MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–20
ISSN 1523-4614 (print) | ISSN 1526-5498 (online)

http://dx.doi.org/10.1287/msom.1120.0381

Information *Hang-overs* in Healthcare Service Systems

The literature on business process design has focused on issues such as bottlenecks, workflow configuration (series vs. parallel), replacing an existing workflow with a shorter one, etc. One important issue that has not received adequate attention is the information-intensive nature of medical service systems. Performance of clinical workflows depends not only on how various steps are carried out but also on when certain information items are collected. We report the results of a long-term empirical study that looked at the implementation of a Radiology Information System (RIS) at a large regional network of radiology clinics. We find that a failure to gather necessary clinical background information in earlier steps significantly delays later steps and causes them to *hang over,* with a significant impact on the total turnaround time of diagnostic reports. We show that information systems can solve this problem by separating the task of gathering information from its usage and relocating that task upstream in the workflow. We argue that such unbundling can lead to shorter report turnaround times—even if it significantly increases the utilization of the bottleneck server. These results have broader implications for the optimal design of other clinical workflows, such as the process of filling prescriptions in pharmacies or the typical surgical pre-anesthesia evaluation in hospitals. Finally, we explain why the impact of addressing hang-over is often non-uniform across clinical modalities, providers, and patient types.

[Keywords: Medical information systems, clinical workflow, workflow design, hang-over, feedback queue.]

1. Introduction

The healthcare sector is plagued with information deficiencies. Health services are information-intensive, but several reports point out instances of missing, inaccurate, incomplete, illegible, and inconsistent medical information. A New England Journal of Medicine study on primary care physicians (Smith 2005) points out, *"Clinicians reported missing clinical information in 13.6% of visits; missing information included laboratory results (6.1% of all visits), letters/dictation (5.4%), radiology results (3.8%), history and physical examination (3.7%), and medications (3.2%). Missing clinical information was … to be at least somewhat likely to adversely affect patients (44%), and to potentially result in delayed care or additional services (59.5%)."* In light of such reports, it has become more important than ever to critically examine the role of information in the delivery of clinical services.

To address this need in part, we investigate how timely gathering of medical information facilitates timely delivery of time-critical services. Process-changing technologies are spreading rapidly within the healthcare field, and this paper provides insights into when benefits (as measured by throughput time and customer satisfaction) are most likely to accrue from investments in them.

We conduct our investigation by analyzing the radiology workflow of an outpatient clinic; this workflow is shown in Figure 1. In the first step, scheduling, a referring physician's office or a patient calls the telephone scheduling center to set up an appointment. Then, in the pre-certification step, the insurance information and coverage are verified. Upon arrival for the exam, the patient registers at the front desk and provides additional personal information. Next, a technologist performs the medical scan. In the critical interpretation step, a radiologist examines the images and background clinical information, diagnoses, and dictates the results. A transcriptionist then prepares a report from the recorded dictation. The report and the images are delivered to the referring physician after the radiologist reviews them and signs off.

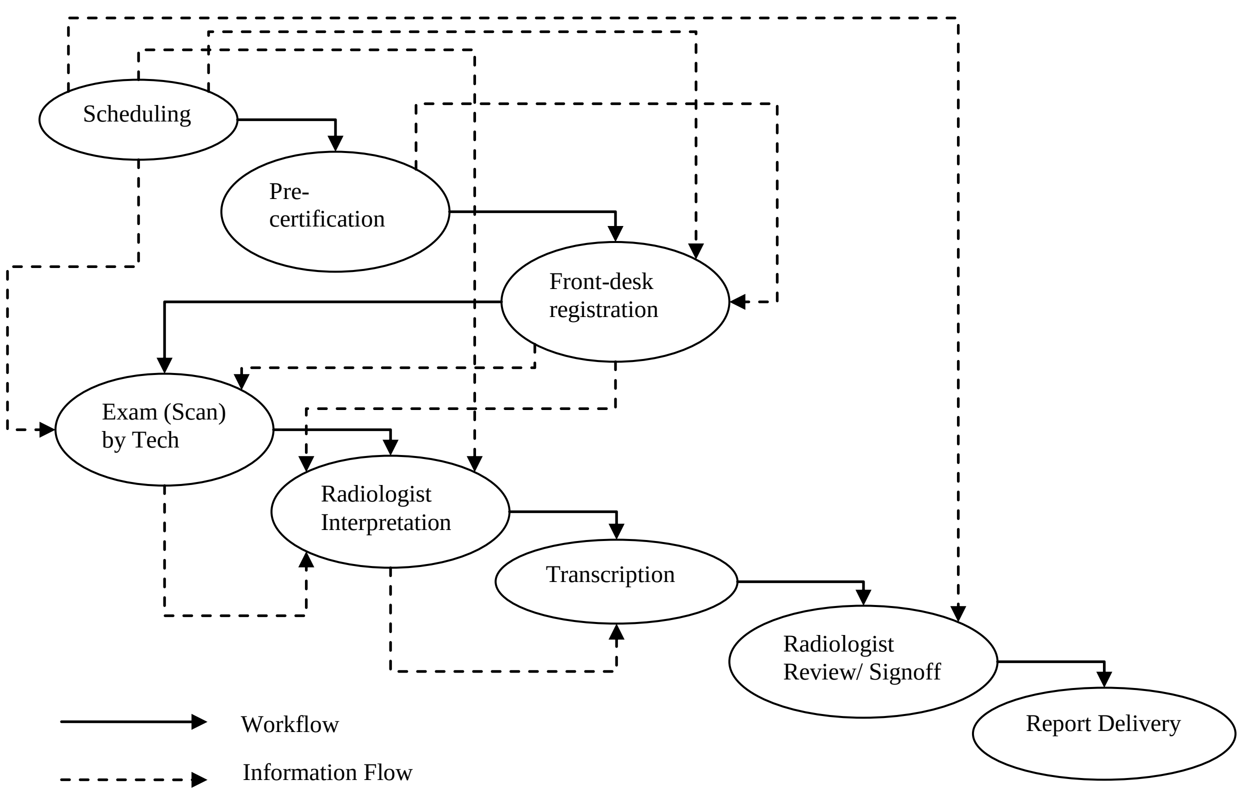

Figure 1. Radiology Workflow; Adapted from (Lahiri and Seidmann 2009)

While the work moves along the solid arrows shown in Figure 1, critical information moves along the dashed arrows. In the front desk registration step, the co-pay amount is calculated and collected based on the insurance information gathered in the scheduling step. The technologist also uses a number of data items gathered in earlier steps, such as information on patient-handling (e.g., the patient may not eat during a certain period before the exam), or on patient-safety (e.g., the patient is claustrophobic, or he is allergic to iodinated contrast media). The radiologist, too, uses information gathered at the time of scheduling (e.g., the requisition from the referring doctor, notes from the referring doctor on patient conditions such as swelling of a tissue, prior studies, and other relevant medical history). In addition, the radiologist depends on information gathered at the time of the exam (e.g., reasons for capturing additional views and the repeat-reject ratio). The transcriptionist uses the reporting template specified by the radiologist while preparing the report. In short, each step processes a lot of information, and it depends critically on earlier steps for its information needs.

The most information-intensive step is the interpretation step. For example, this step for a mammography exam has significant needs. The radiologist must know the patient's demographic information, the history of carcinoma in her family, images from her current and past mammograms, reasons for different BIRADS levels[1] reported in her past mammograms, her past breast surgeries, implants, biopsies, etc. in order to interpret the current images accurately. If all clinically relevant information is not available, the radiologist must put the interpretation and dictation on hold. In other words, the interpretation will *hang over* until the missing information is supplied.

Hang-over issues are particularly common in healthcare. When CVS did not collect information related to drug allergies or insurance from patients at the time of prescription drop-off (McAfee 2005), it caused the step of prescription-filling to hang-over. This hang-over problem led to such long delays that many CVS customers defected to other pharmacies. CVS finally addressed the problem by "manning" the drop-off counters.

It is important to note that hang-overs may also have serious health consequences. A Yale University School of Medicine study (Holt et al. 2007), which looked at nearly 1,800 surgeries, finds that nearly 23% of surgeries were delayed because of missing information, putting patients at severe risk. According to the authors, "*Day-of-surgery delays caused by missing information remain relatively common despite pre-anesthesia evaluation.*" They also point out, "*Pre-anesthesia evaluation clinics are a common component of hospital institutions, but information obtained during the presurgical visit may not be available on the day of surgery. As a result, missing patient information may continue to cause delays on the day of surgery.*" Evaluating health consequences is beyond the scope of this research. However, this research still makes two important contributions: (a) it provides a framework for measuring the hang-over problem and analyzing its impacts, and (b) it discusses how clinical workflows can be redesigned using enterprise information systems in order to mitigate the hang-over problem.

[1] The BIRADS level of a mammogram is a number between 1 and 5; 1 indicates negative and 5 indicates seriously malignant. If any prior study reported a BIRADS level above 1, the radiologist would want to know why it did so.

There are reasons hang-over is so common in healthcare. For other services, such as finance, retail, travel, and entertainment, the question is often about the amount and quality of information. For example, Varian (2004) points out that better information on consumer demand leads to efficient inventory management, effective pricing, and successful development of new services. Clinical information, on the other hand, is not optional—it is essential. Until all the required medical information is gathered, clinical tasks cannot be carried out. The debate is, therefore, about when the relevant information should be gathered, and not about whether it should be gathered at all. As we discover in this study, the task of gathering clinical information is often postponed to downstream steps. This postponement causes the hang-over problem, delaying time-critical clinical steps.

Outpatient clinics, such as the one that we study, depend heavily on referrals and, as a result, constantly endeavor to satisfy consumers—both patients and referring doctors—by reporting results quickly. We, therefore, use the turnaround time, i.e., the total time (waiting plus service) taken for interpretation and preparation of the report, as our primary metric. We study the turnaround time and its components for two modalities, digital mammography and MRI. We also study other ancillary metrics, such as the average call length at the call center responsible for scheduling. In addition, we analyze patient and referring physician satisfaction surveys.

A major empirical finding is that using a newly installed Radiology Information System (RIS) to gather information in a disciplined fashion upstream in the workflow produced substantial operational benefits for mammography. On average, the turnaround time declined by 50%, primarily due to faster interpretation. Interestingly, we did not find any significant impact on turnaround times of MRI exams.[2] We argue that the difference between these two modalities is attributable to the difference between the respective improvements in their hang-over rates. We explain the implications of this differential impact. Further, we argue that a lower average turnaround time is possible even when the utilization of the bottleneck server goes up; this finding has implications for healthcare providers—both for those already using a clinical information system and those planning to acquire one—who often only see the additional burden with technology solutions without understanding their full system benefits.

The rest of the paper is organized as follows. In §2, we present a literature review. We describe our data collection efforts in §3, following which we present our empirical findings in §4. We discuss these findings in §5 and then supplement the discussion with an analytical model in §6. We conclude by summarizing our contributions and outlining directions for future research in §7.

[2] Proper and comprehensive patient information may improve resource management and quality even if it does not really speed up reporting of results. Durations of some patient-related tasks may be dependent on patient information (e.g., height, BMI) that may assist in better planning of slots, patient sequencing, etc. Adverse medical effects may also go down as a result of better information. We do not investigate these additional benefits.

