

Improving Strategic Management of Hospitals: Addressing Functional Interdependencies within Medical Care

Dr. Masanori Akiyama

Daniel Goldsmith

Dr. Michael Siegel

Sloan School of Management

Massachusetts Institute of Technology

Center for Digital Business, Room NE20-336

Cambridge, MA 02142

Voice: (617) 253-2937

Fax: (617) 258-7579

{poas, goldsmith, msiegel} @ mit.edu

Abstract

To better understand the performance of hospital operations in response to IT-enabled improvement, we report the results of a system dynamics model designed to improve core medical processes. Utilizing system dynamics modeling and emerging HIS data, we demonstrate how current behavior within the hospital leads to a 'stove-pipe' effect, in which each functional group employs policies that are rational at the group level, but that lead to inefficiencies at the hospital level. We recommend management improvements in both materials and staff utilization to address the stove-pipe effect, and estimate the resultant cost-saving. We believe that the major gains in health information systems use will accompany new information gathering capabilities, as these capabilities result in collections of data that can be used to greatly improve patient safety, hospital operations, and medical decision support.

Introduction

This work is designed to better understand how hospitals respond and adapt to the introduction of Health Information Systems (HIS). The current work extends previous hospital research on nurses and the injection process to include a wider set of stakeholders. This work examines the entire supply chain of patient injections with consideration ward, pharmacy and doctor impacts on operations. The motivation for this research stems from the observation that while the need for new HIS in hospital environments has been well documented, system managers, as well as medical practitioners, have both recorded their disappointment with many HIS implementations (Kohn, 2000).

Part of the reason for suboptimal performance is that many approaches to HIS fail to take full advantage of the new opportunities provided by data collections systems as a tool to: a) understand, measure, and track hospital operations, b) identify and implement high-leverage improvements, and c) provide opportunities for hospital staff to train and learn

more effectively. Our research suggests avenues to utilize the rich data set provided by HIS to improve hospital efficiency, patient safety, and the receptiveness of staff to IT enabled-improvements. Comprehensive HIS allow for the development of new tools to support effective integration and success operation of new systems.

This research draws from the analysis of the HIS in place at several major Japanese hospitals, named POAS (Point of Act System). As described by Akiyama (2001, 2007), the underlying concept of POAS is to enable records of “‘who did what to whom, where, when, using what, and for what reason.’ In short, real-time input becomes possible at the point of action.” Under the system, logs of medical actions and inventories are created throughout the course of operations, recorded using bar-code scanning technology and nurses equipped with PDAS (personal digital assistants). The system operates continuously at the hospital, handling 100 transactions per second, or more than 360,000 transactions per hour, and has been in continuous operation for four years. For example, the system collects information on every interaction between order, drug, nurse, and patient. Utilizing this data, we can revisit the challenges associated with HIS and understand system wide behavior.

As described by Akiyama, Siegel, and Goldsmith (2007), soon after implementation, POAS facilitated vast improvement in multiple areas of hospital operations, with estimated savings reaching millions of dollars. In addition to POAS-enabled cost savings, the system also led to impressive improvement in patient safety. Prior to the implementation of POAS, there existed nearly a 40 percent chance that there would be a misadministration of an injection prescription, due to the absence of an automated method of checking injections and the lack of real-time communication. After POAS, this percentage was cut dramatically; an alarm would sound prior to the injection if any problems existed (such as a correct patient being presented with an incorrect medication), and the staff would be able to fix the mistake prior to injection. In the years following the initial implementation of POAS, patient-safety benefits continued to be realized, and by all measures, improvement remained robust. However, concern was raised about the sustainability of the system’s financial performance. System managers were concerned about how to obtain further improvements in of the hospital’s financial performance. Of particular concern were the areas of overlap between functional groups within the hospital. For example, for a patient to receive an injection, information and materials must be effectively shared by doctors, pharmacists, and nurses. The ability for system managers to help manage these interactions was thought to be a key determinant to overall system efficiency, as measured in staff and materials utilization. Further, POAS improvements had created a rich-data set of operational data that was being underutilized in hospital improvement. We combine analysis of POAS data sets with site visits to POAS hospitals, systems dynamics modeling, and expert input to derive important process changes across the hospital system.

The Silo Effect in Hospital Ward Management

The basic injection process provides a useful way of thinking about the challenges associated with hospital ward management. The injection process at POAS hospitals,

also described in Akiyama, Siegel, and Goldsmith (2007), refers to the different paths an injection order can take, culminating with either in a successful injection or in a changed or cancelled order.

The key actors in the injection process are doctors, pharmacists, and nurses, and Figure 1 presents a ward view of operations. This figure is a conceptual simplification of actual hospital operations, but highlights each actor’s key procedures. The normal path of injection goes from left to right: a doctor issues a prescription order, the pharmacist package the set of drugs required (referred to as an Rp) and checks the order for correctness then the nursing staff mixes the Rp components together and injects it into the patient. The solid arrows below each actor’s name signal the observables for each actor: for example, pharmacists rarely have any information about the downstream processes of the nursing staff. Their viewpoint is restricted to their core operations, preparing and checking orders. In other words, the current injection process creates a “stovepipe effect,” by which each key actor operates within a functional silo. While hospital managers may view the entire process, doctors, nurses, and pharmacists are all bounded in visibility by the specific breadth of their function.

