

Abstract

Racz, Melanie Beth. The Effect of Cross-Training and Scheduling in an Inbound Call Center using Simulation. (Dr. Stephen Roberts and Dr. Xiuli Chao, Co-Chairs)

This thesis presents an analysis of the benefits of cross-training between the claims division and calls division in a large health insurance call center by building a discrete-event simulation model of the call center. The simulation model was built in Rockwell Software's Arena v. 7.0 using a modified trace-style simulation. The simulation is directly driven by eleven months of data in half-hour interval summaries; the randomness of the events that occur within these half-hour intervals has been interpreted so as to fit the average speed of answer to that given in the data.

The model of the actual system is then modified to incorporate cross-training. The effect on average speed of answer and claims output of cross-training claims agents that can answer phones when needed is analyzed as well as the effect of cross-training calls agents to process claims when call volumes are low. In addition, the total number of call center agents needed when cross-training is introduced is also examined. The schedules of the call center agents are then manipulated in an attempt to lower the average speed of answer during the busiest periods of the day.

The analysis shows that cross-training holds the potential to not only reduce the average speed of answer and increase the rate at which claims are processed, but can also

significantly reduce the wage-related costs. Also, simply changing the schedules of the call center agents without adding resource capacity had a great positive impact on the average speed of answer, lowering it immensely during the periods which formerly had the highest average speed of answer. This result is an indication that carefully planned scheduling, along with strong schedule adherence, can have a significant positive effect on call center performance.

The Effect of Cross-Training and Scheduling in an Inbound Call Center using Simulation

by
Melanie Beth Racz

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the requirements for the
Degree of Master of Science

Operations Research and Industrial Engineering

Raleigh, North Carolina

2004

APPROVED BY:

Co-chair of Advisory Committee

Co-chair of Advisory Committee

Biography

Melanie Racz is a graduate student at North Carolina State University with co-majors in Operations Research and Industrial Engineering. She will receive her Master of Science degree in Summer 2004. She received her Bachelor of Science in Mathematical Sciences with a concentration in Operations Research from the University of North Carolina at Chapel Hill in May 2002, graduating with Highest Achievement. Ms. Racz is a member of Phi Beta Kappa and Pi Mu Epsilon. She is from Cary, North Carolina where her parents reside.

Acknowledgements

I would like to thank Dr. Stephen Roberts, Professor of Industrial Engineering at North Carolina State University, for his constant support, guidance, and knowledge as well as his unending patience. I would also like to express my deep gratitude to Richard Green, who works in the Workforce Management Group at the company on which this thesis was based, for taking the time to answer all of my questions and provide me with the information I needed to complete this project. Additional thanks are due to George Bourne and Gary Weidhass, also employees in the Workforce Management Group, who provided me with data and answered my questions concerning the data.

Table of Contents

LIST OF TABLES	VI
LIST OF FIGURES	VII
GLOSSARY	IX
1. INTRODUCTION	1
1.1 Problem Description	1
1.2 Objectives	5
1.3 Approach to Analysis	6
2. LITERATURE REVIEW	7
2.1 Erlang C: The Traditional Approach to Call Center Management	7
2.2 Simulation in Call Centers	8
2.3 Scheduling and Resource Management	11
2.4 Cross-Training and Skill Transfer	12
2.5 Input Analysis	13
2.6 Employee Satisfaction	14
3. MODEL DESCRIPTION	16
3.1 Data Collected and Use of Data	17
3.1.1 Preparation of Data	17
3.1.2 Use of Data in Model	18
3.2 Overview of Basic Model	20
3.2.1 Arrival Process	22
3.2.2 Time Period Counter	26
3.2.3 Process of Answering Calls	27
3.2.4 Abandons Arrival and Processing	32
3.2.5 Employee Skill Sets, Schedules, and Failures	33
3.2.6 Statistics Collection	34
3.3 Description of Model with Cross-Training	38

3.4	<i>Description of Model with Scheduling</i>	41
4.	ANALYSIS.....	42
4.1	<i>Verification</i>	43
4.2	<i>Validation</i>	45
4.3	<i>Analysis of Cross-Training Model versus Basic Model</i>	47
4.3.1	<i>Claims Output without Cross-Training</i>	47
4.3.2	<i>Claims Output with Cross-Training</i>	49
4.3.3	<i>JMP Analysis of Cross-Training Effects on Customer ASA and Claims Output</i>	50
4.4	<i>Analysis of Scheduling Model versus Basic Model</i>	70
4.5	<i>Analysis of Cross-Training and Scheduling Together</i>	72
5.	CONCLUSION.....	78
	REFERENCE LIST	81
	APPENDIX	83

List of Tables

Table 1: StateSet 'Agent States'	38
Table 2: New Lunch Schedules	41
Table 3: Intervals of the Day	43
Table 4: Claims Volume Data by Month.....	48
Table 5: Claims Volume Output Totals without Cross-Training.....	48
Table 6: Claims Volume Output Totals with Cross-Training.....	49
Table 7: Taguchi Design.....	51
Table 8: PAN Data used in JMP Analysis	53
Table 9: Use of Parameter Estimates	68
Table 10: Parameter Estimates.....	69
Table 11: Claims Output from Model with Scheduling and Cross-Training.....	72