2. Literature

This work is one of the first interdisciplinary studies to investigate the performance impact of workflow changes initiated by an enterprise information system on medical service organizations. The need to make medical service organizations efficient has recently become a top priority in many nations. In the United States, President Obama's administration has committed $20 billion to computerize medical records (Koman 2009). Yet, there exists limited research with regard to what works and what does not work in this domain. A case in point is Kaiser Permanente, which spent several hundred million dollars on an electronic medical record system only to scrap that system later and start all over again (King 2009). Citing Kaiser Permanente's difficulties, Agarwal et al. (2010), in their recent survey of the literature related to healthcare information technology, have called for further research into issues related to implementation. According to them, "*Regardless of functionality, these systems will have little impact on performance if they are not well integrated into the daily workflows of care providers, as illustrated in the implementation challenges faced even by large and highly successful health-care organizations like Kaiser Permanente.*" We partly address this gap in the literature by empirically and analytically linking the benefits of technology with its ability to improve information processing in workflows.

The literature on radiology workflows is very closely related. Research in this area primarily looks at filmless radiology and the impact of a Picture Archiving and Communications System or PACS (Siegel et al. 2004). An RIS is often used in conjunction with a PACS and an Electronic Healthcare System (EHS). PACS consists of digital image acquisition systems, digital storage for images, and display workstations that can be used to enhance and compare the digitally stored images. EHS houses patient information and provides billing and collection functionalities. RIS, on the other hand, automates a patient's journey through the health-imaging workflow by using what is commonly termed a global work list. The work list displays and tracks all active exams and their status in real time. RIS also allows filtering of the work list to fit operational needs and to prevent unauthorized access. For example, an RIS can filter the work list presented to a technologist (also called a tech) to include only exams scheduled in the room to which he is assigned. Commonly observed benefits of PACS are elimination of film-processing (Duerinckx and Grant 1998, Reiner et al. 2002), digital analysis of images, digital delivery of results, digital storage of clinical history, cost savings (Ayal and Seidmann 2009a, Ayal and Seidmann 2009b, Chesson 2006), higher productivity, and greater user satisfaction (Srinivasan et al. 2006). Our contribution to this stream is that we focus on RIS and its informational role.

Another area relevant to this research is work system design (Denning and Medina-Mora 1995, Buzacott 1996, Seidmann and Sundararajan 1997, Powell 2000, Pinker and Shumsky 2003, Hasija et al. 2005, Green et al. 2006, Karmarkar and Apte 2007, Dobson et al. 2009, Lu et al. 2009, etc.). Buzacott (1996) examines different workflow configurations and compares their performances. Powell (2000) considers the optimal assignment of tasks to workers, with the twin objectives of minimizing the processing time and staffing cost. Dobson et al. (2009) find that a large non-clinical staff can increase hand-offs between clinical and non-clinical staff and adversely affect the productivity of medical practices. Green et al. (2006) propose optimal policies for managing demand at hospital diagnostic facilities with the objective of achieving faster service. These studies, however, do not specifically talk about issues related to processing of information. But there are others who do. For example, Pinker and Shumsky (2003) show that incentivizing workflow participants is critical to addressing information asymmetry in tiered workflows in which work moves from one tier of participants to another, for example, from a primary care physician to a specialist. In a more recent study, Dobson and Sainathan (2011) model the workflow of an emergency room; they focus on the role of information on patients' waiting cost, which is needed for triaging. Despite the call for an increased focus on information-intensive services in recent years (Karmarkar and Apte 2007), there exists little research that discusses the hang-over problem, its impact on clinical workflows, and ways to address it. The existing research primarily focuses on configurations (e.g., series vs. parallel), or on economic aspects (e.g., designing proper incentives or contracts to address information asymmetry). It does not explicitly address the workflow performance issues that we examine here, particularly the role of shifting information-gathering tasks upstream in addressing the hang-over effect. Besides, our work approach is mostly empirical, while the existing research is mostly analytical.

Finally, our research extends the extant literature on the business value of information technology (IT). Brynjolfsson and Yang (1996), Dedrick et al. (2000), and Banker and Kauffman (2004) provide useful reviews of the existing literature that quantifies this value. Broadly speaking, the business impact of IT has been studied at three levels: firm, industry, and economy. Our work belongs to the firm-level literature. A number of studies in the firm-level stream measure the impact of IT on productivity, following the emergence of the so-called productivity paradox (Brynjolfsson 1993, Dewan and Kraemer 1998). Most of these studies, if not all, refute the paradox as originally stated and show that investments in IT lead to substantial returns. One sub-stream evaluates IT as a factor of production and assesses its economic impact through aggregate measures like the return on IT capital (Brynjolfsson and Hitt 1996, Dewan and Min 1997, Lehr and Lichtenberg 1999, Menon et al. 2000, Hitt et al. 2002). A second sub-stream evaluates direct impacts of IT on measures like lead time, inventory turnover, processing speed, etc. (Mukhopadhyay et al. 1995, McAfee 2002, Cotteleer and Bendoly 2006). Our work belongs to this second sub-stream. We are, however, among the first to present empirical evidence that positively links the performance impact of IT on a workflow with its ability to unbundle gathering of information from its usage.

3. Data

One reason medical information is often not gathered up front, despite a severe downstream need, is that it is highly case sensitive. For example, in radiology, information requirements differ from one modality to another (e.g., MRI and mammography have different clinical background information needs) and from one patient to another (depending on age, medical history, allergies, etc.). It is, therefore, challenging for non-clinical staff members, such as the scheduling staff, to gather all necessary information in earlier steps. A telephone-scheduler may not know what information is relevant for a mammography patient who had breast implants that were later removed. The same staff member may also not know when to ask for information on a patient's sensitivity to iodinated contrast media, compromised kidney functions, or to inquire about other risk factors. During later steps, clinical staff members—trained nurses, technologists, or radiologists—may identify these missing data items and request their collection.

Sophisticated data-gathering features offered by clinical information systems, such as an RIS, partly address this data-gathering challenge. For example, a modern RIS can be configured to automatically ask for specific clinical information on iodine allergies, kidney function, contrast media sensitivity, claustrophobia, etc. at the time of scheduling of an MRI exam using paramagnetic contrast media. An RIS thus makes it possible to *unbundle* the task of gathering clinical information from its use. Equipped with an RIS, a radiology practice can gather more information in earlier steps.

The ability of an RIS to improve information gathering is not limited to scheduling alone; for example, during the exam step, a properly configured RIS prompts automatically for reasons for any additional views captured. Interestingly, RIS is not alone in offering such information-gathering capabilities. Other medical information systems, such as an Anesthesia Information Management System (AIMS), provide the same functionality. A recent bulletin published by the Institute for Safety in Office-based Surgery claims (Ruskin 2011): "*This ability to collect, store, and organize large amounts of data makes AIMS systems ideal for quality assurance and clinical research.*"

Borg Imaging LLC, the focal firm in our study, is a large network of free-standing health-imaging centers. At the time of the study, it billed 125,000 imaging procedures a year from seven locations in upstate New York. Both before and after RIS, it had the same 13 radiologists and 140 staff members. Previously, it had switched to filmless radiology by adopting PACS. Our study thus does not capture the effects of PACS implementation or of any other organizational changes. Rather, it examines the process changes that were initiated at the time of the RIS installation.

The company acquired the Kodak Carestream™ RIS[3] to replace its homegrown system. The new RIS went live in October 2006. Many process changes were also initiated to take advantage of the new sys-

[3] Carestream™ is now owned by Onex.

tem's information-gathering capabilities. A summary of these process changes and their expected impacts can be found in Table 1.

Telephone schedulers were given case-based scripts to collect a number of data items in a systematic fashion. These data items include: insurance information to facilitate proper co-pay collection, order notes from the referring physician describing the patient's problems, information on other consulting physicians, prior outside studies, special handling instructions (information on implants, swelling, allergies, etc.), special preparation instructions (for example, the patient may not eat before the exam), patient safety instructions (typically related to pregnancy, claustrophobia, and other medical conditions), and family and demographic details. The technologists were also given instructions to record additional information at the time of the exam: they were asked to document reasons for taking additional views, the repeat-reject ratio, etc. In addition, the new RIS was integrated fully with the PACS so that prior inside studies could be accessed easily. The overall objective was to increase the availability of information throughout the workflow, particularly at the bottleneck clinical interpretation step.

Table 1. Process Changes Initiated Using RIS and Their Expected Impacts

Relocation of Information Gathering	**Expected Negative Impact**	**Expected Positive Impact**
Disciplined gathering of background clinical information during **telephone scheduling** rather than deferring gathering of study-sensitive information to the interpretation step	Adds to the workload of the schedulers; potentially adds to the radiologists' workload by requiring them to navigate through a number of screens and tabs to view the collected information	Faster turnaround by radiologists (as delays resulting from missing information are reduced), leading to faster service in the interpretation step as well as greater referring physician satisfaction
Disciplined gathering of exam-specific clinical information at the time of the **exam** rather than deferring in part to the interpretation step	Adds to the workload of the technologists; potentially adds to the radiologists' workload by requiring them to navigate through a number of screens and tabs to view the collected information	Faster turnaround by radiologists (as delays resulting from missing information are reduced), leading to faster service in the interpretation step as well as greater referring physician satisfaction
Disciplined gathering of reporting template information in the **interpretation** step rather than deferring to the transcription step	Adds to the radiologists' workload and reduces flexibility in reporting	Faster transcription turnaround (as the need to search for a suitable template is reduced) as well as greater referring physician satisfaction
Disciplined gathering of insurance information at the time of **telephone scheduling** rather than deferring the job to the registration step	Adds to the workload of the schedulers, and also keeps the staff of the referring physician longer on phone	Efficiency and accuracy in billing and collection, leading to greater referring physician and patient satisfaction

While the changes described above were implemented, all other possible experiments with the workflow were postponed in order to facilitate a clear comparison of the data collected prior to the RIS installation with the data collected afterward. In other words, we sought to establish a "controlled" setting in order to prevent any unwanted "noise" from interfering with the comparison of pre-RIS and post-RIS performance measures. One downside of our controlled setting is that it does not identify "delayed" impacts, which requires analysis of longitudinal data (Devaraj and Kohli 2000).

We studied four locations in Borg's network. The workload of the imaging centers did not show any significant seasonality effect. The pre-RIS sample was collected in June 2006, before technologists, radiologists, transcriptionists, and other users started their preparations for the new system. Following the "go-live" in the last week of October 2006, the data collection resumed. Since we wanted to prevent "learning effects" (McAfee 2002, Ayal and Seidmann 2009a) from interfering with our analysis, we did not include data collected through February 2007 in our final analysis; rather, we used that data to ensure that the learning phase was over before we collected our final sample in March 2007. Most of the data items were collected through automated means, while the rest (exception documents and satisfaction surveys) came from paper records—a summary of these measures is shown in Table 2. The pre- and post-RIS sample sizes differ in some cases; in each case, however, they are adequately large for statistical analysis.

