of risk-associated trajectories occur when the correct information is present, we suggest that other factors such as inadequate understanding of the risks or time pressures may also contribute to the risk of infection.

B. Limitations

The model is not validated for the purpose of prediction – as such, we would urge caution in interpreting the overall likelihoods for risk as being definitive for the in-patient transfer process in the future, or for other hospitals. We feel this is reasonable because the purpose of the method is to analyse the breadth of possible risks in routine processes rather than to predict specific behaviours. A full validation of the model used in the case study would include further calibration for a number of scenarios (we have calibrated for only the infection control scenario due to limitations of the observed data) and testing against a new set of observations representing the same scenarios. This would demonstrate that the model is capable of predicting behaviour, and would thus provide evidence that it could be used as a tool for testing the effects of new organisational process interventions *in silico*.

When we implement the stochastic process that takes likelihoods from the observed study to inform the likelihood of an agent taking a particular action (or making a specific decision),

that stochastic process is independent of the simulation. This means that we have not captured the potential for individual errors to systematically occur together. For example, the model does not take into account whether a ward nurse is more or less likely to also check a patient's full name if he or she contemporaneously checks the identification band of that patient. Similarly, we expect variability amongst porters, which is not represented by the model we have constructed.

There is much in the way of contextual information that is not captured in regards to misidentification. For example, other information such as gender, ethnicity, other appearance-related information and prior contact may all be used to provide partial visual identification of a patient, which may be an informal and efficient method of identification that is already employed to augment the policy. In the case of infection control, other cues besides the information transmitted by the coordinator may influence a porter's infection control choices. Given that the simulated agents are able to perceive only the policy-based portion of information present in the real world, the model is likely to be a conservative estimate of the flow of risk.

C. General conclusions

The inter-ward patient transfer process produces a wide variety of possible behaviour trajectories. The proportion of simulations that were unique (considering different combinations of violations that occurred) was 97.4% for the misidentification scenario and 96.4% for the infection control scenario. This result is significant because it suggests that the presence of violations (in the form of workarounds, unintended violations and communication errors) increases the complexity of a work process, thus creating unique scenarios for which policy based on prescribed work practices cannot account.

This particular process reflects a well-defined workflow with numerous checkpoints and redundant operations that are included with the aim of minimising the potential for risk. Yet despite this, the potential for risk, either risk not mitigated by the process (in the misidentification scenario) or the risk generated by the process (in the infection control scenario) persists. Our analysis suggests that augmenting the formal process with additional procedures, of the type often suggested in the literature (50-52), may result in further unintended consequences (53, 54) due to the increased complexity of the procedure and the workarounds that evolve as a consequence. The complexity is increased by increasing the number of trajectory forks, thus increasing the number of possible trajectories. We suggest an

alternative– to reduce the complexity of the task and enhance the effectiveness of the remaining steps.

Using the in-patient transfer process as an example, we have demonstrated that an agent-based modelling approach to prospective risk analysis may assist in targeting policy changes according to the flow of risk in a process, complementing existing methods. The simulation framework presented here is likely to generalise to any work process for which (a) a set of policies or guidelines are defined in enough detail to permit the modelling of a prescribed work practice; (b) the process is observable such that it is possible to collect information about violations from repeated observations; and (c) involves multiple agents in the form of humans and the information systems that support them (be they verbal, written or software-driven). The method models the effect of combined violations as divergences from a prescribed work practice, indicating the likelihood of individual trajectories that may or may not lead to iatrogenic harm.

Acknowledgements

The authors acknowledge the support of the Australian Research Council (LP0775532) and the National Health and Medical Research Council (NHMRC) Program Grant 568612. The authors would also like to acknowledge two anonymous reviewers whose comments greatly improved this article.

References

1. de Vries EN, Ramrattan MA, Smorenburg SM, Gouma DJ, Boermeester MA. The incidence and nature of in-hospital adverse events: a systematic review. Qual Saf Health Care. 2008 Jun 1;17(3):216-23.
2. Zegers M, de Bruijne MC, Wagner C, Hoonhout LHF, Waaijman R, Smits M. Adverse events and potentially preventable deaths in Dutch hospitals: results of a retrospective patient record review study. Qual Saf Health Care. 2009 Aug 1;18(4):297-302.
3. Runciman WB, Moller J. Iatrogenic injury in Australia. Australian Patient Safety Foundation; 2001.
4. Westbrook JI, Braithwaite J, Georgiou A, Ampt A, Creswick N, Coiera E, et al. Multimethod Evaluation of Information and Communication Technologies in Health in the Context of Wicked Problems and Sociotechnical Theory. J Am Med Inform Assoc. 2007 Nov 1;14(6):746-55.