Also of note in Figure 1 are the flows of information and materials, as represented by the dashed arrows. Information and materials can flow downstream—as in the normal injection process described above—but it can also flow upstream, as in the case of changed orders. In upstream operations, a nurse will check with the doctors to see if the order has been changed prior to injection. If the order change is processed prior to the mixing phase, the Rp components can be returned to stock by the pharmacists and are generally reusable for future patient orders. It is important to note that if the Rp has already been mixed prior to the change order, the change must result in the disposal of the Rp.

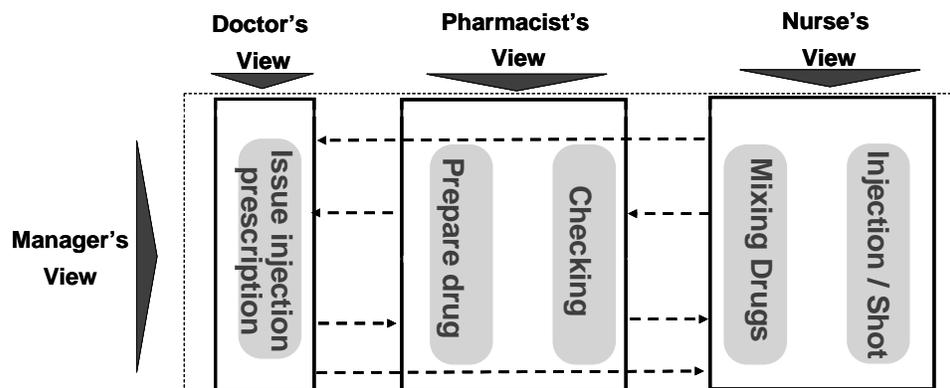


Figure 1. Differing Views of Hospital Operations

The difference between orders that can be reused and those that must be thrown out has been shown to cause significant variations in the amount the hospital spends on drug inventory, as well as the efficiency with which the staff processes an order. In both instances change orders result in inefficiencies in the system. The “Manager’s View” sees the effects across the silos and can better understand these “global” inefficiencies.

In this research, we highlight the central role of pharmacists in efficient management of the injection process, and examine the ability for pharmacists to manage operations given their limited visibility of overall operations. We show that pharmacists have developed self-defeating policies as a result of the stove-pipe effects that maximize downstream flows while ignoring the consequences of upstream flows. We then calculate the costs of these policies, and recommend a series of policy interventions to ameliorate the negative impacts.

Pharmacy Operations

In the section below we present the modeling formulations of ward operations that have emerged from our research. The theory reported here draws on extant literature in process improvement, system dynamics modeling, as well as our study of one hospital setting (Forrester, 1958; Repenning & Sterman, 2002). In particular, we unpack pharmacy operations as a balancing process, by which pharmacists must balance competing demands for their time. Ultimately, we see that pharmacists pursue strategies that they erroneously believe reduce their work burden and lead to waste in medication. While the consequences of pharmacist's actions are difficult to see from their perspective, given their limited visibility of other operations, and lead to significant downstream inefficiency and may even impact patient safety.

Despite these negative consequences, we demonstrate that the current strategies pharmacists employ to manage their work load are *intendedly rational*. The problem is that they are developed from a flawed mental model of hospital operations. Instead of faulting pharmacists for the negative outcomes we observed, we show how structural limitations in hospital operations create an environment in which well intentioned policies at the group level (i.e. pharmacists) can subvert overall operational efficiency. To arrive at this insight, we present a series of causal diagrams to capture hospital dynamics, construct a simulation model to quantify system behavior over time, and simulate the model, presenting the results from the perspective of multiple hospital actors.

The core pharmacy operations are represented in the balancing loop B1 in Figure 2. The rate of new orders from doctors increases the work that the pharmacists must complete. Depending on the delay in pharmacy, the orders will be processed at a given rate: the shorter the delay, all else equal, the rate of operations will increase. The result of pharmacy operations, preparing and checking drugs, here referred to as filling orders, closes a loop by diminishing the amount of work left to do. The “work accomplishment” structure in B1 captures the embedded rationality of pharmacists: their goal is to reduce the amount of work left to accomplishment.

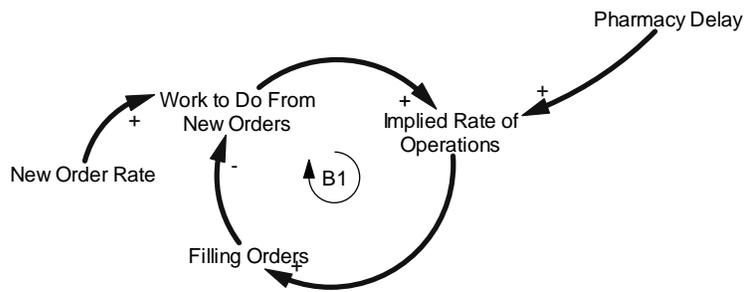


Figure 2. Intended Rational of Pharmacy Operations

Figure 3 presents a constraint on the rate of hospital operations: the capacity of pharmacy, a concept that combines staff and equipment into an overall metric of how fast, at its maximum, the pharmacy would be able to process orders. First, the capacity of the pharmacy is represented as a fundamental constraint on operations, shown in loop B2. The capacity of the pharmacy sets a feasible rate of operations, given an amount of work to do, representing the upper limit of pharmacy operations. The capacity of the pharmacy also has a role in determining the normal pharmacy delay: the time in which, on average, it takes the pharmacy to fill an order. While not the only influence of pharmacy delay (which include such other factors as work habits and expectations), the capacity of the pharmacy influences the pharmacy delay (another way to think of the pharmacy delay is the average time an order will spend within the purview of the pharmacy.)