Table 2. Summary of Data Collection; an asterisk (*) next to the post-RIS mean indicates that it is significantly different from its pre-RIS counterpart at a 5% level

Measure (Pre-RIS, Post-RIS Sample Sizes)	*Pre-RIS*				*Post-RIS*			
	Mean		**SD**		**Mean**		**SD**	
Scheduling-Call Length (234 calls, 338 calls)								
Existing patients, in <u>Minutes</u>	2.49		1.32		2.73*		1.23	
New patients, in <u>Minutes</u>	3.10		1.18		4.44*		1.59	
Scheduling-call abandonment rate or % of calls abandoned daily (69 work days, 25 work days)	2.22%				3.59%*			
Patient Satisfaction Surveys (137 responses, 76 responses)				See Table 5				
Referring Physician Satisfaction Surveys (83 responses, 102 responses)				See Table 6				
	Mammography		*MRI*		*Mammography*		*MRI*	
	Mean	**SD**	**Mean**	**SD**	**Mean**	**SD**	**Mean**	**SD**
Technologist's Administrative Time in <u>Minutes</u> (112 exams, 124 exams)	6.03	3.85	4.39	2.21	2.77*	1.48	3.94	1.78
Fraction of interpretations delayed as reported by exception documents (208 exams, 238 exams)	7.92%		3.74%		2.56%*		2.48%	
Report turnaround time or RTAT, sum of the 3 TATs shown below, in <u>Hours</u> (217 exams, 1325 exams)	4.06	2.34	3.11	1.87	2.17*	1.43	3.20	1.85
Interpretation TAT in <u>Hours</u>	1.58	1.84	1.62	1.06	0.76*	0.60	1.67	0.98
Transcription TAT in <u>Hours</u>	0.44	0.42	0.54	0.36	0.36	0.37	0.50	0.57
Review TAT in <u>Hours</u>	2.04	1.05	0.95	1.41	1.06*	1.24	1.03	1.35

4. Findings

In this section, we present a statistical comparison of the pre- and post-RIS values of different metrics, such as the components of the report turnaround time, length of scheduling calls, technologist's administrative time per exam, and patient and referring physician satisfaction. The objective here is to find out whether the expected impacts mentioned in Table 1 were realized for both mammography[4] and MRI. We begin by discussing the impacts of disciplined information gathering on the report turnaround time (RTAT) and its components. Then, we discuss the impacts on the scheduling and exam steps. Finally, we present a summary of patient and customer satisfaction surveys.

4.1. Turnaround Times

We measured three components of the RTAT (see Figure 1): (1) the interpretation TAT (the time elapsed from the end of exam to the end of dictation), (2) the transcription TAT (the time elapsed from the end of dictation to the end of transcription), and (3) the review TAT (the time elapsed from the end of transcription to the end of review). These components of the RTAT include both the queuing and service times.

We discovered that the distribution of the RTAT and its components did not follow the *normal* distribution. Their probability density functions were asymmetric and one-tailed. Hence, we used the *lognormal survival function* to compare the pre- and post-RIS averages of the RTAT and its components.[5] Put in simpler terms, we regressed the logarithm of the RTAT and its components on a dummy variable, *RIS*, which was 1 for post-RIS and 0 for pre-RIS, as well as a control variable, *Mammo*, which was 1 for mammography and 0 for MRI. The regression equations, with ε denoting the error term of each, are as follows:

$$\textit{Logarithm of Interpretation TAT} = \beta_0 + \beta_1\, \textit{Mammo} + \beta_2\, \textit{RIS} + \beta_3\, \textit{Mammo} \times \textit{RIS} + \varepsilon \quad (1)$$

$$\textit{Logarithm of Transcription TAT} = \beta_0 + \beta_1\, \textit{Mammo} + \beta_2\, \textit{RIS} + \beta_3\, \textit{Mammo} \times \textit{RIS} + \varepsilon \quad (2)$$

$$\textit{Logarithm of Review TAT} = \beta_0 + \beta_1\, \textit{Mammo} + \beta_2\, \textit{RIS} + \beta_3\, \textit{Mammo} \times \textit{RIS} + \varepsilon \quad (3)$$

The first and third equations were used to measure the impact on the radiologist's performance, and the second one to measure the impact on the transcriptionist's performance. We first estimate (1) – (3) using ordinary least squares (OLS). This OLS model is called **Model A** henceforth. We also estimate another model, which is described below.

Model B: This model uses additional control variables for locations. We collected data from four locations: White Spruce, South Clinton, Lattimore 1.5, and Lattimore Open. In this model, there is a dum-

[4] We studied 'diagnostic' mammography, i.e., patients with breast disease symptoms (ACR 2008), and ignored screenings that were interpreted offsite in batches. For diagnostic mammography, timely reporting is highly critical.
[5] Other survival functions, such as the *gamma* distribution, lead to nearly identical coefficient estimates.

my variable for each of the last three locations. There is no dummy variable for the control location, White Spruce; i.e., all three location dummies were set to zero for White Spruce.

Table 3 shows estimates for (1) – (3) using Models A and B, respectively. The table also shows the p-value or significance level for each estimate immediately below it. The very last row of Table 3 shows the p-value for the test H_0: $\beta_2 + \beta_3 = 0$. A large p-value for this test would indicate that the RIS and accompanying workflow changes had no impact on that component of turnaround time for mammography. No such test is needed for MRI; according to our model setup, a large p-value for β_2 would mean that the RIS had no impact on this modality.

Both models in Table 3 show very similar estimates for our critical model parameters, β_2 and β_3. According to either of these models, the RIS installation, surprisingly, had little impact on MRI (the p-value for β_2 is quite large for each equation). Also, as shown in Table 2, the average RTAT increased slightly from 3.11 to 3.20 hours/study, which is not significant at a 5% level.

Table 3. Estimates of (1), (2), and (3); Model A: without control variables for locations; Model B: with dummy control variables for locations

	Equation 1: Interpretation TAT		Equation 2: Transcription TAT		Equation 3: Review TAT	
Model	A	B	A	B	A	B
Intercept (β_0)	0.2874 (<0.0001)	0.3489 (<0.0001)	-0.8357 (<0.0001)	-0.8194 (<0.0001)	-0.8299 (<0.0001)	-1.1027 (<0.0001)
Mammo (β_1)	-0.3255 (0.0357)	-0.3377 (0.0400)	-0.3746 (0.0778)	-0.3516 (0.1203)	1.3528 (0.0001)	1.5963 (<0.0001)
RIS (β_2)	0.0771 (0.3056)	0.0680 (0.3689)	-0.2525 (0.1147)	-0.2579 (0.1136)	0.0294 (0.8624)	0.0047 (0.9781)
Mammo x RIS (β_3)	-0.5215 (0.0015)	-0.5093 (0.0019)	0.0614 (0.7850)	0.0688 (0.7592)	-1.2254 (0.0009)	-1.2025 (0.0011)
South Clinton		-0.0851 (0.2636)		-0.0680 (0.5162)		0.0505 (0.7673)
Lattimore1.5		-0.1415 (0.0409)		-0.0428 (0.6532)		0.5130 (0.0010)
Lattimore Open		-0.0014 (0.9864)		0.0205 (0.8585)		0.4631 (0.0134)
H_0: $\beta_2+\beta_3= 0$	(0.0003)	(0.0022)	(0.3422)	(0.3418)	(0.0002)	(0.0002)
R-squared	30.32%	30.98%	2.15%	2.26%	2.61%	4.47%

On the other hand, as confirmed by the estimates of (1) and (3) in Table 3, the average interpretation and review turnaround times for mammography declined significantly (the p-value for the test H_0: $\beta_2 + \beta_3 = 0$ is extremely small for both (1) and (3), and the point estimate of $\beta_2 + \beta_3$ is negative in each case), indicating an acceleration in the radiologist's performance. Also, as outlined in Table 2, the sum of the average interpretation TAT and the average review TAT declined by 49.72%, from 3.62 (=1.58+2.04) to 1.82 (=0.76+1.06) hours/study, leading to a statistically significant decline in the average RTAT from 4.06 to 2.17 hours/study.

As already mentioned, one of the critical aspects of the RIS implementation was the unbundling of the collection of background clinical information from the image interpretation task as a way of speeding up the radiologist's station. It appears from Table 3 that the benefits of moving the data-gathering tasks produced the expected impact for mammography only. For MRI, we did not notice any similar acceleration in the interpretation and review steps. The same technology and process changes thus worked for one modality but not for the other. This difference across modalities has cardinal implications for effective design of healthcare workflows using enterprise information systems. We explain this difference and its implications in §5.1.

4.2. Scheduling

In order to maximize the throughput rates of the clinical staff, call center operators were given case-based scripts that guided them in the collection of pertinent clinical background data from all new and existing patients. We used the switchboard data collection software used by Borg's call centers to measure the impact on scheduling-call lengths and call-abandonment rates.

We used a similar *lognormal* model to estimate the impact on the scheduling-call length, for very similar reasons. As before, the dummy variable *RIS* was set to 1 for post-RIS and 0 for pre-RIS. The dummy variable *New_Patient* was set to 1 for a call if the call was for a patient new to Borg. The regression equation is:

$$\textit{Logarithm of call length} = \beta_0 + \beta_1\, \textit{RIS} + \beta_2\, \textit{New\_Patient} + \beta_3\, \textit{New\_Patient} \times \textit{RIS} + \varepsilon \qquad (4)$$

Since the switchboard data collection software did not specify which calls were for MRI and which ones were for mammography, we were unable to break down the impact by modality. It is important to note that this dataset does not include any other modalities; we collected it from the call center that handled calls for MRI and mammography only.

According to Table 4, the average call length expectedly increased post-RIS for both new and existing patients, and the increases are significant at a 5% level. According to Table 2, the average call length for existing patients went up by roughly 9%, from 2.49 minutes to 2.73 minutes. We noticed a larger increase (43%) for new patients: the average climbed from 3.1 minutes to 4.44 minutes.

Table 4. Estimate of (4)

Estimate of	
Intercept (β_0)	0.8127 (<0.0001)
RIS (β_1)	0.1018

	(0.0067)
New_Patient (β_2)	0.2365
	(0.0010)
New_Patient × RIS (β_3)	0.2788
	(0.0052)
R-square	11.81%

Further, as shown in Table 2, the call-abandonment rate (the percentage of customers hanging up because of a long wait, measured daily) increased from 2.22% pre-RIS to 3.59% post-RIS. This increase is significant at a 5% level, and it also reflects a significant adverse impact on schedulers. However, as we discuss later, an adverse impact on scheduling does not necessarily imply an adverse impact on the system.

4.3. Technologist's Administrative Time

The technologist's administrative time is the time that the technologist, who performs the medical scan, spends on reviewing and preparing necessary documentation. The RIS installation called for an additional burden on technologists by requiring them to fully document all process exceptions, such as additional views, retakes, non-standard exposures or positions, or other unusual indications. Despite this additional reporting burden, as noted in Table 2, their administrative time per study went down post-RIS. For mammography, it declined from 6.03 to 2.77 minutes (significant according to the 2-population t-test, p-value < 0.0001), and, for MRI, from 4.39 to 3.94 minutes (not significant, p-value = 0.3834).

4.4. Patient and Referring Physician Satisfaction

The findings presented so far mainly involve tangible operational metrics. However, since outpatient clinics like Borg depend heavily on referrals and word-of-mouth communication, it was imperative that we examined intangible aspects such as the customer satisfaction. In the case of outpatient radiology clinics, the word "customer" comprises both referring physicians and patients. We now present the impacts on these two customer groups separately.

Borg periodically handed out paper surveys to referring physicians and patients. The post-RIS surveys included in our analysis were handed out almost a year after the go-live. This one-year lag was introduced to ensure that post-RIS surveys captured the reality of the post-RIS situation and not lingering impressions of the pre-RIS environment. Both surveys had the same format. Survey-takers were asked to rate Borg on different dimensions on a 7-point scale, with 1 being the lowest score and 7 being the highest. Borg typically received high ratings in these surveys, and they considered a rating other than 6 or 7 to be unacceptable. We, therefore, begin by looking at the fraction of respondents who rated Borg 6 or 7, before and after RIS.