5. Iedema R. New approaches to researching patient safety. Soc Sci Med. 2009;69(12):1701-4.
6. Anderson JG, Ramanujam R, Hensel D, Anderson MM, Sirio CA. The need for organizational change in patient safety initiatives. Int J Med Inform. 2006;75(12):809-17.
7. Lee R, Cooke D, Richards M. A system analysis of a suboptimal surgical experience. Patient Saf Surg. 2009;3(1):1.
8. McKnight LK, Stetson PD, Bakken S, Curran C, Cimino JJ. Perceived information needs and communication difficulties of inpatient physicians and nurses. J Am Med Inform Assoc. 2002 Nov 1;9(90061):S64-9.
9. Parnes B, Fernald D, Quintela J, Araya-Guerra R, Westfall J, Harris D, et al. Stopping the error cascade: a report on ameliorators from the ASIPS collaborative. Qual Saf Health Care. 2007 Feb 1;16(1):12-6.
10. Anderson JG, Jay SJ, Anderson M, Hunt TJ. Evaluating the capability of information technology to prevent adverse drug events: a computer simulation approach. J Am Med Inform Assoc. 2002 Sep 1;9(5):479-90.
11. Leape LL, Brennan TA, Laird N, Lawthers AG, Localio AR, Barnes BA, et al. The nature of adverse events in hospitalized patients. N Engl J Med. 1991;324(6):377-84.
12. Wilson RM, Harrison BT, Gibberd RW, Hamilton JD. An analysis of the causes of adverse events from the Quality in Australian Health Care Study. Med J Aust. 1999;170(9):411-5.
13. Parker J, Coiera E. Improving clinical communication: a view from psychology. J Am Med Inform Assoc. 2000 Sep 1;7(5):453-61.
14. Woolf SH, Kuzel AJ, Dovey SM, Phillips RL, Jr. A string of mistakes: the importance of cascade analysis in describing, counting, and preventing medical errors. Ann Fam Med. 2004 Jul 1;2(4):317-26.
15. Reason J. Human error: models and management. BMJ. 2000 Mar 18;320(7237):768-70.
16. Rasmussen J, Nixon P, Warner F. Human error and the problem of causality in analysis of accidents [and discussion]. Philos Trans R Soc Lond B Biol Sci. 1990;327(1241):449-62.
17. Marx DA, Slonim AD. Assessing patient safety risk before the injury occurs: an introduction to sociotechnical probabilistic risk modelling in health care. Qual Saf Health Care. 2003 Dec 1;12(suppl 2):ii33-ii8.

18. van Leeuwen JF, Nauta MJ, de Kaste D, Odekerken-Rombouts YMCF, Oldenhof MT, Vredenbregt MJ, et al. Risk analysis by FMEA as an element of analytical validation. J Pharm Biomed Anal. 2009;50(5):1085-7.
19. Bonnabry P, Cingria L, Sadeghipour F, Ing H, Fonzo-Christe C, Pfister RE. Use of a systematic risk analysis method to improve safety in the production of paediatric parenteral nutrition solutions. Qual Saf Health Care. 2005 Apr 1;14(2):93-8.
20. Burgmeier J. Failure mode and effect analysis: an application in reducing risk in blood transfusion. Jt Comm J Qual Improv. 2002;28:331-9.
21. Robinson DL, Heigham M, Clark J. Using failure mode and effects analysis for safe administration of chemotherapy to hospitalized children with cancer. Jt Comm J Qual Patient Saf. 2006;32(3):161-6.
22. Chiozza ML, Ponzetti C. FMEA: A model for reducing medical errors. Clin Chim Acta. 2009;404(1):75-8.
23. Baird DR, Henry M, Liddell KG, Mitchell CM, Sneddon JG. Post-operative endophthalmitis: the application of hazard analysis critical control points (HACCP) to an infection control problem. J Hosp Infect. 2001;49(1):14-22.
24. Bonan B, Martelli N, Berhoune M, Maestroni M-L, Havard L, Prognon P. The application of hazard analysis and critical control points and risk management in the preparation of anti-cancer drugs. Int J Qual Health Care. 2009;21:44-50.
25. Griffith C, Obee P, Cooper R. The clinical application of hazard analysis critical control points (HACCP). Am J Infect Control. 2005;33(5):e39.
26. Wreathall J, Nemeth C. Assessing risk: the role of probabilistic risk assessment (PRA) in patient safety improvement. Qual Saf Health Care. 2004 Jun 1;13(3):206-12.
27. Garnerin P, Pellet-Meier B, Chopard P, Perneger T, Bonnabry P. Measuring human-error probabilities in drug preparation: a pilot simulation study. Eur J Clin Pharmacol. 2007;63(8):769-76.
28. Koppel R, Wetterneck T, Telles JL, Karsh B-T. Workarounds to barcode medication administration systems: Their occurrences, causes, and threats to patient safety. J Am Med Inform Assoc. 2008 Jul 1;15:408-23.
29. Spear SJ, Schmidhofer M. Ambiguity and workarounds as contributors to medical error. Ann Intern Med. 2005 Apr 19;142(8):627-30.
30. Wooldridge M, Jennings NR. Intelligent agents: Theory and practice. The Knowledge Engineering Review. 1995;10:115-52.