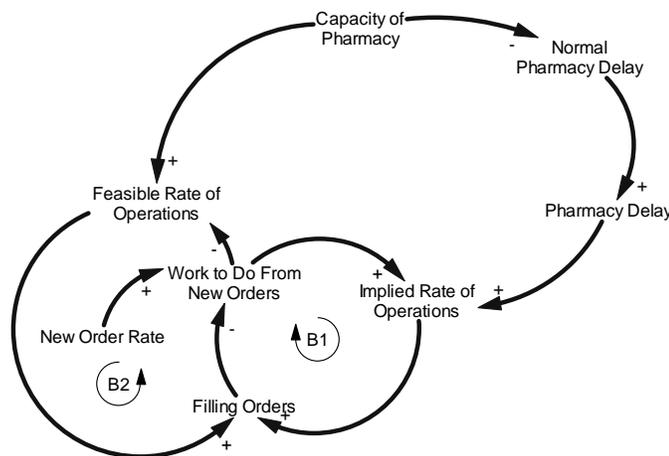


Figure 3. Capacity Constraints

To examine these basic pharmacy dynamics in the context of a sample hospital, we first consider data from POAS on the changes in work burden on pharmacists over time. For several reasons, including structural changes in the Japanese healthcare environment, both the number of patients and the number of orders rose over a period of several years. (Figure 4) Orders, however, rose faster than patients, and the increase created a rise in

the orders per patient. During the same time the overall staff in the pharmacy was constant. Therefore, pharmacists were faced with increasing work burdens.

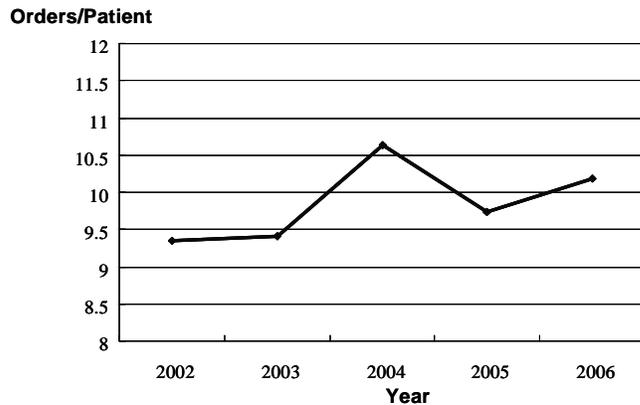


Figure 4. Changes Order per Patient Over Time

As a result of the changes in the new order rate over time, the pressure on pharmacists increases. We have captured the concept of work pressure, which can be thought of a ratio of implied rate of operations divided by the feasible rate of operations. That is, the closer a pharmacist gets to the maximum feasible rate he or she can work, the more pressure they feel.

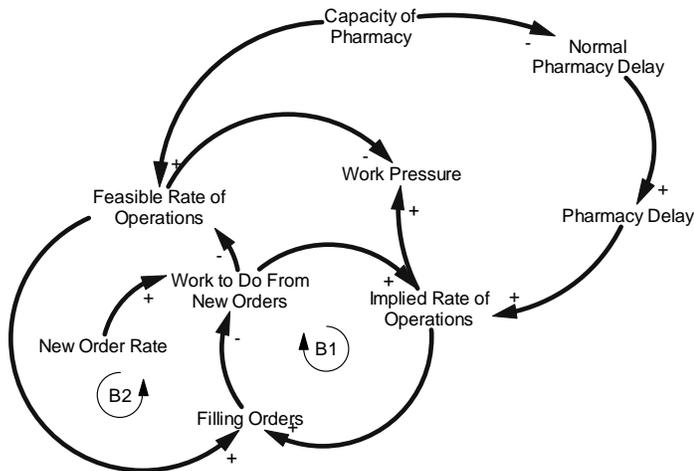


Figure 5. Work Pressure

We represent the rise in orders shown in Figure 6 below with an arrow, and have circled several possible options available to managers and pharmacists to manage the increased work burden. The three options are to: a) increase the capacity of the pharmacy, by, for example, hiring more staff; b) lengthen the normal pharmacy delay, likely requiring patients to wait longer for orders to be complete; and c) increase the implied rate of operations by finding some way to work faster with the same capacity.

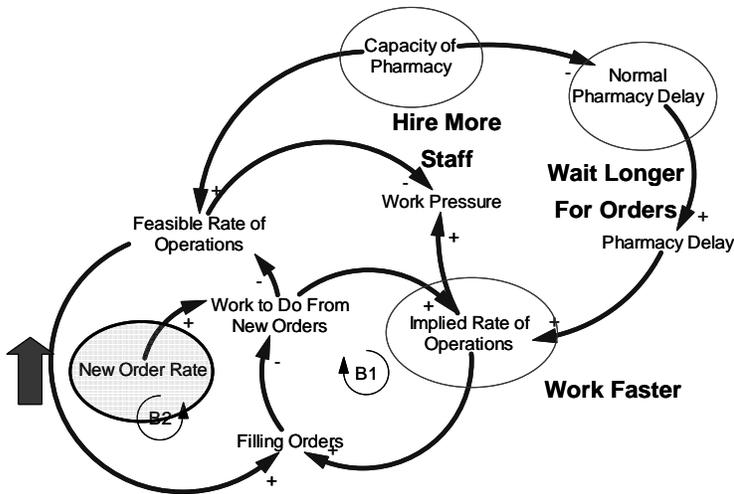


Figure 6. Options to Manage Increases in Work Pressure

While all the options conceptually reduce work pressure, both hiring more staff and requiring a longer wait for orders are likely difficult to accomplish and may be, in fact, infeasible. Increasing the work rate, therefore, is the most appealing option. Further, it is under control of the pharmacists themselves, and allows them regulate work pressure without the intervention of other actors. While managers must become involved to hire more staff, and nurses and patients will be affected by longer delays, pharmacists, by developing a way to work faster, can balance the work burden with little perceived costs. One way to accomplish this is to complete orders in batches, an operation represented by loop B3. (Figure 7)