List of Figures

Figure 1: Assignment of Calls	1
Figure 2: Submodel Flowchart.....	21
Figure 3: Call Processing Flowchart.....	21
Figure 4: File Module	20
Figure 5: Call Arrival Flowchart.....	22
Figure 6: Control Entity Create Module	23
Figure 7: Call Interarrival Delay Module	23
Figure 8: End of Interval Decide Module	25
Figure 9: Record Module	25
Figure 10: Time Period Counter Flowchart.....	26
Figure 11: End of Day Decide Module.....	27
Figure 12: Call Routing Module	27
Figure 13: Queuing of Calls in Seize Agent Module.....	28
Figure 14: Processing of Call Delay Module.....	29
Figure 15: VIM PDF and CDF	30
Figure 16: Delay of Agent for Wrap Module	31
Figure 17: Abandons Arrival and Processing Flowchart.....	32
Figure 18: Abandons Separate Module.....	33
Figure 19: A Service Level Variable Assignment	35
Figure 20: Recording of the ASA Statistic	36
Figure 21: Claims Submodel Flowchart	39
Figure 22: Claims Submodel Flowchart	40
Figure 23: Customer ASA by Scenario	52
Figure 24: Normal Quantile Plots of Responses.....	54
Figure 25: Period 3 ASA Least Squares Regression Model Results	55
Figure 26: Period 3 ASA Least Squares Regression Model Results	56
Figure 27: Period 7 ASA Least Squares Regression Model Results	57
Figure 28: Period 7 ASA Least Squares Regression Model Results	58

Figure 29: Period 10 ASA Least Squares Regression Model Results	59
Figure 30: Period 10 ASA Least Squares Regression Model Results	60
Figure 31: Period 13 ASA Least Squares Regression Model Results	61
Figure 32: Period 13 ASA Least Squares Regression Model Results	62
Figure 33: Period 17 ASA Least Squares Regression Model Results	63
Figure 34: Period 17 ASA Least Squares Regression Model Results	64
Figure 35: Number of Claims Out Least Squares Regression Model Results.....	65
Figure 36: Number of Claims Out Least Squares Regression Model Results.....	66
Figure 37: ASA Results of Scheduling versus No Scheduling.....	71
Figure 38: Customer ASA with Cross-Training One Calls and One Claims Agent Along with Improved Scheduling.....	73
Figure 39: Customer ASA with Cross-Training Two Calls and Two Claims Agents Along with Improved Scheduling.....	74
Figure 40: Customer ASA with Cross-Training Three Calls and Three Claims Agents Along with Improved Scheduling.....	75
Figure 41: Customer ASA with Cross-Training Four Calls and Four Claims Agents Along with Improved Scheduling.....	75
Figure 42: Customer ASA with Cross-Training Five Calls and Five Claims Agents Along with Improved Scheduling.....	76
Figure 43: Customer ASA with Cross-Training Six Calls and Six Claims Agents Along with Improved Scheduling.....	76
Figure 44: Customer ASA with Cross-Training Seven Calls and Seven Claims Agents Along with Improved Scheduling.....	77

Glossary

258: The Provider skill

262: The Employee skill

263: The Broker skill

264: The Provider Benefits& Eligibility skill

266: The Customer skill

ASA: The average speed of answer for all calls answered within a half-hour period of the day

ACD Calls: The number of calls answered within a half-hour period of the day

ACD Time: The average waiting time for all calls in a half-hour period of the day

ACW Time: The average time the answering representative spends in wrap directly after a call completes

Confidence Intervals: All confidence intervals are 95% confidence intervals

Erlang C: A queuing formula developed by A.K. Erlang in 1917; used to determine the number of resources needed in queuing situations using the number of entities waiting to be served, the average service time, and a target for service level or ASA(Cleveland and Julia Mayben 1997).

Lower CI: The lower 95% confidence limits generated from the Arena Category Overview

Split/Skill: A VRU grouping or routing; also referred to by “skill”

SL: Service level; the percent of calls answered within a target time period

Upper CI: The upper 95% confidence limits generated from the Arena Category Overview

Wrap: The process of “wrapping up” or documenting a call immediately after completing the call

1. Introduction

1.1 Problem Description

This thesis will attempt to demonstrate, using a trace simulation, the benefits of cross-functionality and scheduling in a large HMO call center at an anonymous for-profit health insurance agency in the United States. The call center has approximately 70 agents (depending on time period), each of which are assigned to handle different types of calls. Calls are divided into categories (known as split/skills) by a Voice-Response Unit (VRU) and fall into the following groups: Provider, Provider Benefits and Eligibility (B&E), Employee, Broker, and Customer. Employee calls are handled by specially trained agents because of privacy issues. Broker calls are also handled by specially trained agents because brokers are a direct source of revenue for the company. All agents handle Customer calls and almost all handle Provider and Provider B&E. In other words, each agent has a set of skills they are assigned to handle. See Figure 1 for an illustration of the call arrival and handling process.