31. Franklin S, Graesser A. Is it an agent, or just a program?: a taxonomy for autonomous agents. In: Müller JP, Wooldridge MJ, Jennings NR, editors. Intelligent Agents III. Berlin: Springer-Verlag; 1997.
32. Sierhuis M, Clancey WJ, van Hoof RJJ. Brahms: A multi-agent modelling environment for simulating work processes and practices. International Journal of Simulation and Process Modelling. 2007;3(3):134-52.
33. Sierhuis M, Jonker C, Riemsdijk Bv, Hindriks K. Towards organization aware agent-based simulation. International Journal of Intelligent Control and Systems. 2009;14(1):62-76.
34. Anderson JR, Matessa M, Lebiere C. ACT-R: A theory of higher level cognition and its relation to visual attention. Hum-Comput Interact. 1997;12(4):439-62.
35. Altmann EM, Trafton JG. Memory for goals: An architectural perspective. Proceedings of the twenty first annual meeting of the Cognitive Science Society. Hillsdale, NJ: Erlbaum; 1999. p. 19-24.
36. Altmann EM, Trafton JG. Memory for goals: An activation-based model. Cognitive Science: A Multidisciplinary Journal. 2002;26(1):39 - 83.
37. Salvucci DD, Taatgen NA. Threaded cognition: An integrated theory of concurrent multitasking. Psychol Rev. 2008;115(1):101-30.
38. John BE, Kieras DE. The GOMS family of user interface analysis techniques: comparison and contrast. ACM Trans Comput-Hum Interact. 1996;3(4):320-51.
39. Hoot NR, LeBlanc LJ, Jones I, Levin SR, Zhou C, Gadd CS, et al. Forecasting emergency department crowding: A discrete event simulation. Ann Emerg Med. 2008;52(2):116-25.
40. Jun JB, Jacobson SH, Swisher JR. Application of discrete-event simulation in health care clinics: A survey. J Oper Res Soc. 1999;50(2):109-23.
41. Eldabi T, Paul RJ, Young T. Simulation modelling in healthcare: reviewing legacies and investigating futures. J Oper Res Soc. 2007;58(2):262-70.
42. Laskowski M, Mukhi S. Agent-based simulation of emergency departments with patient diversion. In: Weerasinghe D, editor. Electronic Healthcare. Berlin: Springer. pp. 25–37
43. Cardoen B, Demeulemeester E. Capacity of clinical pathways—a strategic multi-level evaluation tool. J Med Syst. 2008;32(6):443-52.
44. Catchpole K, Sellers R, Goldman A, McCulloch P, Hignett S. Patient handovers within the hospital: translating knowledge from motor racing to healthcare. Qual Saf Health Care. 2010 Aug 1;19(4):318-22.