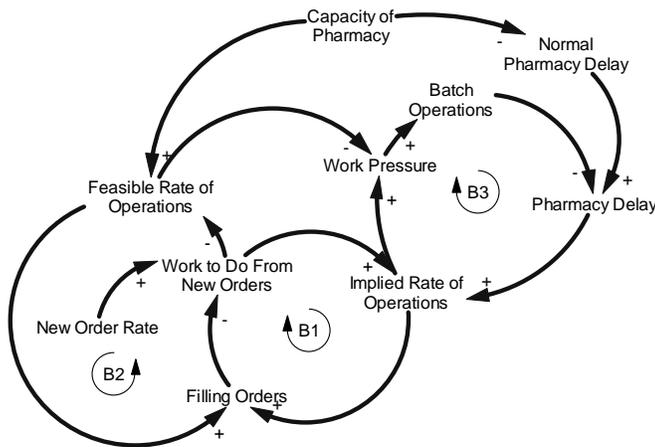


Figure 7. Batching

Batch processing combines orders into batches which are processed in delivered in large groups and allow pharmacists to increase the rate at which orders are processed. While new orders represent the majority of the work pharmacists must accomplish, order changes also creates work for pharmacists to accomplish. The rate of order changes also creates another stock of work for pharmacists to accomplish. The demands from order

changes, however, creates competing demands, represent by loops B4 and R1. Working to fill order changes reduces the pile of order change, but reduces the rate at which new orders can be processed. From interviews with pharmacists conducted at a sample hospital, the pharmacists were clear that they considered new orders the most important to fill, and the backlog of new orders to be the most salient indicator of how effective they were in managing their work load. Pharmacists would likely pursue policies to favor loop R1 over loop B4. (Figure 8)

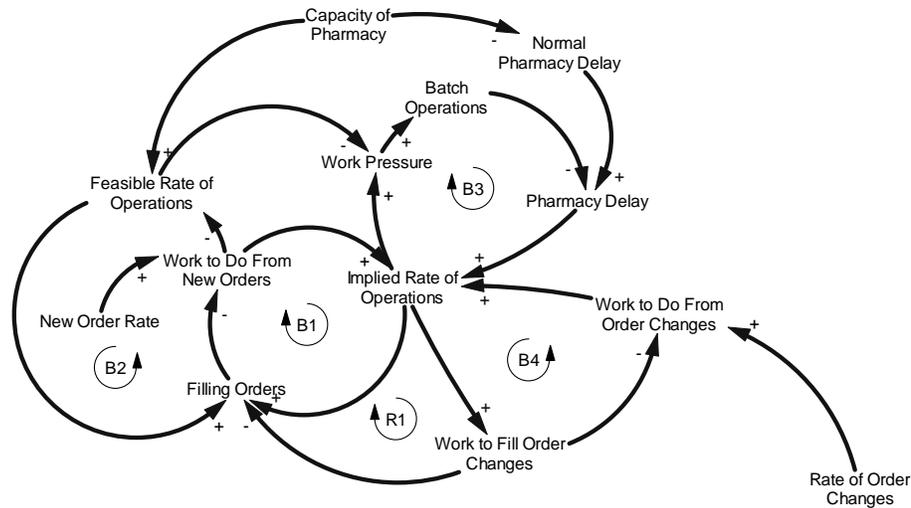


Figure 8. Balancing Work Pressures

Also, that batching operations in loop B3 also have a virtuous property of reducing work from redos, as shown in loop B5 of Figure 9. Batching reduces the average time the order stays in the pharmacy by speeding up the rate at which they fill orders. As orders spend less time in the pharmacy, fewer changes are, on average, made while the order is at the pharmacy, reducing the amount of work pharmacists must do from order changes in the pharmacy. Batching, therefore, appears to be not only an effective strategy to manage new orders, it also reduces the amount of work from redos. A pharmacist commented on this, effect saying “We would like to get redos in the pharmacy down to zero.” (Figure 9)