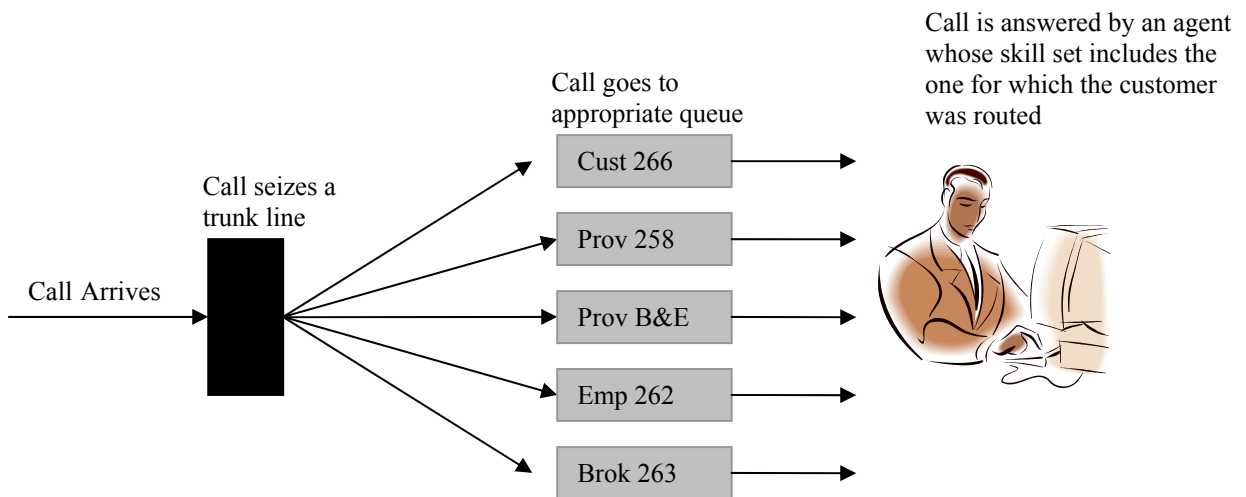


Figure 1 Assignment of Calls

The HMO group has two divisions: the call center division and the claims processing division. The associates that handle claims have comparable salaries and skill levels to those that handle calls. However, none of the claims associates are trained to take calls and none of the calls agents are trained to process claims, i.e. HMO is not cross-functional. There is currently no scheduling tool in place; call center agents are assigned lunch times, but there appears to be little monitoring of agent activity to see that they take their lunches during their scheduled time. The capability to monitor this activity exists through a software package called CentreVu Supervisor, which reports what state associates are currently logged in. Its accuracy is limited because associates are responsible for accurately reporting their state when they are not taking a call. The lunch times assigned are based on what has worked best in the managers' experiences; they are not based in any calculated fashion on forecasted call volumes. As a result, their service level plummets midday and management frequently requests permission to hire more associates to alleviate the midday crisis period.

During a call, an agent is in a CentreVu state called ACD, according to Cleveland and Mayben (1997). The time spent in ACD for each call is recorded and at the end of each 30 minute interval the average ACD Time is recorded along with the number of calls offered and the number of calls answered during the interval. Upon completion of a call, agents document, or wrap-up the call. This process is called "wrap" and agents log into the wrap state in CentreVu. The time spent in wrap is also recorded and averaged at the end of the interval. The model for this paper uses actual call data from 8/5/2002 to 7/25/2003 that reports by interval (i.e. 8:00am-8:30am, 8:30am-9:00am, ... , 6:00pm-6:30pm) for each individual day the average speed of answer (or ASA, the average number of seconds that pass before a call is answered), the number of calls offered, the number of calls processed (or ACD calls), the average ACD time, the average wrap (or ACW) time, the average time until abandon, and the number of abandoned calls. Using a trace-style simulation, it reads each line of data once every 30 simulation minutes and uses the exact data to directly assign values to various attributes and variables that are

used to drive the arrival process, call answering process, abandons, and statistics collection.

At this company, ASA is used as a service/performance measurement. However, Cleveland and Mayben (1997) show that a better measurement is the percent of calls (X) that are answered within Y seconds. This measurement is what will be referred to as service level. For example, an 80/20 service level is 80% of calls answered within 20 seconds. A secondary service objective is to serve the caller's needs in the shortest time, measured in CentreVu by ACD Time.

Because ASA time is calculated from the calls that get answered only and not the calls that abandon, abandons will be completely ignored in the calculation of ASA in the model. The model will, therefore, be verified by looking at ASA data. There is also data available on SL calculated using 30 seconds as the base. This data became available on 5/5/2003, so it will also be used to verify the model. However, the SL data does include abandons when calculating the statistic. To reflect this discrepancy, abandons were included in the model by creating entities called "Abandon" whose only job is to update the service level statistics collection variables. Because ASA can overestimate how long most people wait and SL hides instances when people wait much longer than 30 seconds, it is desirable to integrate the two statistics in the validation process. However, as will be explained in Chapter 4, it was not possible given the time constraints on this project to match the two statistics as closely as had been preferred. It was possible to obtain a fit that closely matched the shape of both graphs, but not the volume.

Operating hours are officially 8:00am-6:00pm. However, the data for the interval 6:00pm-6:30pm shows that several calls are answered after the call center supposedly closes. An HMO manager stated that some people stay until 6:30pm to wrap things up and occasionally calls come in during that time. Even though no agents are scheduled until 6:00pm, the model has been modified to reflect this tendency by scheduling 8 of the 44 agents to stay from 9:30am-6:30pm (an 8.5 hour day instead of the usual 8 hour day).

Agents start work at half-hour intervals from 8:00am to 9:30am and go until 4:30pm to 6:00pm, respectively. In some other call centers agents are assigned two 15-minute breaks, but in this particular call center they take short breaks (restroom, water, personal phone calls, etc.) when needed and their two breaks are added to their lunch time; this way, they can take a full hour lunch instead of a 30-minute lunch, which would be the case with scheduled 15-minute breaks. To account for such small breaks as well as absences, the number of agents was reduced by approximately 20%, similar to what was done in Pichitlamken, et al (2003). The usual amount is 68-70 agents. Using 4 months worth of login data (records of every time an agent logs into or out of the system and the date and the unique agent number) from April 2003 to July 2003, the number of agents found to be absent each day on average during a given month was approximately 13-19%. Because the months used are the warmer, more traveled months, 15% was chosen, totaling a 35% reduction in staff. Hence, the simulation model will use 44 full-time agents instead of 70.

In the data fitting process it was discovered that an equivalent reduction in the later periods of the day resulted in a high ASA. To remedy this, several part-time 'agents' were added; they are active only for the last few periods of the day. In fitting the data for the Employee and Broker skills, several resources were added that were only scheduled for various short parts of the day. In addition, because the resources in the model are much more efficient than real-life resources and adding or subtracting one from the list of available resources has such a large effect on the ASA and SL in the model, 6 of the regularly scheduled (full-time equivalent or FTE) resources were given failures to reduce their availability. Many of the added part-time Employee and Broker resources were given failures as well. Not accounting for lost time due to failures, the total number of FTEs in the model is approximately 48.

In general, the term scheduling includes daily activity such as meetings and paperwork. This thesis will focus on scheduling as applied to lunch assignments. An ideal schedule would be found by looking at call volumes and assigning 60-minute lunch times in

quarter-hour or half-hour intervals according to call needs, as opposed to “best-guess” assignments by management.