As per Table 5, the fraction of referring physicians who rated Borg 6 or 7 increased significantly in some dimensions but not all. The perception of the scheduling process improved despite the increase in the average call length. The biggest improvement was in dimension 6 of Table 5: the fraction of physicians who were satisfied with timely delivery of reports increased from 78% to 91%. Some other dimensions also improved significantly, and so did the overall satisfaction. Since physicians' responses to different dimensions were highly correlated, we felt that it was necessary to use Principal Component Analysis (PCA). When we searched for orthogonal principal components, we found only one component that explained more than 10% of variations; this component, which we named Factor 1, explained nearly 54% of all variations alone. The correlation of Factor 1 with the dummy variable *RIS* turned out to be both positive (= 0.125) and significant (the p-value for the one-tail test is 0.0445). This positive correlation also suggests that streamlining information collection translated not only to tangible operational benefits but also to improvements in intangible aspects such as referring physicians' satisfaction.

Table 5. Referring Physician Satisfaction: Fraction of respondents who rated Borg 6 or 7 on a 7-point scale; an asterisk (*) indicates that the post-RIS fraction is statistically different from the pre-RIS level at a 5% level. Next to each fraction, the corresponding average rating is noted within parentheses.

Overall Satisfaction	**Pre-RIS**	**Post-RIS**
Overall satisfaction of the referring physician	0.86 (6.37)	0.95* (6.48)
Rating of patient satisfaction by the referring physician's office	0.86 (6.44)	0.92* (6.53)
Dimensions of service		
1. Professionalism of scheduling staff	0.88 (6.44)	0.96* (6.54)
2. Expertise of scheduling staff	0.85 (6.39)	0.94* (6.53)
3. Flexibility of scheduling process	0.81 (6.22)	0.88* (6.47)
4. Efficiency of the scheduling process	0.88 (6.41)	0.93 (6.52)
5. Simplicity of the scheduling process	0.89 (6.41)	0.94 (6.56)
6. Timely delivery of clinical reports	0.78 (6.25)	0.91* (6.40)
7. Image delivery options offered at the time of scheduling	0.86 (6.40)	0.89 (6.52)
8. Responsive addressing of errors and corrections	0.87 (6.54)	0.94* (6.72)
9. Report delivery in requested format	0.86 (6.45)	0.90 (6.58)

The fractions of patients who rated Borg 6 or 7 are shown in Table 6. We noticed that the impact of RIS on patient satisfaction was minimal across the board. We confirmed this largely insignificant impact using Principal Component Analysis (the details of which we are omitting for the sake of brevity). Significant changes were observed only in the dimensions related to registration and co-pay collection at the front desk (i.e., dimensions 3 and 7 in Table 6).

Table 6. Patient Satisfaction: Fraction of respondents who rated Borg 6 or 7 in a 7-point scale; an asterisk (*) indicates that the post-RIS fraction is statistically different from the pre-RIS level at a 5% level. Next to each fraction, the corresponding average rating is noted within parentheses.

Overall Satisfaction	**Pre-RIS**	**Post-RIS**
Overall satisfaction of the patient	0.94 (6.60)	0.95 (6.63)
Likely to recommend Borg to others	0.89 (6.53)	0.89 (6.54)
Dimensions of service		
1. Accuracy of the scheduling process	0.90 (6.54)	0.92 (6.62)
2. Respectful attention at the front desk	0.97 (6.85)	0.97 (6.86)
3. Efficiency of check-in and registration process	0.94 (6.70)	0.99* (6.88)
4. Acceptability of wait time (between registration and exam)	0.92 (6.53)	0.92 (6.68)

5. Procedure explanation provided by Borg's staff	0.97 (6.87)	0.97 (6.84)
6. Perception of the journey through various phases of the appointment	0.96 (6.79)	0.96 (6.80)
7. Cost information provided by Borg's staff	0.83 (6.38)	0.93* (6.49)

5. Discussion

We start the discussion of our empirical findings by looking at the impact of RIS-initiated process changes on the radiologist's turnaround times. We then discuss the impacts on transcriptionists, schedulers, technologists, and satisfaction surveys.

5.1. Radiologist Interpretation and Review

To understand the impact on turnaround times, we must first examine the way a radiologist does her work. What would a radiologist do if the information required at the time of interpretation has not been gathered in earlier steps? The radiologist would not risk the safety of the patient by trying to guess. She would suspend the interpretation of that exam and "park" it until all missing information is collected by the assisting clinical staff. The radiologist would also document the missing information and defer the dictation until the missing information is delivered to her.

The implication of such "parking" is that, while an exam for which the complete set of information is available is interpreted in a timely fashion, the interpretation of those without necessary information can suffer a large delay, ranging from hours to days. This delay includes the time that the assisting administrative staff needs to follow up with the patient, the referring physician, or the technologist for the missing information. It can make the exam's report turnaround time (RTAT) very high and can also drive up the average RTAT for the clinic. We, therefore, conjecture that a reduction in the percentage of studies with incomplete information entering the radiologist's queue is critical to achieving a lower average RTAT.

Our research into the process-exception documentation at Borg (see Table 3 for the sample size information) identified many interesting instances of hang-over exams. When we reviewed the process exception documents that were created before the implementation of the commercial RIS, we found various reasons for hang-over, which are listed in Table 7. We felt that a good percentage of these, if not all, were results of suboptimal information collection in upstream steps and could be eliminated by relocating various information-gathering tasks using the new RIS.

Table 7. Reasons Borg's Radiologists Delayed Interpretation

Sample Reasons for Delaying Interpretation
(a) "Missing (the details of) last physical"
(b) "DB error at PACS"
(c) "Additional views (not documented / explained by tech)"
(d) "Missing (previous) study"
(e) "No (referral) MD notes on swelling"
(e) "Unknown physical scar"
(f) "An unexplained surgical scar"
(g) "Call back patient (for retake)"

(h) "Prior exam (found), but wrong images (provided)"

To test the effectiveness of the RIS-initiated process changes, we compared the hang-over rate (as obtained from the process-exception documents) before the installation with the rate after. The percentage of interpretations delayed for mammography dropped from a pre-RIS level of 7.92% to a post-RIS level of 2.56% (significant at a 5% level). For MRI, however, the percentage decreased only slightly, from 3.74% to 2.48% (not significant at a 5% level). Thus, the data provides empirical support for our *conjecture* that the differential impact on the interpretation turnaround time is correlated with the impact on the hang-over rate. We also double-checked the actual RTAT distributions to confirm this finding. The actual distributions for mammography are shown in Figure 2, along with the respective fitted lognormal densities. Apparently, the RIS shortened the long right tail of the distribution.

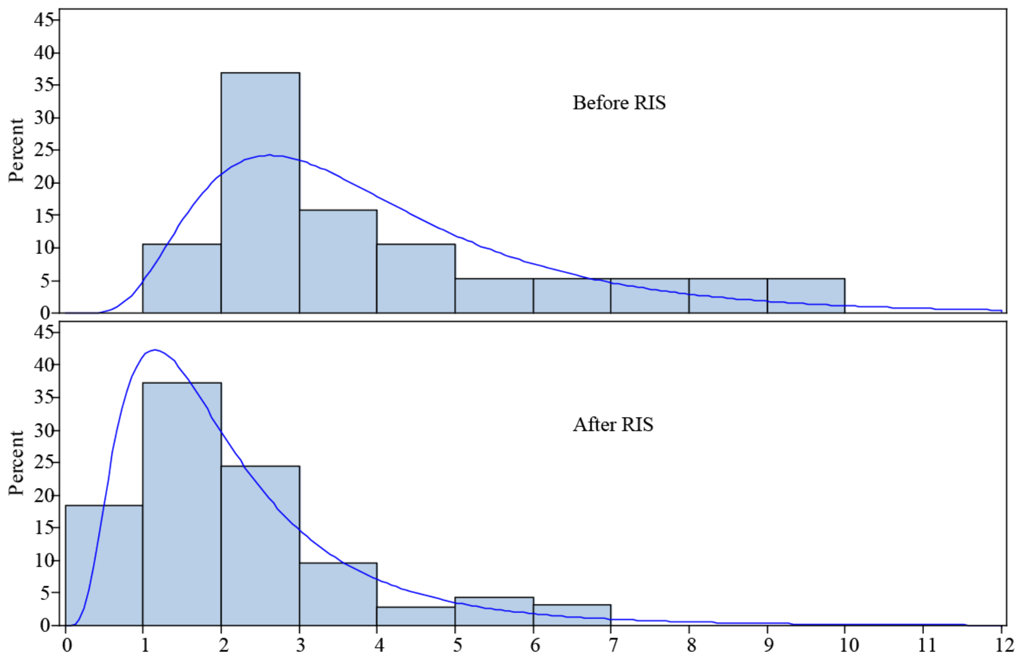

Figure 2. Pre-RIS and Post-RIS Distributions of Mammography RTAT in Hours: Histograms along with Fitted Lognormal Densities

We gain several fundamental managerial insights from this finding. First, reducing hang-over is critical to improving workflow performance. Second, a higher hang-over rate prior to workflow redesign means a greater opportunity to realize an improvement. Mammography was more prone to hang-over, so it benefited more. There were reasons mammography was more hang-over-prone. Consider, for example, an MRI scan of a recently fractured elbow. Its interpretation would need less background information than what would be needed for a diagnostic mammography exam. For the former, there would be no need for information such as the history of carcinoma in the patient's family, past mammograms, the history of breast surgeries, biopsies, etc. Third, an influential driver of the benefits of upstream information gathering is the size of the hang-over-related delay, i.e., the time involved in gathering missing information. If this delay is larger, the benefits of reducing hang-over would also be greater, because the improvements in operational metrics arise in part from mitigating the impact of this delay.

Another major implication is that enterprise system users and vendors who rely on measuring service times at the bottleneck workstation (e.g., the radiologist's cycle time in the case of RIS) often do not fully understand the benefits of their information systems. The utilization of the server is indicative of system performance, only when the system is linear. For systems with hang-over loops, mitigating hang-overs and information-gathering delays is essential to improving performance and achieving superior customer service. We elaborate on this issue in §6.

5.2. Transcription

Despite requiring radiologists to pre-specify reporting templates (which are called *standard normals* by the medical staff) at the time of interpretation, we did not observe any significant savings in the transcriptionists' turnaround time (see estimates of (2) in Table 3). Both the estimate of β_2, which shows the impact on MRI, and that of $\beta_2 + \beta_3$, which captures the impact on mammography, turned out to be insignificant. There were no impacts on either modality. Upon further investigation, we learned that Borg had already been using homegrown reporting templates that were similar to the templates provided by the RIS, which was probably the reason using RIS to gather the reporting template information in the clinical interpretation step had little impact on the average transcription TAT.

5.3. Scheduling

The RIS undoubtedly delivered substantial benefits with regard to turnaround times, leading to significant performance improvements for certain participants (such as the radiologists for the mammography modality). Yet, it had an adverse impact on the schedulers. Both the call length and the abandonment rate increased. The increase in the call length was larger for new patients because the new system required that additional demographic and financial (e.g., insurance) information be collected from them up front.

This adverse impact on scheduling is a reminder that the improvement in the radiologist's performance was obtained partially at the expense of those upstream, namely, the scheduling staff. The main lesson is, therefore, that leveraging an enterprise-wide technology like RIS can have negative side-effects, and mitigating them may require additional training as well as realignment of performance standards and incentive structures. Additionally, information-gathering efforts upstream may also require a reallocation of staff from downstream steps (e.g., check-in) to avoid service issues later on. Another salient lesson is that structured information-gathering capabilities of electronic systems create the possibility of shifting work from more expensive resources, such as radiologists and the administrative and clinical staff assisting them, to less expensive resources, such as the call-center staff. With technology, it thus becomes possible for the clinical staff to do what they are always advised to do—which is to "delegate all duties that don't require a physician's license" (McKee 2004).