45. NSW Health [Internet]. Correct Patient, Correct Procedure and Correct Site. 2007 October 30th [cited 2010 September 20th]; Available from: http://www.health.nsw.gov.au/policies/pd/2007/PD2007_079.html
46. NSW Health [Internet]. Hand Hygiene Policy. 2010 September 13th [cited 2010 September 20th]; Available from: http://www.health.nsw.gov.au/policies/pd/2010/PD2010_058.html
47. Ong M-S, Coiera E. Safety through redundancy: a case study of in-hospital patient transfers. Qual Saf Health Care. 2010:19:1-7.
48. Barlas Y. Formal aspects of model validity and validation in system dynamics. System Dynamics Review. 1996:12(3):183-210
49. West E. Organisational sources of safety and danger: sociological contributions to the study of adverse events. Qual Health Care. 2000 Jun 1;9(2):120-6.
50. Singh H, Naik A, Rao R, Petersen L. Reducing diagnostic errors through effective communication: Harnessing the power of information technology. J Gen Intern Med. 2008;23(4):489-94.
51. Kopec D, Kabir MH, Reinharth D, Rothschild O, Castiglione JA. Human errors in medical practice: Systematic classification and reduction with automated information systems. J Med Syst. 2003;27(4):297-313.
52. Bates DW, Gawande AA. Improving safety with information technology. N Engl J Med. 2003 Jun 19;348(25):2526-34.
53. Koppel R, Metlay JP, Cohen A, Abaluck B. Localio R, Kimmel SE, et al. Role of computerized physician order entry systems in facilitating medication errors. JAMA. 2005 Mar 9;293(10):1197-203.
54. Ash JS, Berg M, Coiera E. Some unintended consequences of information technology in health care: The nature of patient care information system-related errors. J Am Med Inform Assoc. 2004 Mar 1;11(2):104-12.

Figure and Table Captions

Figure 1. Processes are represented as a series of trajectories, where each opportunity for violation is indicated by a fork. A proportion of trajectories are indicated to diverge far enough away from the prescribed work practice to create the risk of an adverse event. Note that the likelihood of each trajectory is not indicated, and this is the focus of the proposed method.

Figure 2. A distribution of trajectories for the misidentification scenario indicates the gradual amelioration of risk for a patient with incorrect details. The boxes indicate the important steps in the policy of the patient transfer process, vertical arrows indicate chronology and horizontal channels indicate information transfer between the six agents. The shaded trajectories indicate the presence of risk at each point in the process, and percentages indicate the proportion of the completed simulations that are associated with risk along the given trajectory.

Figure 3. A distribution of trajectories for the infection control scenario indicates the increasing risk associated with lack of adequate infection control in the in-patient transfer process. The boxes indicate the important steps in the policy of the patient transfer process, vertical arrows indicate chronology and horizontal channels indicate information transfer between the six agents. The shaded trajectories indicate the presence of risk at each point in the process, and percentages indicate the proportion of the completed simulations that are associated with risk along the given trajectory.

Table 1. The number of simulations required to find rare trajectories (at least once) given a specific likelihood and assuming a binomial distribution (across simulations) where a success is defined by the presence of the end-of-chain risk.

Table 1

Likelihood of a trajectory	Representation in the simulation	Simulations required to find trajectory (95% confidence)
0.500	Reference likelihood	5
0.444	Observed level of inadequate infection control	6
0.100	Reference likelihood	28
0.010	Reference likelihood	298
5.98×10^{-3}	Precision chosen in experiments	500
0.001	Reference likelihood	2994

Figure 1.