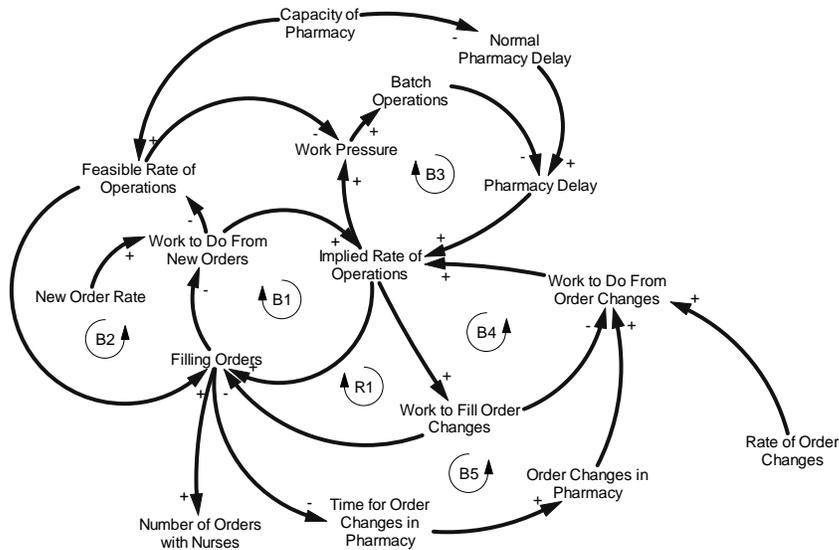


Figure 9. Reducing Changes at Pharmacy

Based on the dynamics demonstrated in the causal diagrams above, we have constructed a simulation model (shown its entirety in Appendix 1). We first simulated the model from the *pharmacist's perspective* for a one-week period, and present the results in time series graphs in Figure 10. We see the positive benefits on batching from the pharmacy perspective. The base case (no batching) is shown in solid lines, and the batching case is shown in dashed lines. As a result of batching, work to do goes down, as the completion rate goes up, and the numbers of redos that occur in the pharmacy also fall.

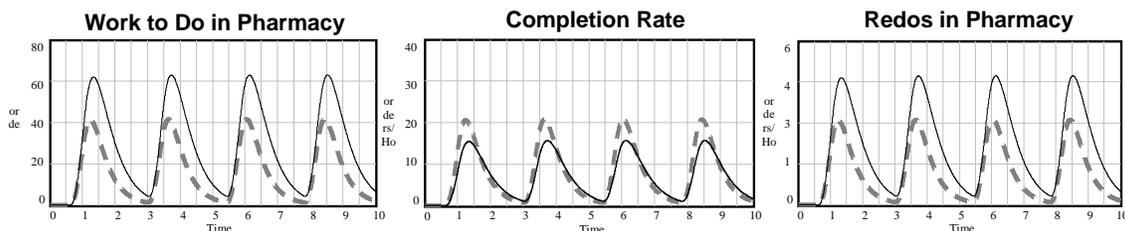


Figure 10. Simulation of Batching in Pharmacy

Conceptually the mental model that leads to the batching strategy is limited. The results presented in figure x are skewed by this partial view of possible outcomes, rather than a whole-systems perspective. Figure 11 explicitly includes feedback from the nurses, which is not represented in the pharmacist's conceptual model of operations. Loops R2a and R2b show the key feedback from redos at the nursing station; in both loops the redos that the pharmacists believe they are avoiding by batching reappear. Pushing orders faster out of the pharmacy does reduce redos orders *while they are in the pharmacy*, but only by delaying the same redos until the nurses are working on them. Loop R2a shows this feedback, with the key delay for these orders to return to the pharmacy explicitly recognized by the double line across the causal link. In other words, the redos from the nurses are the very same redos the pharmacy would have processed if the order had stayed in the pharmacy longer (even though pharmacists don't see this.)

Loop R2b shows how this feedback is observed from the pharmacy perspective. The redos from the pharmacy are not treated as redo orders; instead, when they appear from the nurses, they *seem* to be new orders. Pharmacists, because of how the orders are processed, lose the traceability of ordering that would let them match up the subsequent order change with the new order. The perceptual differences between loops R2a and R2b, which correspond to different mental models of pharmacy operations, are crucial to understanding pharmacy behavior.