The simulation will model the current situation using the data available and evaluate staffing costs and the level of customer service in terms of service level. The model will be modified to include cross-training (cross-functionality) and scheduling, with costs and service levels re-evaluated. The model is built on the Rockwell Software Arena version 7.01 (Copyright © 2002-2003 Rockwell Software, Inc.) platform. Note that the terms ‘interval’ and ‘period’ are used interchangeably throughout this paper.

1.2 Objectives

1. Examine the effects of cross-training between claims and calls:
 - a. What is the cost of cross-training versus the current setup, as measured by service levels and monetarily through staffing costs?
 - b. What is the effect of cross-training on ASA?
 - c. Given the salary increase and training cost of a cross-functional associate, as well as the possible service level increase resulting from cross-training, what is the ideal percentage of staff that should be cross-functional? Is this number different for the claims and calls teams?
 - d. How does cross-training affect the number of associates needed to staff the call center?
2. Determine the effect of scheduling breaks and lunches based on call volumes and of monitoring schedule adherence:
 - a. How is ASA affected?
 - b. Does scheduling change the number of agents required to staff the call center?

1.3 Approach to Analysis

The model is initially built to accurately represent the current situation in the call center. For comparison basis, current ASA and current service levels are recorded. Cross-training is added to the model and these measurements are retaken. The current schedules are modified by changing one resource schedule at a time in an attempt to smooth the ASA curve. Each change is evaluated both separately implemented and simultaneously implemented. Employee satisfaction is evaluated using literature available in similar situations.

Equilibrium cost of cross-training (i.e. running cost after initial training costs) is evaluated using employee salary only since other factors are difficult to quantify. The change in the number of claims processed with cross-training will also be evaluated. The analysis will be performed using the Process Analyzer in Arena with a focus on the Customer ASA as a performance measurement. In evaluating the potential benefits of scheduling, the smoothness of the Customer ASA curve will be examined as well as any change in ASA for each period.

2. Literature Review

As call center technology advances and call centers become more complex, there is a growing need for more accurate workforce management (WFM) systems. The traditional approach to call center management has been Erlang C calculations and experience-based, rather than numbers-based, decision-making. Because, according to Vijay Mehrotra and David Profozich (1997) as well as Brad Cleveland and Julia Mayben (1997), many of the assumptions involved in Erlang C calculations are no longer valid, more attention has been given in the last decade to call center modeling approaches. Queuing and simulation research related to call center environments has given insight into scheduling and resource management as well as input analysis. Information on cross-training in the literature is limited, but there have been several papers outlining the usefulness of simulation in call centers and providing detailed case studies of call center simulations.

2.1 Erlang C: The Traditional Approach to Call Center Management

Brad Cleveland and Julia Mayben discuss the properties of Erlang C that make it a tool that should not be heavily relied upon. The Erlang C, or Erlang delay, model was developed in 1917 by A.K. Erlang (p. 85). It is a queuing formula developed in 1917 used to determine the number of resources needed; its applications span a variety of industries and circumstances, but in the call center environment it is used to determine resource needs to meet given service level or ASA targets. In queuing theory terminology, Erlang C is an M/M/s model (exponential interarrival times, exponential service time, s servers). The number of blocked customers is given by the formula, in which the variable ρ is the server utilization and a is the average number in queue:

$$C(s, a) = \frac{\frac{a^s}{s!(1-\rho)}}{\sum_{k=0}^{s-1} \frac{a^k}{k!} + \frac{a^s}{s!(1-\rho)}}, \text{ where } \rho = \begin{cases} \frac{a}{s} & \text{if } a < s \\ 1 & \text{if } a \geq s \end{cases}$$

According to Cleveland and Mayben, the formula assumes that there are no abandons and that any calls that would have been lost in the real world are actually delayed until they are answered. It also assumes that there is infinite trunking capacity and that no callers will have busy signals; in the case of the call center being studied by this thesis, this assumption holds fairly true with the real system, but it is certainly not true in all call centers.

Because of these assumptions, it is a fairly accurate predictor when service level is good and there are few abandons and busy signals. However, it can overestimate staffing needs during busier periods. In addition, Cleveland and Mayben also state that Erlang C assumes that arrival traffic does not increase or decrease beyond random fluctuation within the time period. It also assumes that there are a fixed number of staff handling calls throughout the time period and that all agents within a group can handle all calls routed to the group. In the call center being studied, not all agents can handle all calls that come in, so Erlang C is not entirely applicable. Assuming that there are a fixed number of agents available also means that Erlang C cannot be used to study the effects of cross-training since cross-training is based on altering staffing levels continuously to efficiently meet demand. In complex call centers, these two assumptions are often not met and so Erlang C can seriously overestimate resource needs.

2.2 *Simulation in Call Centers*

There are several papers in the literature that relate experiences with simulation in call centers. Their purpose is to explain the advantages of simulation, to provide a guide for call center simulation projects, or to give the results of case studies. Mehrotra, et al (1997) describe why simulation works well for call centers, arguing that simulation provides a kind of “laboratory setting” for testing new strategies as well as flexibility of

inputs and outputs. Klungle (1997) adds that skills-based routing has added a large amount of complexity to call centers and has made Erlang C obsolete. Usage of Erlang C models in a complex call center will often result in overstaffing. Because of the ability of simulation to account for all aspects of a call center, Klungle argues that simulation is the “best alternative” in call center modeling. Klungle (1999) goes on to describe his experiences in simulating the AAA Michigan claims call center. His performance measures are service level, average speed of answer, agent utilization, abandonment rate, average length of call, and percent of calls answered without waiting. A shortest processing time rule for call routing was shown to improve queue length, average waiting time, and abandonment rate when compared to the traditional routing approach of grouping by functional type.