5.4. Technologist's Administrative Time

A reason for the surprising decline in the technologist's administrative time is likely that much of the clinical data concerning contrast media sensitivity, metal parts or implants in the patient's body, and the like were collected at the time of scheduling. Moreover, the RIS eliminated the need to copy patient information from one paper form to another. Thus, the data on administrative times indicate that the leverage from disciplined data collection and elimination of duplicate data entry compensated the technologists adequately for the additional reporting that was assigned to them following the RIS installation (such as documenting in detail reasons for capturing additional views).

Here as well, the modality that had a greater inefficiency initially, mammography, benefited more. As already mentioned, we observed a statistically significant decline in the administrative time for mammography, but only an insignificant change for MRI. Thus, we again find that the performance of a downstream step is more sensitive to upstream information gathering for some modalities than it is for others.

5.5. Patient and Referring Physician Satisfaction

One concern prior to the study was whether the improvement in the average report turnaround time would translate to improved referring physician satisfaction. The literature often assumes linear waiting costs and concludes that any decline in the average waiting time also implies a lower waiting cost for the customer. However, this assumption may not necessarily hold in the case of healthcare. For example, some referring physicians are often willing to wait up to a certain point, but they find any waiting beyond that point unacceptable. This is why examining the survey data was important.

As indicated by Table 5, there was a statistically significant improvement in the percentage of referring physicians who rated Borg 6 or 7 in dimension 6, *timely delivery of clinical reports*. The implication is that the improved operational metrics indeed translated to gains in satisfaction. Besides, our PCA also showed a significant positive correlation between the principal component and the dummy variable *RIS*, further corroborating the beneficial impact of the new RIS.

We observed small changes in patient satisfaction. One possible reason for such limited changes was that the satisfaction levels were very high to begin with. Also, many of the dimensions were not directly affected by the RIS. The RIS, however, had an impact on gathering of insurance information at the scheduling step, which is why we observed a net positive impact on dimension 3, *efficiency of check-in and registration process.*

Likewise, we saw an improvement in dimension 7, *cost information provided by Borg's staff.* This improvement was again a result of better gathering of insurance information up front. We collected additional data on co-payments to confirm this finding. Prior to the RIS installation, front desk representatives often overcharged if they did not have the complete insurance information for a patient. This policy meant that Borg had to issue refund checks to many patients later on, which was detrimental to their satisfaction. The total amount refunded typically exceeded $12,000 per month before the RIS installation, but it dropped dramatically to about $3,000 a month afterward. Table 8 shows how it dropped following the go-live in late October.

Table 8. Total Co-pay Refund and Number of Refund Checks Issued after RIS; Go-live date: Last week of October 2006; December 2006 – February 2007 data contains exams scheduled both pre- and post-RIS, and March 2007 data contains only those scheduled post-RIS.

Month	**Amount refunded**	**Number of refund checks issued**
December 2006	$12,386.96	395
January 2007	$10,751.01	429
February 2007	$8,473.27	304
March 2007	$2,947.47	156

During December-February, refunds were issued for many exams that were scheduled before the RIS installation. Refund checks issued in March 2007 were almost entirely for exams that were scheduled post-RIS. Evidently, gathering insurance information up front had a favorable effect on the accuracy of co-pay collection.

6. Server Utilization vs. Turnaround Time

Unexpectedly, the RIS implementation at the clinic faced resistance from many radiologists; they complained that the RIS increased their workload as they had to navigate through several screens to retrieve the pertinent information for each study. The requirement to navigate through the screens of the new RIS might have indeed affected the radiologists adversely.[6] Some radiologists complained, "We don't need so many screens—look, all these fields are not even used." Others claimed that they were burdened with too many unnecessary mouse-clicks. Many radiologists also pointed out, "We can't see both sides of the requisition at the same time—we need to page down to see the other side." Our observations of the new RIS corroborated many of their concerns regarding its user interface. Since the radiologists were paid on a "per case" basis, it was only natural that they resisted changes that could increase their cycle times and hurt their incomes. Hence, using analytical tools, we explored the consequences of adversely affecting the utilization of the bottleneck server. Specifically, we investigated whether faster turnaround could be

[6] Since the radiologist cycle (service) time data are extremely sensitive, we were not allowed to collect that data in detail. The actual impact on the cycle time is therefore unknown.

achieved even when the utilization of the radiologist's workstation and the average queue length there increased.

Our model is that of a feedback queue, similar to "rework" models in the literature. Hang-over is, however, different from rework. Rework is done to address quality issues such as defects (Jewkes 1994), and it has nothing to do with information-gathering tasks in the workflow. Rework models in the operations literature typically use Bernoulli feedbacks, i.e., the number of rounds of service required for an arriving customer follows a geometric distribution (Seidmann and Nof 1985). The feedback process is often assumed to be instantaneous; i.e., work items that do not meet the quality standard are fed back immediately to the service queue. Hang-over means postponing service because the necessary information is not available. It requires just two rounds of service: one round to identify the missing information set, and another to complete the work. The feedback process is not instantaneous, as the service cannot be completed until all missing data items are gathered. In order to account for these differences, we propose a new queuing model. Using the model, we explain why server utilization does not fully reveal the impact of addressing the hang-over problem: we show that server utilization can increase even as the average turnaround time decreases.

6.1. FCFS Queue

We model the radiologist's workstation as a single-server "First Come First Served" (FCFS) queue with delayed feedback (Figure 3). We employ this structure because, except for medical emergencies, a radiologist would process cases in the order received. We assume a Poisson arrival rate of λ per hour. The symbol p denotes the *no hang-over rate*; i.e., a fraction p of the exams does not suffer hang-overs, and the remaining fraction $(1 - p)$ does. We also refer to the fraction $(1 - p)$ as the *hang-over rate*.

The service time distributions are all assumed to be *exponential*: for exams requiring only one round of service, the rate is μ_{11}; for hang-over exams, the rate for the first round is μ_{12}, and the rate for the second round is μ_{22}. An exam that suffers a hang-over is assumed to rejoin the queue after a random delay having a mean of $1/d$. Henceforth, we refer to $1/d$ as the *hang-over-related delay* or the *feedback delay*. In general, $\mu_{11} \neq \mu_{12} \neq \mu_{22}$, and the system may not be product-form (Lahiri and Seidmann 2010).

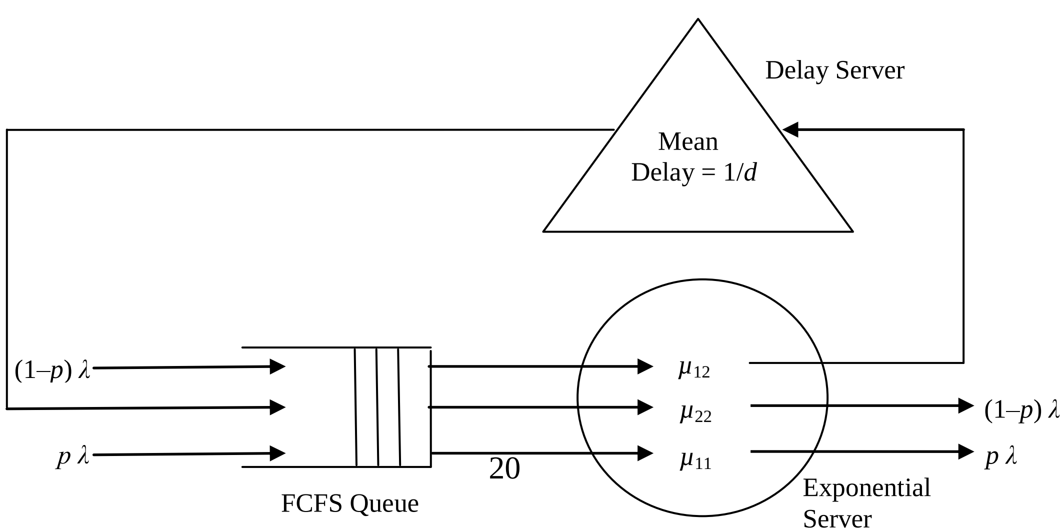

Figure 3. Queuing Model of the Radiologist's Workstation (Clinical Interpretation Step)

The sojourn time—the interpretation TAT, or the time elapsed from the point an exam joins the FCFS queue to the point it leaves the server after service—of an exam requiring one round of service is denoted S_{11}. The *successive* sojourn times of an exam requiring two rounds of service are denoted by S_{12} and S_{22}. The expected first-round sojourn time of a randomly selected exam is denoted by S_1; therefore, $E[S_1] = pE[S_{11}] + (1 - p)E[S_{12}]$. The expected sojourn time (or the average interpretation TAT) of a randomly selected exam, denoted $E[S]$, is, therefore, $E[S_1] + (1 - p)\,(E[S_{22}] + 1/d)$.

Besides the expected sojourn times, we are also interested in the utilization, ρ, which equals $\rho_{11} + \rho_{12} + \rho_{22}$. Here, $\rho_{11} = \lambda p/\mu_{11}$, $\rho_{12} = \lambda(1 - p)/\mu_{12}$, and $\rho_{22} = \lambda(1 - p)/\mu_{22}$. Because the system is not product-form, we use operational analysis (Denning and Buzen 1978).

Proposition 1 *For the FCFS queue,* $E[S_{11}]$, $E[S_{12}]$, *and* $E[S_{22}]$ *are as follows:*

$$E[S_{12}] = \frac{\frac{1}{\mu_{12}} + \rho_{11}\left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right) + \frac{\rho_{22}}{\mu_{22}}}{(1+\rho_{22})(1-\rho)},\ E[S_{11}] = E[S_{12}] + \left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right), \text{and } E[S_{22}] = \rho E[S_{12}] + \frac{1}{\mu_{22}}.$$

Proof of Proposition 1. The expected sojourn times can be derived by directly solving (5) – (7) below, and by applying the identity: $1 - (\rho_{11} + \rho_{12}) - \rho\rho_{22} = (1 + \rho_{22})\,(1 - \rho)$.

$$E[S_{12}] = \frac{1}{\mu_{12}} + \rho_{11}E[S_{11}] + \rho_{12}E[S_{12}] + \rho_{22}E[S_{22}] \tag{5}$$

$$E[S_{11}] = \frac{1}{\mu_{11}} + \rho_{11}E[S_{11}] + \rho_{12}E[S_{12}] + \rho_{22}E[S_{22}] = E[S_{12}] + \left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right) \tag{6}$$

$$E[S_{22}] = \frac{1}{\mu_{22}} + \rho_{11}E[S_{12}] + \rho_{12}E[S_{12}] + \rho_{22}E[S_{12}] = \rho E[S_{12}] + \frac{1}{\mu_{22}} \tag{7}$$

Note that (5) and (6) follow from the PASTA (Poisson Arrivals See Time Averages) property and Little's Law. Upon arrival, a new exam sees $\lambda pE[S_{11}]$ no-hang-over exams waiting for service or being serviced. The expected time needed to complete them is $\lambda pE[S_{11}]/\mu_{11} = \rho_{11}E[S_{11}]$. The last two terms in each equation are similar. Hence, the sum $\rho_{11}E[S_{11}] + \rho_{12}E[S_{12}] + \rho_{22}E[S_{22}]$ equals the expected waiting time for the first round of service.