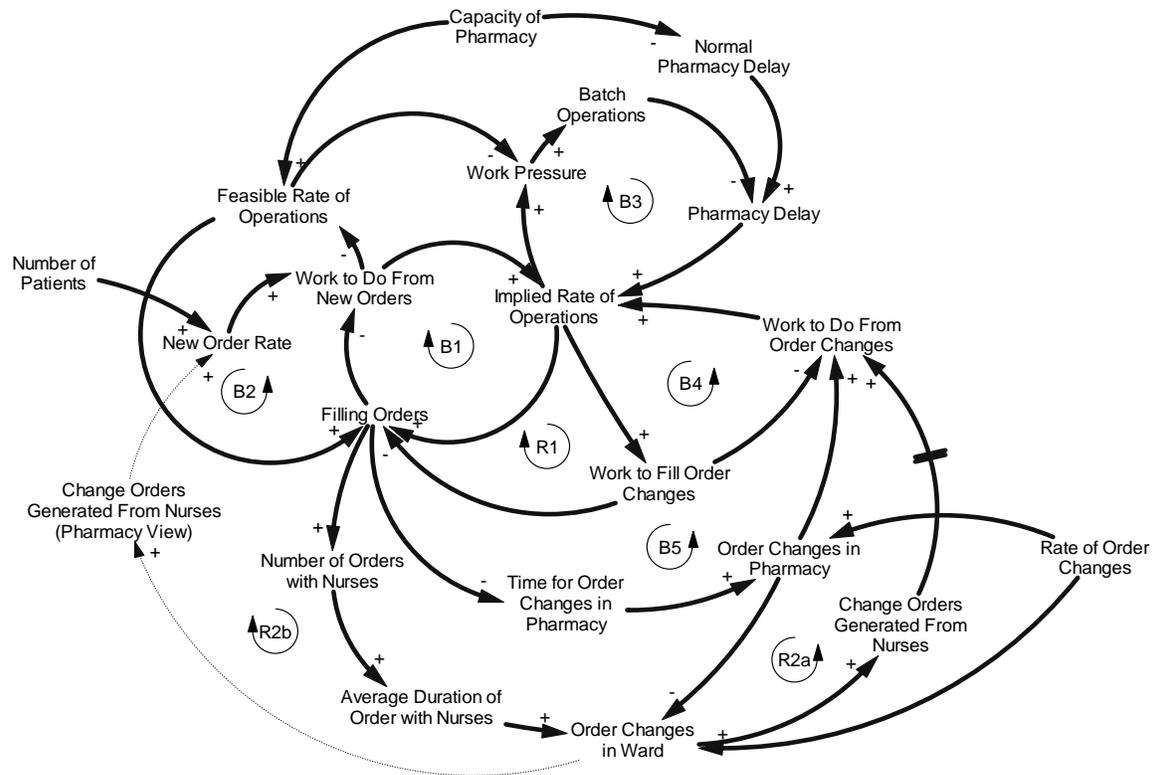


Figure 11. Batching Reduces Changes at Pharmacy

The Role of Pharmacy Operations on Downstream Processes

We now include nurses into our causal framework and observe the downstream consequences of batching. To do so, we return to the possible outcomes of an injection order; it can be a) successfully injected, b) changed before mixing while at the pharmacy, c) changed before mixing while at the nurse station, or d) changed after mixing at the bedside. There are important differences between outcomes: most importantly, once the order is mixed, it must be wasted, as it can't be used for other patients. Additionally, an order returned from the nurse station has accumulated more time being processed than one changed at the pharmacy. These differences are demonstrated in Figure 12, a stock and flow structure that captures the four different outcomes (shown as outflows from each stock.) The first series of constructs relate to the flow of orders and material. The flows, denoted by straight arrows with values, are the rate at which orders for injections

are successfully moved (referred to as Rps, a bundled collection of injection materials) between stations in the hospital. Figure 12 also shows three of the system's stocks, denoted by a rectangle, which are computed as the integration of the stock's inflows less its outflows. The stocks are the accumulation of orders and waiting to be processed at three stages, the pharmacy, the nurse station, and the patient's bedside.

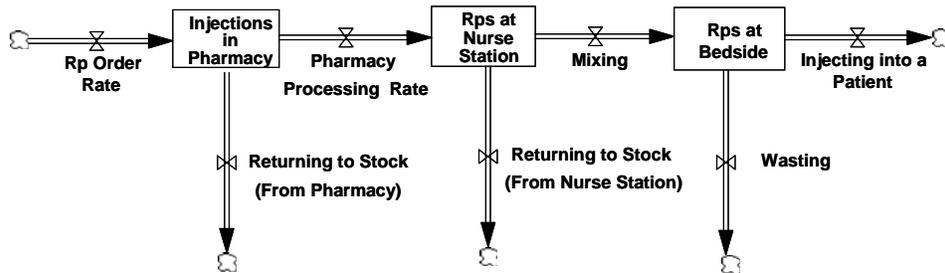


Figure 12. Batching Reduces Changes at Pharmacy

In Figure 13 below, the consequence of batching are captured with arrows indicating the implied changes in the variables; beginning with the pharmacy processing rate, as the rate increases, injections in the pharmacy decrease, as do returns from the pharmacy. At the same time, Rps at the nurse station rise (given a constant mixing rate). The dashed rectangle shows the limits of the pharmacy's perspective.

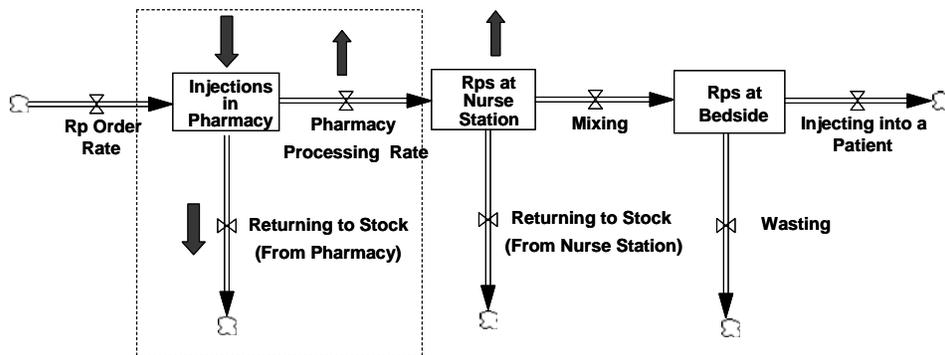


Figure 13. Batching Reduces Changes at Pharmacy

Next, we capture the perspective of the nurses in Figure 14, and again show the implied behavior with arrows. It has been demonstrated in previous research that nurses are likely to respond to rising orders at the nurse station also with a batching strategy.¹

¹ The insights shown are also reflected as additions to the causal diagram in Appendix 2.

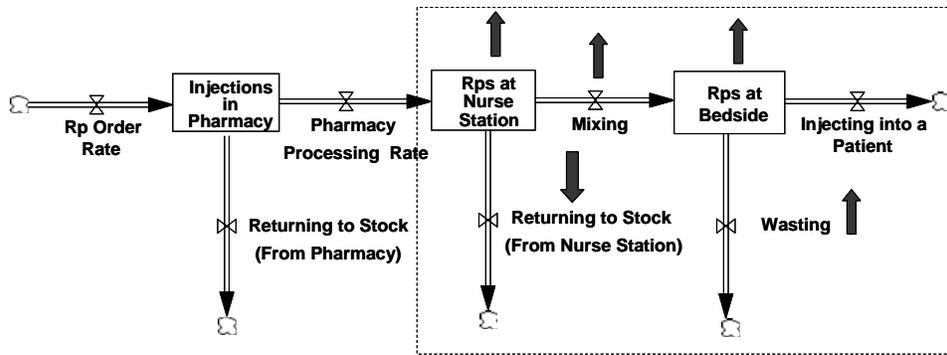


Figure 14. Batching Reduces Changes at Pharmacy

As intended under POAS, each Rp is scheduled in advance to be mixed at a specific time, which corresponds to the staggered schedule of each injection. However, nurses often mix Rps in large batches throughout the day, which clusters the nurse’s workload and increases the blocks of available downtime. This balances the increase in the pharmacy processing rate, and accelerates the rate at which Rps are moved from the nurse station to the bedside (i.e., in the mixed form). This however, leads more change orders to occur while the orders are mixed, and these orders must be thrown out.