Bapat and Pruitte (1998) provide a detailed description of the benefits of simulation in call centers by comparing the advantages and disadvantages to those of traditional call center management approaches. To add to Klungle’s (1997) and Mehrotra’s (1997) complaints about Erlang calculations causing overstaffing, Bapat and Pruitte point out that “studies have...shown that 60% to 70% of the costs in call centers today are associated with staffing and human resources.” As a result, WFM systems that use Erlang to provide scheduling recommendations are limited in effectiveness; the growing industry trend in WFM systems is to include simulation tools. According to Bapat and Pruitte, the issues faced by call centers today that can be completely addressed by simulation techniques include:

- Efficient call handling processes
- Service level
- Skill-based routing
- Simultaneous queuing
- Priority queuing
- Agent preferences and proficiencies
- Agent schedules

Tanir and Booth (1999) agree with these issues and demonstrate the effectiveness of simulation by modeling the complete customer experience in a Bell Canada call center. To limit the granularity of the data, they make the following assumptions:

- The workforce is 50% of the real one since the model does not capture training times, vacation, differences in competency level of agents, and managerial time.
- Financial input is limited to an average hourly wage for agents
- Initial call volumes are based ACD data collected at half-hour intervals

They use a top-down modeling approach by first identifying the high-level processes, then adding the lower-level ones as the model develops. By adjusting the size of the workforce, they showed that a lack of resources at key times in the day can be detrimental to the daily service level. This thesis will attempt to show that this sensitivity can be minimized by proper resource scheduling and cross-training, making more agents available at the busiest parts of the day.

Another important example of call center simulation uses recent research done in the area of input analysis by Avramidis, et al (2003). Pichitlamken, et al (2003) model a call center with inbound and outbound traffic and two types of agents: inbound-only and blend. Their objective is to meet a service level requirement. The challenges they outline are similar to those faced in the model of this paper:

- The format of the data available is in 30 minute interval aggregates
- Arrival rates vary from day to day and within days
- Only the maximum patience time is observed for those who abandon, not for those who are served
- Actual agent availability is difficult to determine due to short breaks, absenteeism, personal calls, etc. They account for this difficulty using a global reduction of the workforce by a fixed percentage (10% to 15%), similar to the approach taken by Tanir and Booth (1999).

They choose an inhomogeneous Poisson process for the arrival rate based on results in Avramidis (2003). Furthermore, the service times were found to be best fit by a gamma distribution, though they found that a lognormal fit was also a good fit.

2.3 Scheduling and Resource Management

Because so much of a call center's cost is in staff needs, much research has been done on agent staffing levels, especially using queuing models. Andrews and Parsons (1993) consider a different method of optimization by focusing on economic considerations rather than meeting a set service level objective. Because their objective was slightly different than the traditional approach, they considered trunk-usage charges as a cost of waiting in addition to the average hourly wage of agents. In contrast, Koole and Sluis (2002) consider a shift scheduling problem where the goal is to meet an overall service level objective as opposed to scheduling to meet a service level at every interval.

Askin and Harker (2001;2003) consider more than just human resources. The dependency between different resources in inbound call centers is investigated in Askin and Harker (2001). A queuing model is developed for both a loss system and a system with reneges that is the "first in the literature to capture the impact of shared information processing resources on phone center performance." The model presented is useful in determining when it is more cost effective to increase capacity through hiring agents or through investing in an information system upgrade. Following the lead of Andrews and Parsons (1993), Askin and Harker (2003) go on to model a system with reneging where optimality is defined as maximizing profits, not meeting a pre-defined service level. The objective is to find the optimal number of servers to be allocated to different call types when there is a common shared limited resource. The resources they consider in addition to human agents are VRUs (Voice Response Units) and information technology resources. Their model shows that ignoring shared resources can have a large impact, likely resulting in overstaffing. For the purposes of this paper, shared resources will not be heavily considered in order to limit the granularity of the data; also, representatives of the company have said there are more than enough trunk lines, so busy signals have not been a problem.

2.4 Cross-Training and Skill Transfer

Cross-training agents carries with it certain considerations for the human factor. The model does not account for the initial training process as it depicts the long-run effect of cross-training. In addition, in determining the length and approach to training, one must examine the nature of the new skills in relation to skills already possessed.

Information in the literature on cross-training in a call center environment is extremely limited, but Leshner and Browne (1993) outline the benefits of cross-training in the insurance industry:

- Increased productivity (average of 20%)
- Shifting of staff to support growth
- Improved customer service
- Increased employee morale
- Increased flexibility
- Provides career development opportunities for employees
- Culture becomes more progressive

Some companies may shy away from cross-training because of the effort involved in the transition. Says Leshner and Browne (1993), “successful organizations find that about 20% of team members’ time during the first year is spent on various training activities.” In addition, it is important to consider skill transfer effects in order to minimize negative transfer. Wickens (1992) states that negative transfer occurs when “the two situations have highly similar stimulus elements but different response or strategic components...” However, “two systems may be considerably different in their display characteristics but can involve positive transfer if there is identity in the response elements.” Given this information about display characteristics, Wickens goes on to say that if situations arise “in which responses between two systems must be different and incompatible...the amount of negative transfer may be reduced by actually *increasing* the display differences.” It is recommended that in cross-training between claims and calls, job functions be broken down into their basic elements to examine where negative transfer may come into play.

Because processing claims and taking calls may have certain skill requirements in common, it is important to examine the possible interaction or lack of interaction of skills necessary to perform each task. In general, according to Swezey and Llaneras (1997), the learning process follows a power function, i.e. “performance improves rapidly early in practice, but soon approaches a state of diminishing returns where each additional increment in performance requires longer and longer practice intervals.” Additionally, “easy tasks often lead to ceiling effects, while excessively difficult tasks often demonstrate an insensitivity to practice.” The difficulty level of learning a new task is affected by “...meaningfulness, concreteness, familiarity, and associations with other information in the memory store.” The difficulty of cross-training will be affected by the level of association, similarity, and interaction between the two tasks. Hence, the training program used to train new associates in either task may not be suitable for associates who are already practiced in the other in that if there is likely no negative transfer, it may need to be abbreviated.