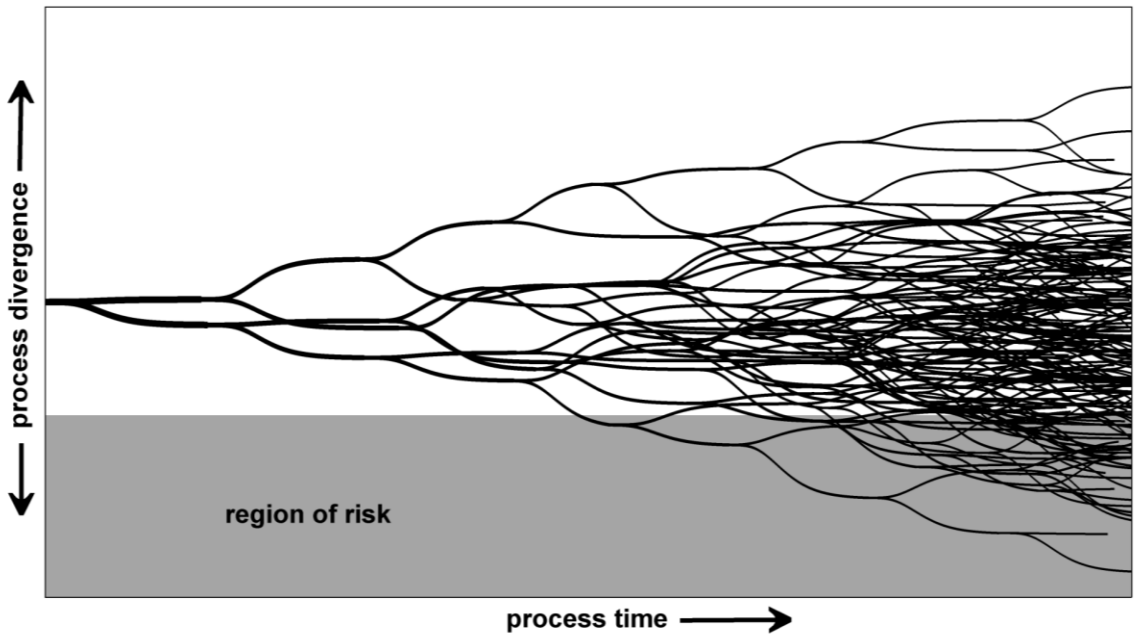

Figure 2.

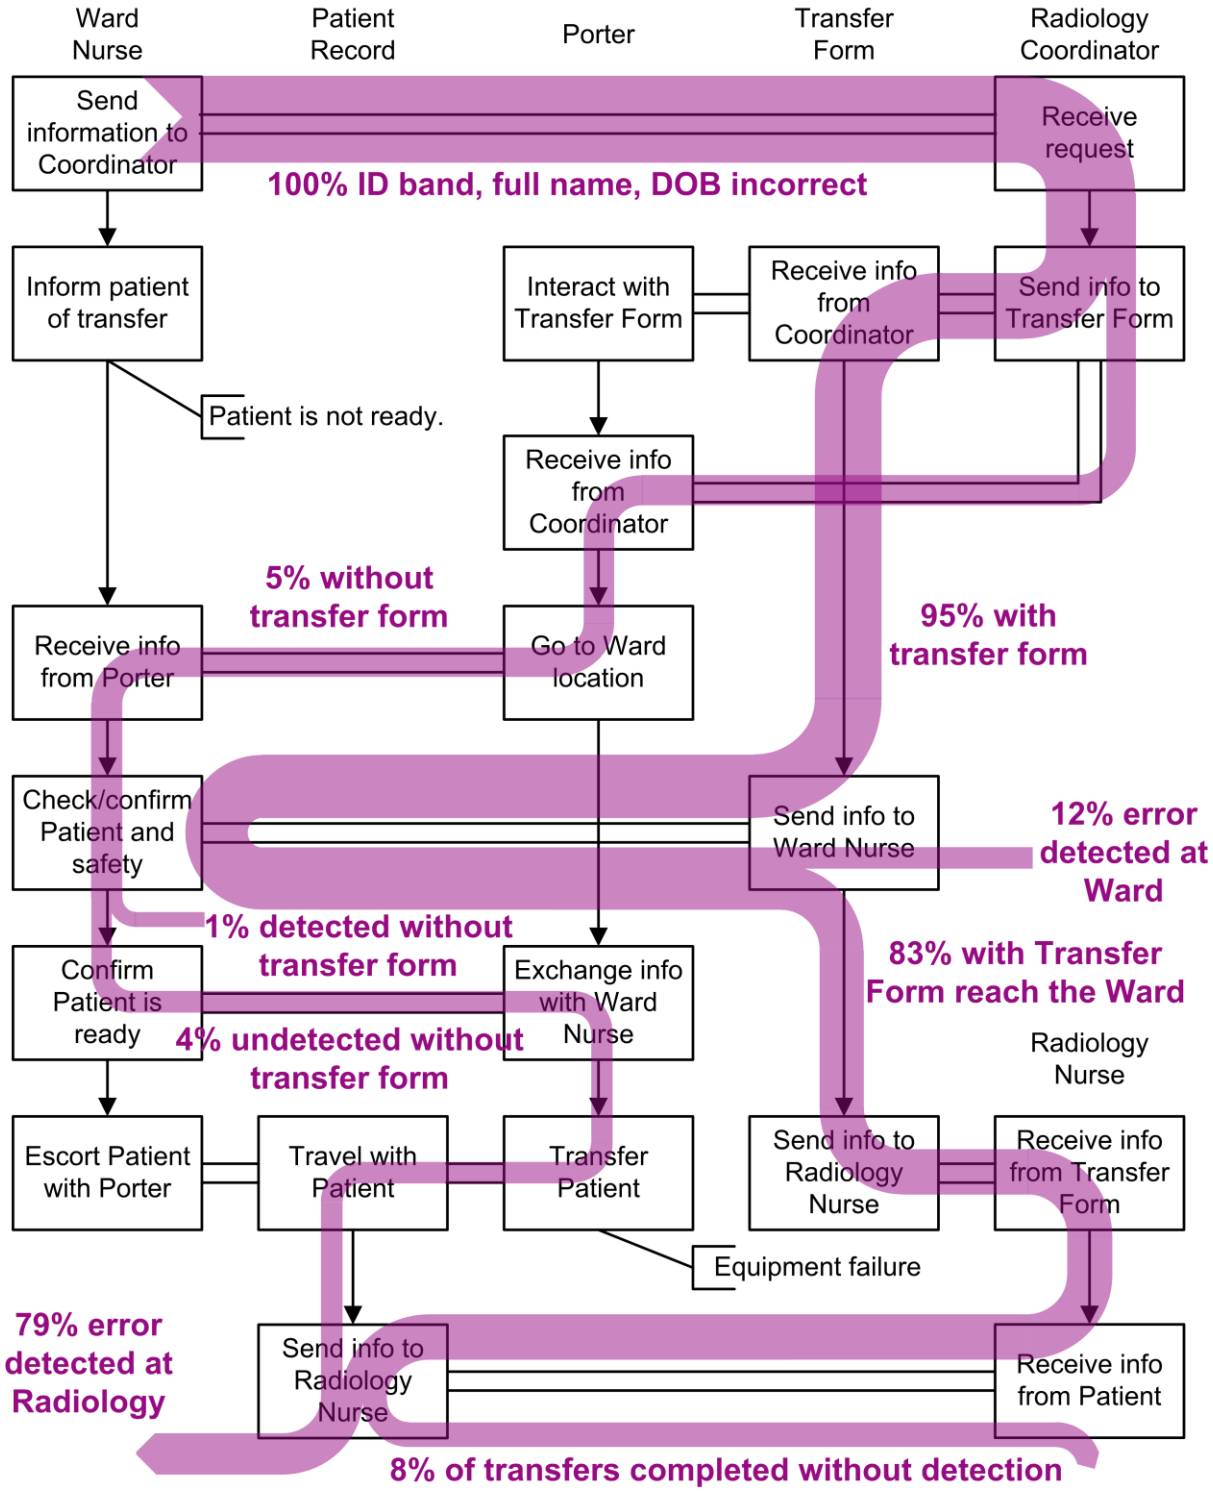
Ward Nurse
Patient Record
Porter
Transfer Form
Radiology Coordinator
Send information to Coordinator
Receive request
100% ID band, full name, DOB incorrect
Inform patient of transfer
Interact with Transfer Form
Receive info from Coordinator
Send info to Transfer Form
Patient is not ready.
Receive info from Coordinator
5% without transfer form
95% with transfer form
Receive info from Porter
Go to Ward location
Check/confirm Patient and safety
Send info to Ward Nurse
12% error detected at Ward
1% detected without transfer form
83% with Transfer Form reach the Ward
Confirm Patient is ready
Exchange info with Ward Nurse
4% undetected without transfer form
Radiology Nurse
Escort Patient with Porter
Travel with Patient
Transfer Patient
Send info to Radiology Nurse
Receive info from Transfer Form
Equipment failure
79% error detected at Radiology
Send info to Radiology Nurse
Receive info from Patient
8% of transfers completed without detection