Simulating our model from the nurses perspective, we see a very different outcome as a result of pharmacy batching, as shown in Figure 15. The base case is the solid red line, and the batching case is shown in solid blue line. Instead of the improvements we witnessed from the pharmacy perspective, pharmacy batching has lead to an increase in the nurse work rate, an increase in the number of orders wasted, and lead to an overall increase in the time that staff must spend on order changes.

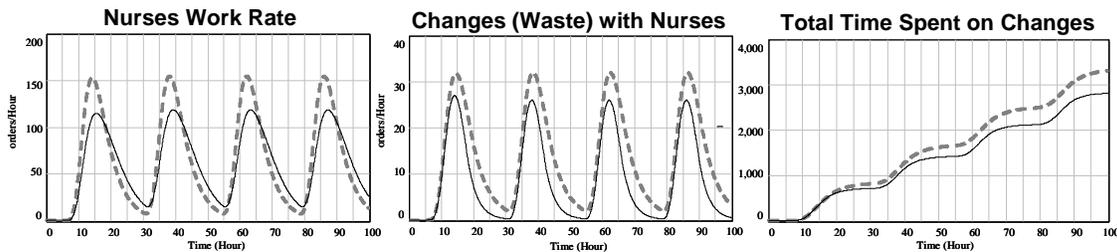


Figure 15. Simulation results from the nurse’s perspective

We support our simulations by using POAS data, and examine the behavior our simulation suggests for a one-week period over a period of years. The amount of orders that are either wasted or returned increases after initial gains were made in from 2002 to 2003. The percentage of orders wasted rose from 1.2 percent in October 2003, to 1.4 percent in October 2006. Additionally, the number of orders returned from nurses increased from 8.9 percent to 10.4 percent over the same period. (Figure 16)

Number of Orders

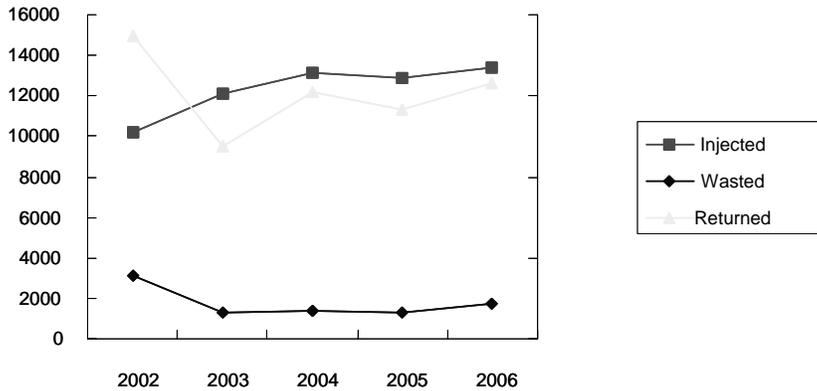


Figure 16. Injection Outcomes

POAS data also verifies the relationship demonstrated in the simulation results shown above. By relating the mixing gap, the time between when an order is mixed and when it is injected, to the percentage of orders wasted, we see a positive relationship between increase in the mixing gap and increases in the percentage of orders that are wasted. (Figure 17)

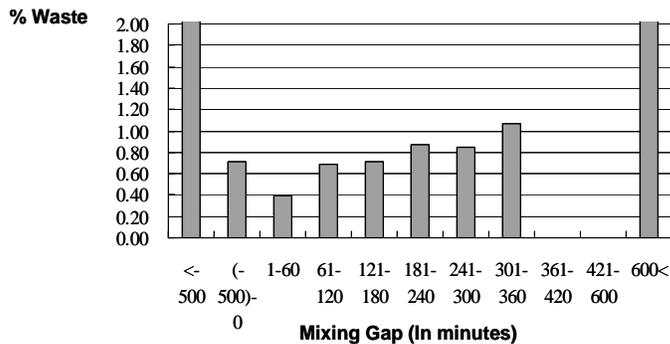


Figure 17. Mixing Gap and Percentage Waste

In addition, POAS data shows a negative mixing gap for a small segment of orders. That is, the orders are shown to be mixed after they have already been injected—a practical impossibility. This data suggests evidence for an additional phenomenon in nurse behavior that may result from pharmacy batching: early injection. Because nurses: a) have the orders earlier, and b) are mixing them earlier, there is evidence to suggest that nurses are injected them prior to the scheduled injection time. Some of these orders are then changed, meaning patients received the wrong injection. Examining POAS data we see that orders are injected early nearly 3.3 percent of the time. While the consequences of early injection require further investigation, the evidence suggests it is at least possible that pharmacy processes are having unintended downstream consequences on patient safety. Further, because of the stove-pipe effect, pharmacists, without policy intervention, would fail to realize the magnitude of consequences of their actions.