2.5 Input Analysis

Traditionally, arrival processes for incoming call centers have been modeled using a general Poisson process. Recently, however, this standard has been studied more closely and shown to be significantly inaccurate. A queuing model of a small bank was created by Brown, et al. (2002) for the purpose of studying arrival and service time distributions. The group had access to one year’s worth of detailed call-by-call data. Analysis showed that arrivals are best described by an inhomogeneous Poisson process and that service times (talk times) are most accurately modeled using a lognormal distribution. They also discovered that arrivals and service times are positively correlated as a function of the time of day, i.e. the busier it gets, the longer the service times become. A possible explanation is that when a caller has to wait a long time for their call to be answered, they are more likely to initially spend some time complaining about the wait and then may ask more questions so that they feel like the wait was worth their time.

More recently, Avramidis, et al. (2003) executed a similar study with a focus on arrival patterns and intra-day, not day-to-day correlations. They outline the following properties of arrival processes:

- P1: The total number of calls in a day “has overdispersion relative to the Poisson distribution”
- P2: “The arrival rate varies considerably with the time of day”
- P3: “There is a strong positive association...between arrival counts in a time partition of a day”
- P4: “There is a significant dependency between arrival counts on successive days”

They create three models that capture P1 through P3. Model M1 uses a doubly stochastic Poisson process for arrivals. Model M2 has a more flexible covariance matrix than model M1 and model M3 has the greatest flexibility in induced moments. Their case study of a Bell Canada call center using these models shows that “simulation-based call center performance measurement is sensitive to the arrival-process model.”

2.6 *Employee Satisfaction*

In addition to training considerations, cross-training also compels one to bear in mind the effect on employee satisfaction. The model built assumes that there will be no change in the behavior of the employees. Because, say Kalimo, et al (1997), job satisfaction is an important factor in absenteeism and turnover as well as (it can be assumed) productivity, it should be a consideration in deciding whether or not to cross-train agents.

According to Kalimo, et al (1997), “job satisfaction... is characterized as a general positive attitude toward work.” More specifically, “work-related determinants of high job satisfaction include task variety, autonomy, skill development, good interpersonal relations, fair pay, and a pleasant work environment.” Kalimo, et al. assert that job dissatisfaction can be caused by low job content that creates repetition and boredom. This effect can be remedied by providing for the employee increased involvement and control over tasks as well as a greater diversification in tasks. Enhancing employee skills

and career development is one way to decrease psychological distress in associates. Given the importance of job diversification and skill development to job satisfaction, it is safe to argue that the increased variety of skills acquired by cross-trained agents will contribute to their having a positive attitude towards work.

3. Model Description

This chapter will describe in detail the basic model, the model with cross-training, the model with scheduling, and the model with both cross-training and scheduling. Chapter 4 deals with the analysis of these models. The replication length is 236 days, the number of days available in the data for the extended hours. Statistics and the system are initialized between replications and there is no warm-up period.

The general modeling approach used is a trace-style simulation with some modifications. Using a trace simulation as a basis was chosen to attempt to model what actually happened during the year from which the data was taken as closely as possible. Because the arrivals are not a regular Poisson Process and the service times are not exponential, a technique that would allow for the specification of such details was needed. Because of the nature of the data, an exact trace simulation could not be used. The data being given in half-hour interval summaries requires interpretation of what happened in the system on a lower level during each of those half-hour periods. What is not known from the data is the exact arrival time of each call, the exact wait time of each call, the talk time of the calls, and the time the agent spent in wrap. Instead, what is known is the sum of the calls answered in a period and the average waiting, talk, and wrap times. Therefore the simulation is “trace-like” in that it uses the data given directly, but it is not completely a trace simulation since variability and randomness are introduced in the arrival process, talk times, and wrap times. This randomness also necessitates replications to reduce the standard deviation of the performance measures- primarily ASA but also SL.

Because the data is given in half-hour periods, the model is largely based on half-hour cycles. Control entities are created every 30 minutes to read the next data point in the arrival processes (both regular call arrivals and abandons); they are also used to read information about time and day of week once every 30 minutes. These control entities gather information from the data about the number of calls that should arrive in the next 30 minutes, the average talk time for the calls, and the average time agents should spend

in wrap. They also read the current period of the day (1=8:00am-8:30am, 2=9:00am-9:30am, etc.) and the current day (1=Monday, 2=Tuesday, etc.). The value of each data point gathered is stored in an attribute or a variable for later use. In the case of the regular call arrival processes, the control entity loops for the rest of the 30 minutes by passing through a delay block and a separate block to simulate the interarrival times and arrivals of calls, respectively. At the end of the period, they are sent through some statistics collection modules and are disposed. The control entity used to gather information about day and time simply loops continuously for the duration of the simulation.

3.1 Data Collected and Use of Data

The data collected has been briefly outlined in the introduction. Before using the data in the modeling or validation, much of it had to be gleaned and manipulated. This section will discuss the preparation of the data as well as how the data is used as input in the model.

3.1.1 Preparation of Data

Apart from the five described earlier, the original data came with several Split/Skills that were not highly used (in fact most had no call volume) and one Split/Skill that was used enough that it needed to be included in the data. An employee at the company recommended that it be included in the Employee group. To combine the skills, the ACD Calls and Abandoned Calls data points were summed for each separate interval. The ASA was multiplied by the respective ACD Calls datum; the two quantities were added and the sum was divided by the sum of the ACD Calls for both skills to obtain an overall ASA. A similar approach was taken for Average Abandon Time, ACD Time, and ACW Time. The skills that were insignificant were deleted from the data.