To obtain (7), we note that the PASTA property does not apply to a hang-over exam returning for service, because it does not return according to a Poisson process. The time when a hang-over exam arrives initially and the time when it returns are on average $E[S_{12}] + 1/d$ apart. During this period $\lambda p(E[S_{12}] + 1/d)$ non-hang-over exams are expected to arrive. However, while the hang-over exam waits in search of information, $\lambda p/d$ of them are expected to leave. Therefore, when it returns, it would see $\lambda pE[S_{12}]$ non-hang-over exams in the system. The expected time needed to serve them would be $\lambda pE[S_{12}]/\mu_{11} = \rho_{11}E[S_{12}]$. A similar argument explains the last two terms as well. $\blacksquare$

It is important that we note that, if $\mu_{11} = \mu_{12} = \mu_{22} = \mu$, Proposition 1 expectedly becomes the well-known product-form solution $E[S_{11}] = E[S_{12}] = E[S_{22}] = 1/(\mu\,(1 - \rho))$, where $\rho = (2 - p)\,\lambda/\mu$.

Before we can use Proposition 1 to illustrate the tradeoff between the utilization and the expected sojourn time, we need to ask what values are reasonable for different model parameters. First, note that, when the interpretation of an exam is put on hold, the time the radiologist spends fetching the image, the time he spends pulling up relevant information and prior images, and all other cognitive efforts associated with the analysis of the exam are partly lost. In addition, the radiologist must write a query to the technologist or the administrative assistant discussing the details of the missing information that must be collected before a final interpretation can be done. This extra work implies inefficiencies or a higher average service time per hang-over exam, which means that $1/\mu_{12} + 1/\mu_{22} \geq 1/\mu_{11}$.

Let $\mu_{11} = \mu$ and $\mu_{12} = \mu_{22} = 2\mu/\alpha$, where $\alpha \geq 1$ is the *inefficiency factor*. Note that $\mu_{12} = \mu_{22}$ does not hold in general, but it is a reasonable simplification here. In the first round, the radiologist analyzes the available information and spells out what information is needed, and, in the second, he analyzes the new information and dictates the diagnosis. Both rounds thus have the same structure: an analysis followed by a summarization of clinical findings. Additionally, we limit ourselves to $\alpha \geq 1$ because we want to capture the inefficiency described in the preceding paragraph. We now express the expected sojourn time as a function of the utilization (ρ) and the hang-over rate ($1 - p$). Let $E[S] = \hat{S}(\rho, p)$. Using Proposition 1, we can show the following:

$$\hat{S}(\rho, p) = (1 - p)/d + \frac{\rho\left(4(\alpha(1-p)+p)^2 + 2(1-p)(\alpha^2 + 2p - 2\alpha p)\rho + (\alpha-2)^2(1-p)p\rho^2\right)}{2\lambda\left(\alpha(1-p)+p\right)(1-\rho)(2p+\alpha(1-p)(2+\rho))}$$

It is worth noting here that, because $\alpha \geq 1$, a higher p (i.e., a lower hang-over rate) implies an identical or lower utilization (ρ), unless the reduction in p comes off at the expense of additional information-retrieval work on the part of the server (such as navigating through a number of computer screens or tabs). If this additional work exactly offsets the beneficial impact of a higher p, the utilization of the server would remain the same. If the additional work more than offsets this reduction, the utilization would increase; otherwise, the utilization would decrease.

We now consider a numerical example (Figure 4) to show how the expected sojourn time varies with the hang-over rate for different server utilization levels. Suppose that the hang-over rate is 10% and the utilization is 85% initially, so that the expected sojourn time is approximately 0.779 hours. Also, suppose that the utilization rises to 90% after the implementation (because of increases in service times, or, equivalently, a decrease in μ). As shown in the figure, the expected sojourn time does not necessarily increase—if the hang-over rate falls to 4% or below (the shaded region in the figure), it decreases. A lower sojourn time can, therefore, be realized despite an increase in the server utilization.[7]

Hang-over Rate	Utilization (ρ)										
(1-p)	85%	86%	87%	88%	89%	90%	91%	92%	93%	94%	95%
10%	0.779	0.811	0.848	0.891	0.942	1.003	1.077	1.171	1.290	1.450	1.674
9%	0.739	0.771	0.808	0.851	0.902	0.963	1.037	1.130	1.250	1.410	1.633
8%	0.699	0.731	0.768	0.811	0.861	0.922	0.997	1.090	1.210	1.369	1.593
7%	0.659	0.691	0.728	0.770	0.821	0.882	0.957	1.050	1.169	1.329	1.552
6%	0.619	0.651	0.687	0.730	0.781	0.842	0.916	1.009	1.129	1.288	1.511
5%	0.579	0.610	0.647	0.690	0.741	0.802	0.876	0.969	1.088	1.248	1.471
4%	0.538	0.570	0.607	0.650	0.701	0.761	0.836	0.928	1.048	1.207	1.430
3%	0.498	0.530	0.567	0.610	0.660	0.721	0.795	0.888	1.007	1.166	1.389
2%	0.458	0.490	0.527	0.569	0.620	0.681	0.755	0.848	0.967	1.126	1.348
1%	0.418	0.450	0.486	0.529	0.580	0.640	0.714	0.807	0.926	1.085	1.308

Figure 4. The Expected Sojourn Time, $\hat{S}(\rho, p)$, in Hours; α (Inefficiency factor) = 1.2, λ (no. of exams process/ hr) = 15/hr, $1/d$ (Mean hang-over delay) = 4 hrs; In the shaded region, the expected sojourn time is lower than 0.779, its initial value.

Using our model, we can also critically examine the three components of the sojourn time—the waiting time in the FCFS queue, the information-gathering delay, and the service time at the server. Figure 5 shows these components for the two scenarios "circled" in Figure 4. It illustrates that mitigating hang-over can lead to a lower expected sojourn time even when doing so adversely affects the FCFS queue at the bottleneck server. In other words, a reduction in the hang-over rate can cause the server (in our case, the radiologist) to experience a larger backlog even when its impact on the expected sojourn time is beneficial. The lesson here is that the impact on the server does not fully reveal the impact of the hang-over effect. It is important to look at the entire workflow end-to-end, particularly the impact of the feedback loop. The implication is also that it may be useful for healthcare organizations wanting to reduce the expected turnaround time to consider IT solutions that can affect the bottleneck server adversely.

[7] In this example, we have used 4 hours as the mean hang-over delay ($1/d$). In reality, this delay can be much longer; e.g., it can be several days before the medical imaging staff tracks a patient again. The qualitative insights obtained remain the same regardless of the value used.

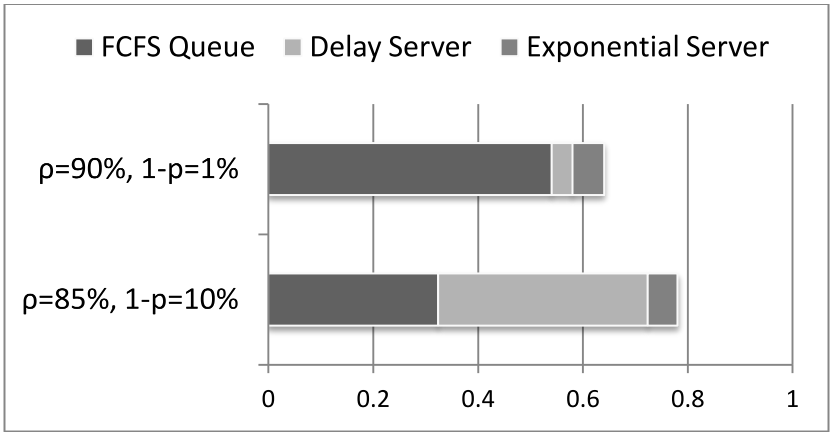

Figure 5. Components of the Sojourn Times for the Two “Circled” Scenarios in Figure 4

A consequence of increased congestion at the bottleneck server would also be a higher expected turnaround time for non-hang-over exams. Therefore, a reduction in the hang-over problem may adversely affect certain customer classes even when it improves the overall performance.

The adverse impact on non-hang-over customers would occur even if server utilization is the same in the pre- and post-RIS scenarios. The intuition is as follows. When the utilization remains the same or increases but the hang-over rate declines, a non-hang-over exam suffers a larger waiting time because it sees a larger number of non-hang-over exams in the FCFS queue on arrival. The server completes these non-hang-over cases (as opposed to parking them) before processing the newly arrived exam, which increases its waiting time in the queue.

The only way to reduce the expected sojourn times for non-hang-over exams would be to reduce the radiologist utilization. However, a reduction is often not possible in many clinics. For example, if radiologists in an outpatient clinic feel that they are getting less work, they would ask for a bigger throughput—as they are paid piece-rate for their work.

6.2. FCFS Queue with Batching

Batching of work is still common in many healthcare organizations where the work is not handed off electronically. In those organizations, work items move to the server in batches, and each batch is handed off to the next step after all items in the batch are processed. This is also the common practice in many outpatient clinics for screening mammography and other less critical modalities. It is, therefore, important that we examine whether our observation, that a lower turnaround is possible despite increases in the bottleneck server’s utilization, also holds for organizations that use batching.

Specifically, we examine the case in which work is moved to the server (e.g., the radiologist’s workstation) in batches of size Q. After the batch receives its first-round service, the completed cases move to the next step (i.e., transcription), while the hang-over cases wait for additional information. We do not assume any batching for cases feeding back: An incomplete case feeds back as and when the information required to fully process it becomes available. We also assume that the server first processes all batches that arrived before its return and then processes it in order.

Let S_Q denote the first-pass sojourn time of a randomly selected arriving batch and S_{Q2} the first-pass sojourn time of a batch containing hang-over exams. As before, S_{22} denotes the second-round sojourn time, i.e., the time elapsed from the point a case rejoins the FCFS queue to the point it leaves the system. The expected sojourn time of this system, $E[S]$, is $E[S_Q] + (1 - p)\,(E[S_{22}] + 1/d)$.

The utilization for this model, ρ, is $\rho_Q + \rho_{22}$, where $\rho_Q = (\lambda/Q)*(Q\,(p/\mu_{11} + (1-p)/\mu_{12})) = \lambda p/\mu_{11} + \lambda(1-p)/\mu_{12} = \rho_{11} + \rho_{12}$, and $\rho_{22} = \lambda(1-p)/\mu_{22}$.