Recommendations

Our analysis supports a recommendation to not only reduce batching in the pharmacy, but to actually increase the pharmacy delay for filling orders. Figure 18 shows three simulation results: a) the base case (solid line); b) the batching case shown above (dashed line); and c) the increased pharmacy delay case (dashed-dot line). In the increase pharmacy delay case, orders spend, on average 50 percent longer in the pharmacy, but doing so creates a meaningful reductions in cumulative work to do for order changes and order returns for nurses (as more changes are caught in the pharmacy before additional nurse work must be done.)

Further, we see opportunities to utilize data in pharmacy operations. For example, POAS data indicates that a small set of medications disproportionately account for waste. Using feedback from data to schedule operations (e.g., filling order, missing), and increasing the delay in the pharmacy for these select medications, would produce significant savings in materials and staff utilization. For example, by having the pharmacists hold the top 5 drugs, which account for nearly 25 percent of waste, as long as possible and avoid early mixing by the nurses, we estimate potential savings of approximately 70 million yen, or 600 thousand US dollars, on an annual basis. Further improvements could be made by dynamically linking operational data to pharmacy operations.

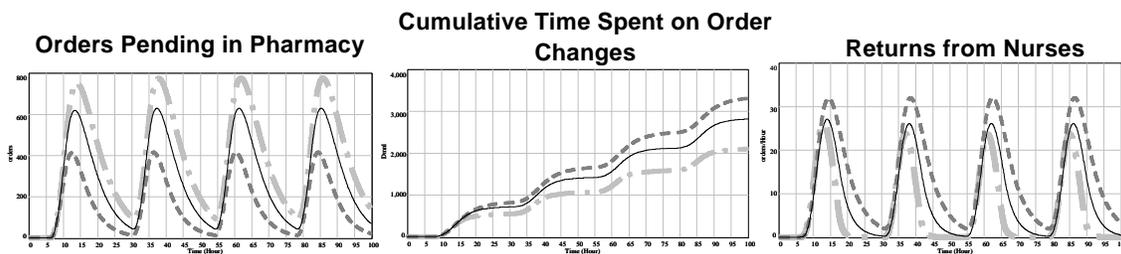


Figure 18. Increasing Pharmacy Delay Improves Metrics

Conclusion and Discussion

The goal of research in this area is to develop management improvements by means of systems modeling and analysis combined with the use of newly available operational data sources, such as POAS. We expect that this style of research will help sustain and advance system-wide improvements in operational efficiency in a hospital setting. Ongoing investigations will focus on designing robust experiments and simulations to increase understanding of the causes and effects for sustained improvement. While implementation challenges to this process improvement exist, we have proposed a high-leverage yet non-disruptive improvement to quickly demonstrate financial benefits. Continued studies, including additional staff interviews, will focus on better understanding the culture involved in ward management to develop a workable solution for more comprehensive improvements in ward management.

Utilizing system dynamics modeling and emerging HIS data, we have demonstrated how current behavior within the hospital leads to a ‘stove-pipe’ effect, in which each

functional group employs policies that are intendedly rational at the group level, but that also lead to inefficiency across operations at the hospital level. Our data suggests that critical determinants of success in efficient hospital operations include the perceptions stakeholders have about the effects of the actions on upstream and downstream processes. Faulty attributions about the drivers of efficiency can trap operations in deteriorating modes of performance, and subvert the momentum gained from IT-enabled processes. We recommend management improvements in both materials and staff utilization to address the stove-pipe effect, and estimate the resultant cost-saving. As part of this analysis we also demonstrate opportunities to merge real-time operational data with feedback modeling to provide dynamic tools for hospital administration, risk management, and education and training. We believe that the major gains in health information systems use will accompany new information gathering capabilities, as these capabilities result in collections of data that can be used to greatly improve patient safety, hospital operations, and medical decision support.

In the future, we envision a system that merges newly available operational data sources (i.e., real-time POAS data), electronic medical records, and operational data into feedback models that create dynamic ward management tools. We propose a system that would provide information to hospital about efficiency metrics, and would provide managers a means to change policies, such as which drugs to exempt from mixing, at any point during on-going hospital operations. This platform would be open for improvements, promoting the development of additional tools that improve operations and manage patient risk.

The data produced by the POAS system combined with computational social science methodologies, particularly system dynamics modeling, is an exciting step forward in understanding the dynamics of managerial improvement in hospitals. With the reported initial results and continued research, we predict that this work will have a significant impact on reducing hospital costs, improving patient safety, and accommodating improvements in hospital staff operations.

References

Akiyama M, Kondo T.. Risk management and measuring productivity with POAS--point of act system. Medinfo. 2007;12(Pt 1):208-12.

Akiyama M. Risk Management and Measuring Productivity with POAS - Point of Act System. A Medical Information System as ERP (Enter-prise Resource Planning) for Hospital Manage-ment. Methods Inf Med. 2007;46(6): 686-93.

Akiyama M., Migration of the Japanese health-care enterprise from a financial to integrated management: strategy and architecture, Medinfo. 10(Pt 1):715-718,2001.

Linda T. Kohn, Janet M. Corrigan, and Molla S. Donaldson, Editors. To Err Is Human: Building a Safer Health System Committee on Quality of Health Care in America, Institute of Medicine. 2000.

Repenning NP, Sterman JD. Capability Traps and Self-Confirming Attribution Errors in the Dynamics of Process Improvement. Administrative Science Quarterly 2002; 47(2):265-295.

Appendix 1. Simulation Model