Number assignments were given to each half-hour interval in a new column so that the data would be easier to filter and use in the model. The assignments were made so that

8:00am-8:30am corresponded to period 1, 8:30am-9:00am to period 2, etc. Using the data functions in Microsoft Excel, the data was interpreted as the day of the week and also assigned a number in a new column. Monday corresponded to day 1, Tuesday to day 2, etc. The data was then divided by skill so that each skill had a separate spreadsheet. The day and interval did not always align for all skills; occasionally intervals would be missing and sometimes there were records for intervals prior to 8:00am or past 6:30pm. When records needed to be added for missing intervals, it was assumed that call volume was zero. Any intervals not between 8:00am and 6:30pm were deleted from the data. The data was cleaned up so that the date and period lined up for all five skills. The file that the data is read from in the model was created from this gleaned and organized database.

The login data used as a starting point for agent skill assignments fitting was analyzed by querying the data for a short time interval (one month) and for unique login identification numbers. Then the number of observations of unique login IDs was examined; if an ID was only observed a few times (as opposed to several dozen times), it was likely to be a manager logging in to make changes in the system (according to an employee at the company) and was disregarded. After removing these ‘manager’ ID records from the data, the login data was also used to determine absence rates; this analysis was performed in a similar fashion using Microsoft Access to query the data by day to get the total number of IDs that logged in each day.

3.1.2 Use of Data in Model

As described in the introduction, much of the data was used to find a starting place for the data fitting process. The schedule that was obtained from the HMO manager was used directly initially but had to be modified during the verification process to fit the data. Also, the login data was used to find a good starting place for the distribution of agent skills but was, again, modified during the verification process to fit the data. The rest of the data was used directly either in verification (ASA and SL averages) or in the simulation itself.

The model reads data from the following categories:

- ACD Calls (number of calls answered per period)
- Average ACD Time (average talk time)
- Average ACW Time (average time spent in Wrap)
- Abandoned Calls
- Average Abandon Time (average time until call abandons)
- Day
- Interval

It utilizes the File module and ReadWrite module in Arena. The data is read from a Microsoft Excel spreadsheet with named ranges. The File module is shown below in Figure 4. The file rewinds when it reaches the end, which occurs at the end of each replication. The Recordsets are the named ranges in the Excel file. Each ReadWrite module assigns an attribute to the entity with a value of the data point read.

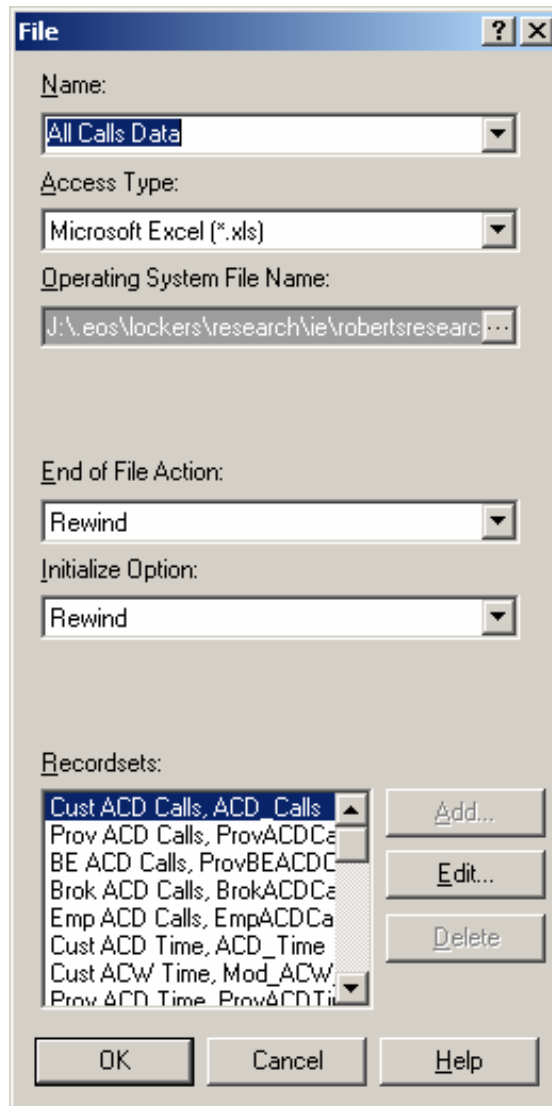


Figure 2 File Module

3.2 *Overview of Basic Model*

In this section, screen shots of the actual Arena flowchart model will be referred to. The underlying details will be explained more thoroughly in the sections to follow. The flowchart of the basic model has three major components: the arrival process submodels, the calls processing submodel, and the time period counter submodel (Figure 3). Because the data is separated by skill, the call arrival process and call handling (Figure 4) process have been split up into five groups, one for each skill.