Figure 3.

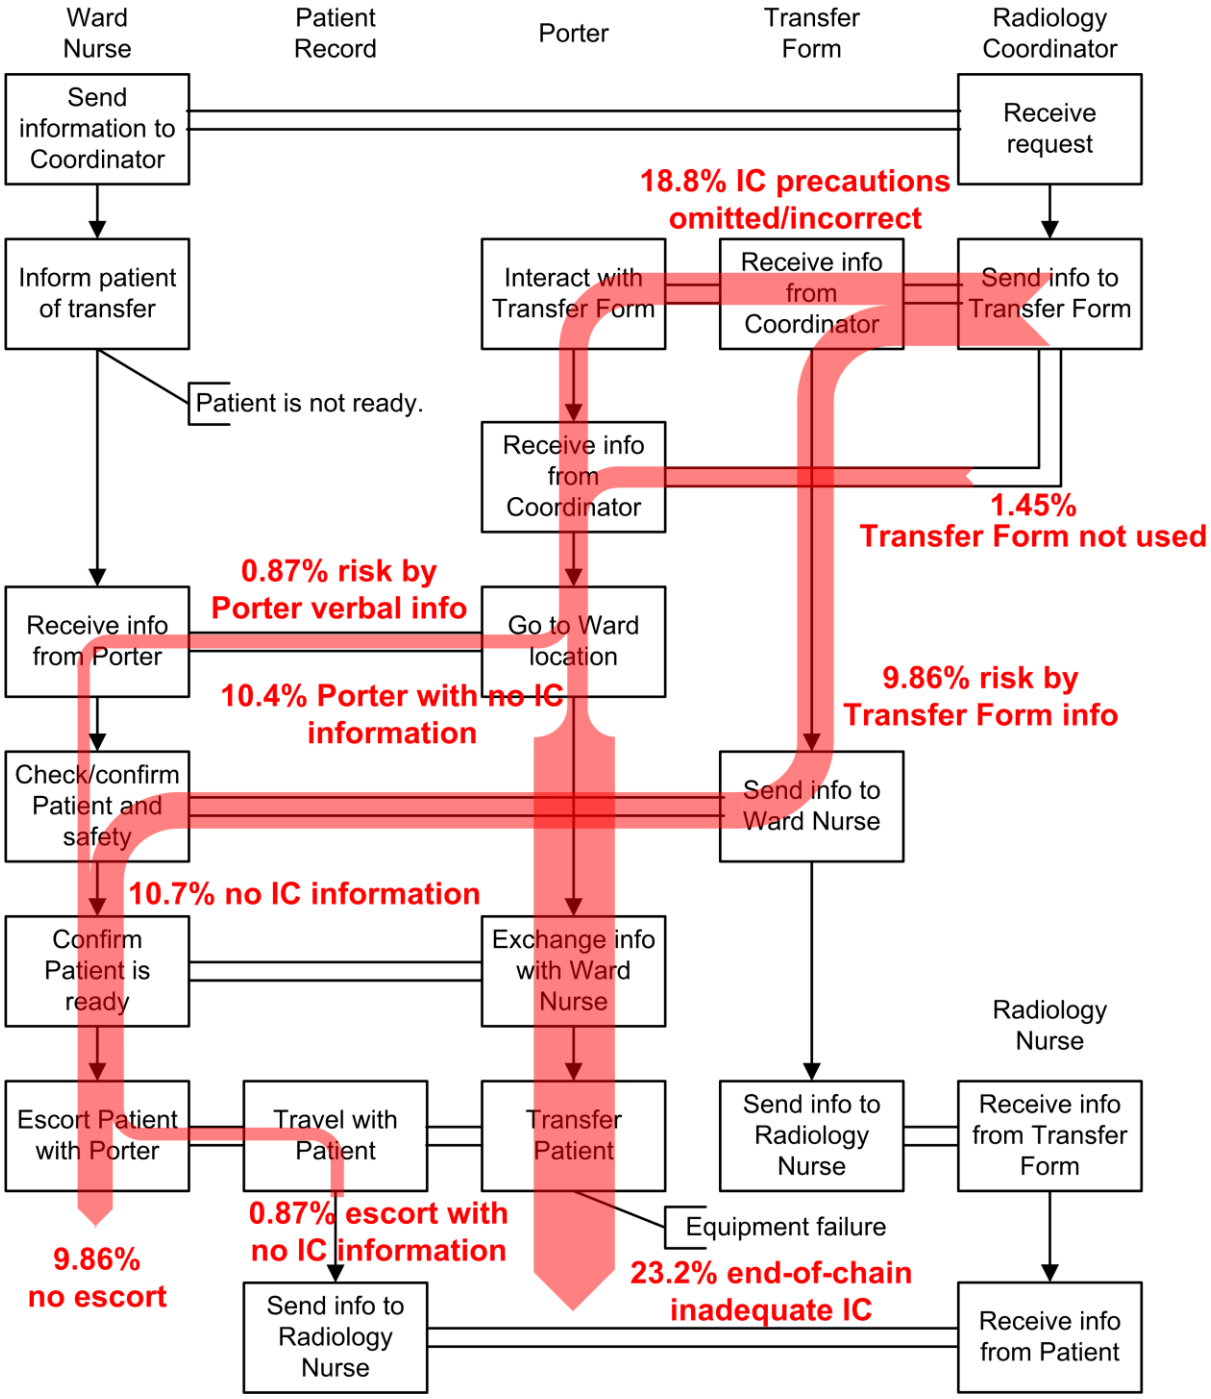
Ward Nurse
Patient Record
Porter
Transfer Form
Radiology Coordinator
Send information to Coordinator
Receive request
18.8% IC precautions omitted/incorrect
Inform patient of transfer
Interact with Transfer Form
Receive info from Coordinator
Send info to Transfer Form
Patient is not ready.
Receive info from Coordinator
1.45% Transfer Form not used
0.87% risk by Porter verbal info
Receive info from Porter
Go to Ward location
9.86% risk by Transfer Form info
10.4% Porter with no IC information
Check/confirm Patient and safety
Send info to Ward Nurse
10.7% no IC information
Confirm Patient is ready
Exchange info with Ward Nurse
Radiology Nurse
Escort Patient with Porter
Travel with Patient
Transfer Patient
Send info to Radiology Nurse
Receive info from Transfer Form
0.87% escort with no IC information
Equipment failure
9.86% no escort
23.2% end-of-chain inadequate IC
Send info to Radiology Nurse
Receive info from Patient