Proposition 2 *For the FCFS queue with batching as described above:*

$$E[S_{Q2}] = E[S_{12}] + \Lambda\left(\frac{1-\rho_Q}{(1+\rho_{22})(1-\rho)}\right),\ E[S_Q] = \frac{(1-p)p^Q Q}{1-p^Q}\left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right) + E[S_{Q2}],$$

$$E[S_{22}] = \rho E[S_{Q2}] + \frac{1}{\mu_{22}},\ \textit{where}\ \Lambda = \frac{Q\left(\frac{p}{\mu_{11}} + \frac{1-p}{\mu_{12}}\right) - p^Q\,\frac{Q}{\mu_{11}}}{1-p^Q} - \frac{1}{\mu_{12}},$$ *and* $E[S_{12}]$ *is as given by Proposition 1.*

Proof of Proposition 2. Let the random variable X_Q denote the service time of a batch. Its first and second moments are as follows.

$$E[X_Q] = Q\left(\frac{p}{\mu_{11}} + \frac{1-p}{\mu_{12}}\right) \tag{8}$$

$$E[X_Q^2] = 2Q\left(\frac{p}{\mu_{11}^2} + \frac{1-p}{\mu_{12}^2}\right)$$

Let W_Q be the waiting time seen by an arriving batch. Applying the Pollazcek-Khinchin formula, we can express it as follows.

$$\begin{aligned} W_Q &= \rho_Q W_Q + \frac{\rho_Q E[X_Q^2]}{2E[X_Q]} + \rho_{22}E[S_{22}] \\ &= \rho_Q W_Q + \left(\frac{\rho_{11}}{\mu_{11}} + \frac{\rho_{12}}{\mu_{12}}\right) + \rho_{22}E[S_{22}] \end{aligned} \tag{9}$$

The first term on the right-hand side of the expression above is the time needed to serve batches waiting in the queue. The second term is the expected remaining service time of a batch in service,

$E[X_Q^2]/(2E[X_Q])$, times the probability that the arriving batch finds one in service upon arrival, ρ_Q. The third term is the time needed to serve hang-over exams that have fed back. Note that $E[S_{22}] \neq 1/\mu_{22} + W_Q$, because hang-over exams do not return according to a Poisson process.

Let us now denote the service time of a batch that contains one or more hang-over exams by X_{Q2}. Since p^Q is the probability that a batch has no hang-over exams, we have the following result.

$$p^Q \frac{Q}{\mu_{11}} + (1 - p^Q) E[X_{Q2}] = E[X_Q] \tag{10}$$

Using (8) and (10), we can solve for $E[X_{Q2}]$ and $E[X_Q]$.

$$E[X_{Q2}] = \frac{Q\left(\frac{p}{\mu_{11}} + \frac{1-p}{\mu_{12}}\right) - p^Q \frac{Q}{\mu_{11}}}{1 - p^Q} = \Lambda + \frac{1}{\mu_{12}} \tag{11}$$

$$E[X_Q] = E[X_{Q2}] + \frac{Qp^Q(1-p)\left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right)}{1 - p^Q} \tag{12}$$

Using (11) and (12), we can express $E[S_Q]$ and $E[S_{Q2}]$ as follows.

$$E[S_{Q2}] = E[X_{Q2}] + W_Q = \Lambda + \frac{1}{\mu_{12}} + W_Q \tag{13}$$

$$E[S_Q] = E[X_Q] + W_Q = \frac{Qp^Q(1-p)\left(\frac{1}{\mu_{11}} - \frac{1}{\mu_{12}}\right)}{1 - p^Q} + E[S_{Q2}] \tag{14}$$

As we did before in the case of (7) in Proposition 1, here as well, we use operational analysis to obtain $E[S_{22}]$.

$$E[S_{22}] = \frac{1}{\mu_{22}} + \rho_Q E[S_{Q2}] + \rho_{22} E[S_{Q2}] = \rho E[S_{Q2}] + \frac{1}{\mu_{22}} \tag{15}$$

We get the expression for $E[S_{Q2}]$ by solving (9), (13), and (15), and by using the identity: $1 - \rho_Q - \rho\rho_{22} = 1 - (\rho_{11} + \rho_{12}) - \rho\rho_{22} = (1 + \rho_{22})(1 - \rho)$. The rest follows from (14) and (15). ∎

It is important to note that, when $Q = 1$, Λ is equal to 0, and, as a result, our two propositions are consistent, i.e., $E[S_{Q2}] = E[S_{12}]$ and $E[S_Q] = E[S_1]$.

We can, as before, express the expected sojourn time, $E[S]$, as a function of ρ and p. Let us again denote this function by $\hat{S}(\rho, p)$. Figure 6 shows how this expected sojourn time, $\hat{S}(\rho, p)$, varies with p for different server utilization levels (ρ). As before, let us assume that the hang-over rate is 10% and the utilization is 85% initially. The expected sojourn time is approximately 1.038 hours before the RIS implementation. We again examine what happens when the utilization rises to 90% following the implementation. As shown in the figure, the expected sojourn time in this case also falls if the new hang-over rate is 4% or less, just as it does when there is no batching.

The hang-over reduction, however, is not as effective with batching as it is without. When the hang-over rate falls to 1% and the utilization rises to 90%, the expected sojourn time becomes 0.872 hours, which is 16% lower than its initial value of 1.038 hours. In the absence of batching, this improvement is 18% (see Figure 4). If the batch size is increased to 10, we would see an even smaller improvement of 14%. Therefore, the implication of Proposition 2 is that hang-over reduction is less effective for workflows that use batching. The intuition is as follows. A large batch size affects the queuing delay at the FCFS server adversely and thereby makes any additional adverse impact on the utilization relatively more costly from a performance standpoint.

Hang-over Rate	**Utilization (ρ)**										
(1-p)	**85%**	**86%**	**87%**	**88%**	**89%**	**90%**	**91%**	**92%**	**93%**	**94%**	**95%**
10%	1.038	1.079	1.126	1.180	1.243	1.319	1.412	1.526	1.674	1.869	2.142
9%	0.994	1.034	1.080	1.133	1.196	1.271	1.361	1.475	1.619	1.812	2.081
8%	0.949	0.989	1.034	1.086	1.148	1.222	1.311	1.422	1.565	1.754	2.018
7%	0.904	0.943	0.988	1.039	1.100	1.172	1.260	1.370	1.510	1.696	1.956
6%	0.859	0.898	0.941	0.992	1.052	1.123	1.209	1.317	1.455	1.637	1.893
5%	0.814	0.852	0.895	0.945	1.003	1.073	1.158	1.264	1.399	1.578	1.829
4%	0.769	0.806	0.848	0.897	0.955	1.023	1.107	1.210	1.343	1.519	1.764
3%	0.723	0.760	0.801	0.849	0.906	0.973	1.055	1.156	1.286	1.459	1.699
2%	0.678	0.714	0.754	0.801	0.857	0.923	1.003	1.102	1.229	1.398	1.633
1%	0.632	0.667	0.707	0.753	0.807	0.872	0.950	1.047	1.172	1.337	1.567

Figure 6. The Expected Sojourn Time, $\hat{S}(\rho, p)$, in Hours; α (Inefficiency factor) = 1.2, λ (no. of exams process/hr) = 15/hr, Q (Batch size) = 5, $1/d$ (Mean hang-over delay) = 4 hrs; In the shaded region, the expected sojourn time is lower than 1.038, its initial value.

In short, the benefits of hang-over reduction are not limited to systems that do not use batching, although the benefits are lower in percentage terms for systems with higher batch sizes. We have already found that a reduction in hang-over may not have the same impact on every modality (e.g., mammography vs. MRI) or on every customer class (e.g., non-hang-over vs. hang-over); we now find that its impact may even vary depending on the batch size.

7. Conclusions

The central question that we attempt to answer in this work is quite common to many health services: Should more effort be spent in implementing systems for gathering patient information upstream in the clinical process to avoid waiting for the necessary medical information downstream? This question is relevant for two reasons. First, health services, although information-intensive, face many information-deficiency problems. Second, process-changing technologies—including those that make upstream data collection easier—are now witnessing widespread adoption within the healthcare field, which makes it important to have insights into when operational benefits are more likely to ensue from them. We address this issue first by empirically examining the impact of a commercial RIS on a radiology workflow. We then supplement our empirical findings with a queuing model of the bottleneck interpretation step. We gain several interesting insights, which have implications not only for radiology practices but also for other health management systems:

1. *Efforts to redesign healthcare workflows should emphasize hang-over reduction:* Any effort to redesign an information-intensive workflow should not be limited to configuration-related aspects, such as moving from a tandem to a parallel setup, consolidating certain steps, or adding case managers. It should consider the way information is gathered and shared in the workflow. Our empirical findings establish that properly designing an information-intensive healthcare workflow requires addressing the hang-over problem, which results from suboptimal information gathering upstream in the workflow. Hang-over problem is indeed a major impediment to achieving superior operational performance. Our empirical findings concerning mammography show that, even when a small percentage of exams hang over, it can significantly hurt the overall operational performance: When the hang-over rate declined from an approximately 8% to about 2.5% following the RIS implementation, it led to a whopping 50% savings in the average report turnaround time. Also, as noted in the introduction, hang-overs often have serious health consequences. Therefore, in the case of healthcare services, reducing hang-over is important from both perspectives—the perspective of providing superior service as well as that of improving medical outcomes. Interestingly, from personal interactions with a European bank, we have learned that hang-overs can cause serious delays in banks as well, particularly in the case of corporate customers for whom all relevant legal documents must be available at the time of certain transactions such as opening a new account. Hence, it is critical that managers make hang-over reduction a priority while redesigning their information-intensive service systems.

2. *Disciplined data collection upstream in the workflow can reduce hang-over:* A central insight of this research that reducing hang-overs require the workflow redesign strategy of <u>unbundling</u> information gathering from downstream clinical steps, such as clinical diagnosis, and moving them to upstream steps, such as scheduling. Significant externalities (interdependence) exist between the various steps in a healthcare workflow, because later steps are critically dependent on earlier steps for their information needs. Not collecting a critical data record in an earlier step significantly hurts later steps by causing them to hang over. Conversely, a slight addition to the workload in certain earlier steps can produce significant benefits downstream. The ability of new commercial enterprise systems to unbundle gathering of infor-

mation from its use provides tools necessary not only for shifting work to cheaper resources (e.g., from radiologists to schedulers) but also for controlling these externalities. In our study, better data gathering using RIS at the time of scheduling not only improved the radiologist's performance but also had significant beneficial impacts on the technologists' administrative time and front desk registration (co-pay collection). It also led to greater customer satisfaction in some cases. An important insight here is that, if a healthcare practice is experiencing significant hang-over delays in downstream clinical steps despite using sophisticated enterprise systems, it is probably not using its systems effectively to gather critical data items up front. Likewise, a healthcare practice acquiring a new enterprise system needs to recognize that the system will not deliver its full value unless it is properly leveraged to improve information gathering in the workflow. A secondary insight is that shifting information-gathering work upstream can adversely affect upstream steps (e.g., scheduling), and mitigating negative side effects on them could require additional realignments of staffing levels and compensation structures.

3. *Disciplined data collection upstream can have differential impacts even within the same organization:* On the upside, we found that disciplined data collection led to statistically significant performance improvements, including substantial improvements at the bottleneck radiologist interpretation step. On the downside, the impact of disciplined information gathering varied: mammography benefited substantially in most cases while the impact on MRI was largely muted. The implication is that the benefits of disciplined data collection can vary across modalities and workflows—even within the same organization. In general, when efficiencies already exist or the demand for information is less, implementing expensive new technologies or processes would be less effective in boosting performance. In the practice that we examined, most MRI exams did not require as much background clinical information as diagnostic mammography exams did—information such as the history of carcinoma in the patient's family, BIRADS levels reported in prior tests, a history of breast surgeries, etc. Consequently, mammography had a higher hang-over rate to begin with, which meant a greater opportunity to achieve substantial reductions in hang-over delays. We empirically found that the improvements in the interpretation turnaround times for the two modalities were indeed correlated with the respective improvements in their hang-over rates. A statistically significant reduction of nearly 5.5% in the hang-over rate for mammography not only shortened the long right tail of the distribution of its report turnaround time but also resulted in a sharply lower average turnaround time and greater referring physician satisfaction. On the other hand, the corresponding reduction for MRI was statistically insignificant, as was the impact on its report turnaround time.