Appendix: examples of agent behaviour

The agent-based model of the in-patient transfer process comprises 4 human agents (the Porter, Ward Nurse, Radiology Nurse and Radiology Coordinator), 2 information system agents (the Patient Record and the Transfer Form), 6 objects and 186 activities. Abstract activities such as the one detailed below are implemented in Brahms to permit the repeated simulation of the process under two different scenarios. In one of those scenarios, the patient is infectious and the system is tested to examine the set of trajectories associated with the mitigation of infection risk via adequate infection control.

Like many other agent-based architectures, a Brahms agent holds a set of beliefs about its environment including the objects and other agents. An agent also holds a set of beliefs about itself, including its current location, and a set of other measurable beliefs. These may be propositional (true or false) or countable (an integer), and may also be unknown. Agents may communicate with other agents, pick up or put down objects. Agents may also modify their beliefs according to changes in the state of the environment, or according to requirements specified by their knowledge of prescribed work practices.

The example below is a model of a porter reading a transfer form while located in the ward. If the porter discovers that the patient is listed as infectious, then the porter will pick up the appropriate infection control, which is known to also be in the ward. The porter may perform this task at any time after arriving at the ward.

Agent: Porter **P**

Other agents and objects: Patient object **Q**, Transfer Form object **T**, Infection Control object **I**

Location: Ward **W**

Abstract activity: Use adequate infection control for infectious patients

Initial beliefs:
- **Q.infectious** = unknown
- **P.location** is **W** = true
- **I.location** is **W =** true
- **P** contains **T** = true
- **P** contains **I** = false

Plan:
- COMMUNICATE with **T** to receive **Q.infectious**

Case A: In the case where **T** communicates that **Q.infectious** = true:
1. SET BELIEF **Q.infectious** = true
2. GET **I**
3. SET BELIEF **P** contains **I** = true

Case B: In the case where **T** communicates that **Q.infectious** = false:
1. SET BELIEF **Q.infectious** to false

In a second example, the ward nurse is performing an identification check on the patient in the ward. This activity happens prior to the patient's transfer to radiology. In this case, the ward nurse is not required to complete these tasks in any particular order. For example, if the ward nurse identifies a difference between the patient's full name and the full name listed on the transfer form, then the ward nurse may immediately determine that the patient is misidentified (in relation to the transfer form) without needing to perform the check on the patient's ID band.

Agent: Ward Nurse **WN**

Other agents and objects: Patient object **Q**, Transfer Form object **T**

Location: Ward **W**

Abstract activity: Check that the patient's ID and full name matches the ID and full name on the transfer form

Initial beliefs:
Q.idband = unknown
Q.fullname = unknown
T.idband = unknown
T.fullname = unknown
WN.location is **W** = true
T.location is **W =** true
WN contains **T** = true
P contains **Q** = true
Q.misidentified = unknown

Plan:
COMMUNICATE with **T** to receive **T.idband**
COMMUNICATE with **T** to receive **T.fullname**
COMMUNICATE with **Q** to receive **Q.idband**
COMMUNICATE with **Q** to receive **Q.fullname**

Case A: In the case where **T.idband** does not equal **Q.idband**
SET BELIEF **Q.misidentified** = true

Case B: In the case where **T.fullname** does not equal **Q.fullname**
SET BELIEF **Q.misidentifed** = true

Case C: In the case where **T.fullname** equals **Q.fullname** and **T.idband** equals **Q.idband** SET BELIEF **Q.misidentified** = false