4. *It is important to look at the impact on the system end-to-end:* Using a queuing system model of the bottleneck interpretation step, we show that mitigating hang-over can reduce the expected turnaround time even when doing so leads to higher utilization for the bottleneck server. Technologies like RIS can burden bottleneck servers by requiring them to navigate through a complex screen sequence. However,

we find that, despite this adverse impact on the server, a technology may lead to faster service overall if it is leveraged properly to reduce hang-over. In the numerical example presented in §6.1, even when the bottleneck server's utilization rises from 85% to 90%, the average turnaround time declines if the hang-over rate falls to 4% or below from its initial level of 10%. Using a bar chart we further show that, while higher server utilization leads to more congestion delays at the server's queue, it is more than offset by reductions in hang-over delays if the new hang-over rate is sufficiently low. The lesson, in short, is that the impact on the bottleneck server's utilization does not accurately reflect the impact of an enterprise information system. Server utilization provides an accurate picture only when a system is linear but not when it has hang-over loops. This insight is particularly important, since healthcare staff often sees only the additional burden with technology solutions without understanding the full system benefits. It is imperative that healthcare organizations consider additional metrics, particularly the hang-over rate, for assessing the full system benefits of a new enterprise system.

5. *Benefits of hang-over reduction are relatively lower for systems with larger batch sizes:* The insight that mitigating hang-over can reduce the expected turnaround time even when doing so leads to a higher utilization for the bottleneck server also applies to workflow systems with batching, although the reduction in percentage terms is lower for a system with a larger batch size. In the numerical example presented in §6.2, when the bottleneck server's utilization rises from 85% to 90% and the hang-over rate falls from 10% to 1%, the resulting reduction in the expected turnaround time is nearly 18% for a system that uses no batching while it is about 16% for a system with a batch size of 5 and just 14% for that with a batch size of 10. An implication is thus that healthcare practices must also consider the effect of batching while estimating potential benefits of a hang-over reduction effort.

Our study has its limitations. We have looked at a large outpatient clinic having several locations. However, studying one organization is often not sufficient. Besides, it is not apparent that our findings can be extended to other types of clinics, for example, to inpatient clinics. Also, we focused on the diagnostic report turnaround time metric. There are other important issues that we do not consider, for example, the impact on the quality of reports. Quality is a huge concern, and it is particularly so for mammography (Adcock 2004). Another limitation is that we compare two modalities, MRI and mammography, but do not have the necessary data to break them down into subcategories. Additional research is essential to address these limitations.

There are many possible directions for future research. More sophisticated analytical models of hang-over, for instance, could allow more general arrival and service time distributions as well as other alternative service disciplines. There is also an immediate need for studying other instances of hang-overs. Our research provides solid evidence on the benefits and limitations of addressing hang-overs in a radiology workflow; however, as already mentioned, hang-overs are also common to many other clinical activities, including filling of prescriptions in pharmacies and pre-anesthesia evaluations in hospitals. Also, there are other information systems akin to RIS, such as Pharmacy Information Systems and Anesthesia Information Management Systems. What is lacking is detailed empirical and analytical research that explains whether these systems have any value with respect to their impacts on operational performance. Researchers should study these other clinical systems and find out what works and what does not. Anyone interested in exploring the benefits of hang-over elimination in other clinical activities will find practical guidelines in this paper.

Acknowledgments

Professor Zvi Lefkovitz, MD, Dr. Seth Borg, MD, Eastman Kodak Company, Borg Imaging LLC, Participants of the 2008 *Radiological Society of North America* Annual Meeting, and many generous colleagues.

References

Adcock, K. A. 2004. Initiative to Improve Mammogram Interpretation. *The Permanente Journal* **8**(2) 12-18.

Agarwal, R., G. Gao, C. DesRoches, A. K. Jha. 2010. Research Commentary — The Digital Transformation of Healthcare: Current Status and the Road Ahead. *Information Systems Research* **21**(4) 796-809.

American College of Radiology. 2008. ACR Practice Guideline for the Performance of Screening and Diagnostic Mammography. URL: http://www.radrounds.com/page/acr-practice-guidelines

Ayal M. and A. Seidmann. 2009a. On the Role of RIS/ PACS in Healthcare: An Empirical Study. *In Proceedings of the 42nd Hawaii International Conference on System Sciences*.

Ayal M. and A. Seidmann. 2009b. An empirical investigation of the value of integrating enterprise information systems: the case of medical imaging informatics. *Journal of Management Information Systems* **26**(2) 43-68.

Banker, R. D., R. J. Kauffman. 2004. The evolution of research on information systems: A fiftieth-year survey of the literature in management science. *Management Science* **50**(3) 281-298.

Brynjolfsson, E. 1993. The productivity paradox of information technology. *Commun. of the ACM* **36**(12) 67-77.

Brynjolfsson, E., L.M. Hitt. 1996. Paradox lost? Firm-level evidence on the returns to information systems spending. *Management Science* **42**(4) 541-559.

Brynjolfsson, E., S. Yang. 1996. Information technology and productivity: A review of the literature. *Advances in Computers* **43** 179-214.

Buzacott, J. A. 1996. Commonalities in business processes: Models and issues. *Management Science* 42(5) 768-782.

Chesson, E. 2006. Radiology is streamlining. *Health Imaging & IT* **August 6**.

Cotteleer, M. J., E. Bendoly. 2006. Order lead-time improvement following enterprise information technology implementation: An empirical study. *MIS Quarterly* **30**(3) 643–660.

Dedrick, J., V. Gurbaxani, K.L. Kraemer. 2000. Information technology and economic performance: A critical review of the empirical evidence. *ACM Computing Surveys* **35**(1) 1-28.

Denning, J., P. Buzen. 1978. The operational analysis of queuing network models. *Computing Surveys*. **10**(3) 225-261.

Denning, J., R. Medina-Mora. 1995. Completing the loops. *Interfaces* **25**(3) 42-57.

Devaraj, S., R Kohli. 2000. Information technology payoff in the health care industry: A longitudinal study. *Journal of Management Information Systems* **16**(4) 41-68.

Dewan, S., K. L. Kraemer. 1998. International dimensions of productivity paradox. *Commun. of the ACM* **41**(8) 56-62.

Dewan, S., C. Min. 1997. Substitution of information technology for other factors of production: A firm-level analysis. *Management Science* **43**(12) 1660-1675.

Dobson, G., E. J. Pinker, R. L. van Horn. 2009. Division of Labor in Medical Office Practices. *Manufacturing & Service Operations Management* **11**(3) 525-537.

Dobson, G., A. Sainathan. 2011. On the impact of analyzing customer information and prioritizing in a service system. *Decision Support Systems* **Forthcoming**.

Duerinckx, A. J., E. G. Grant. 1998. Cost of PACS and computed radiography in the United States. *Radiology* **208**(2) 554-555.

Green, L. V., S. Savin, B. Wang. 2006. Managing Patient Service in a Diagnostic Medical Facility. *Operations Research* **54**(1) 11-25.

Hasija, S., E. J. Pinker, R. A. Shumsky. 2005. Staffing and routing in a two-tier call center. *International Journal of Operations Research* **1**(1-2) 8-29.

Holt, N. F., D. G. Silverman, R. Prasad, J. Dziura, K. J. Ruskin, 2007. Pre-anesthesia clinics, information management, and operating room delays: results of a survey of practicing anesthesiologists. *Anesthesia and Analgesia* 104(3) 615-618.

Hitt, L.M., D.J. Wu, X. Zhou. 2002. Investment in Enterprise Resource Planning: Business Impact and Productivity Measures. *Journal of Management Information Systems* **19**(1) 71-98.

Jewkes, E. 1994. A Queuing Analysis of Priority-based Scheduling Rules for a Single-Stage Manufacturing System with Repair. *IIE Transactions* **26**(4) 80-86.

Karmarkar, U. S., U. M. Apte. 2007. Operations management in the information economy: Information products, processes, and chains. *Journal of Operations Management* **25**(2) 438-453.

King, R. 2009. How Kaiser Permanente Went Paperless. *BusinessWeek.com* **April 7**.

Koman, R. 2009. Tech stimulus: $20B for health care records. *ZDNet.com* **January 26**.

Lahiri, A., A. Seidmann. 2009. Analyzing the Differential Impact of Radiology Information Systems Across Radiology Modalities. *Journal of the American College of Radiology* **6**(20) 705-714.

Lahiri, A., A. Seidmann. 2010. The Hang-over Effect in Information-Intensive Service Systems. *In Proceedings of the 43rd Hawaii International Conference on System Sciences*.

Lehr, B., F. Lichtenberg. 1999. Information technology and its impact on productivity: Firm-level evidence from government and private data sources 1977-1993. *The Canadian Journal of Economics* **32**(2) 335-362.

Lu, L. X., J. A. van Mieghem, R. K. Savaskan. 2009. Incentives for Quality through Endogenous Routing. *Manufacturing & Service Operations Management* **11**(2) 254-273.

May, G. A., D. D. Deer, D. Dackiewicz. 2000. Impact of digital radiography on clinical workflow. *Journal of Digital Imaging* **13**(2 Suppl 1) 76-78.

McAfee, A. 2005. Pharmacy Service Improvement at CVS. *Harvard Business School Press*.

McAfee, A. 2002. The impact of enterprise information technology adoption on operational performance: An empirical investigation. *Production and Operations Management* **11**(1) 1-21.

McKee, K. 2004. Ways to work smarter. *Medical Economics* **81**(22) 36-38.

Menon, N. M., B. Lee, L. Eldenburg. 2000. Productivity of information systems in the health-care industry. *Information Systems Research* **11**(1) 83-92.

Mukhopadhyay, T., S. Kekre, S. Kalathur. 1995. Business value of information technology: A study of electronic data interchange. *MIS Quarterly* **19**(2) 137-156.

Powell, S. G. 2000. Specialization, teamwork, and production efficiency. *International Journal of Production Economics* **67** 205-218.

Reiner, B. I., E. L. Siegel, J. A. Carrino. 2002. Workflow Optimization: Current trends and future directions. *Journal of Digital Imaging* **15**(3) 141-152.

Ruskin, K. J. 2011. Anesthesia Information Management Systems for the Office Practitioner. *Institute for Safety in Office-based Surgery* URL http://isobsurgery.org/?page_id=195

Seidmann, A., A. Sundararajan. 1997. The effects of task and information asymmetry on business process design. *International Journal of Production Economics* **50** 117-128.

Seidmann, A., S. Y. Nof. 1985. Unitary manufacturing cell design with random product feedback flow. *IIE Transactions* **17**(2) 188-193.

Shumsky, R. A., E. J. Pinker. 2003. Gatekeepers and Referrals in Services. *Management Science* **49**(7) 839-856.

Siegel, E. L., B. I. Reiner, K. M. Siddiqui. 2004. Ten filmless years and ten lessons: a tenth anniversary retrospective from the Baltimore VA Medical Center. *Journal of the American College of Radiology* **1** 824-833.

Smith, P.C., R. Araya-Guerra, C. Bublitz, B. Parnes, L. M. Dickinson, R. van Vorst, J.M. Westfall, W. D. Pace. 2005. Missing clinical information during primary care visits. *Journal of the American Medical Association* **293**(5) 565-571.

Srinivasan, M., E. Liederman, N. Baluyot, R. Jacoby. 2006. Saving time, improving satisfaction: the impact of a digital radiology system on physician workflow and system efficiency. *Journal of Healthcare Information Management* **2** 123-131.

Varian, H. 2004. Information technology may have cured low service-sector productivity. *The New York Times* **February 12**.