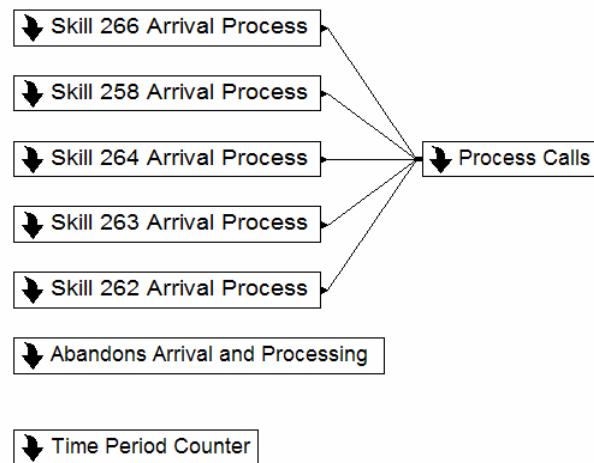


Figure 3 Submodel Flowchart

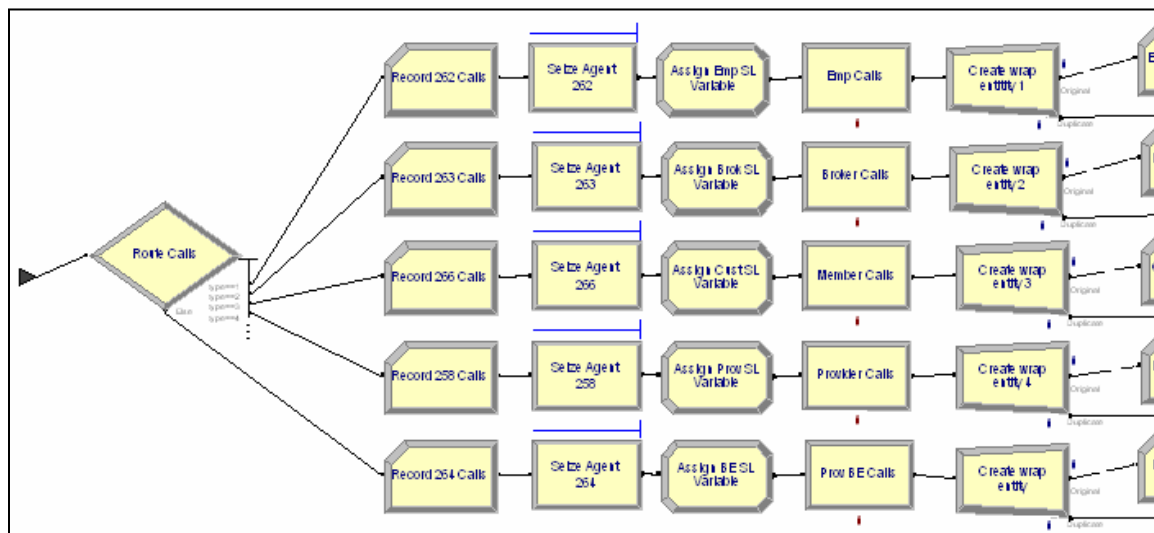


Figure 4 Call Processing Flowchart

3.2.1 Arrival Process

Since each arrival process is the same for all skills, this description and any screenshots provided will be from the Customer 266 split/skill. Below (Figure 5) is a view of the arrival process flowchart. It is too long to fit in one screenshot; the right half is below the left half.

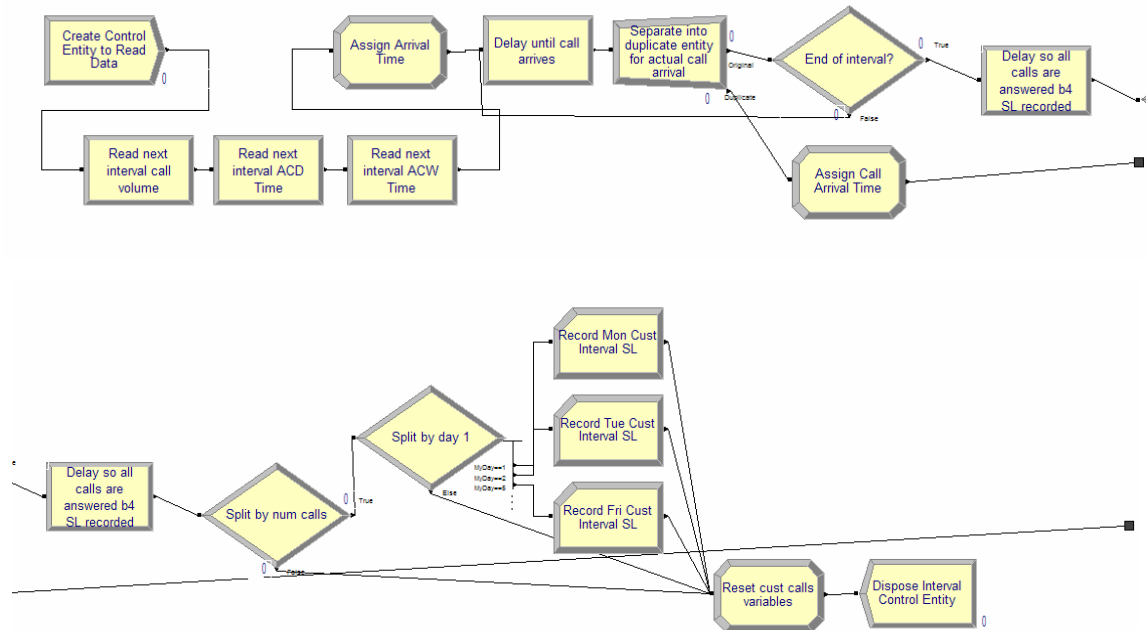
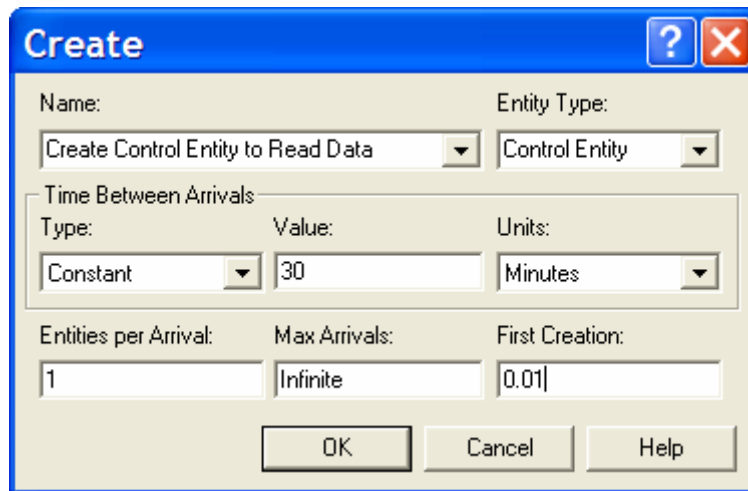


Figure 5 Call Arrival Flowchart

The arrival process uses a control entity (Figure 6) that is created every 30 minutes to read the ACD calls, ACD time, and ACW time data. The first creation is at time 0.01 so that the creation of the Time Period Counter entity does not occur simultaneously with the control entity. The Control Entity is assigned the attributes ‘ACD calls’, ‘ACD time’, and ‘ACW time’ through the three Read modules. It then goes to an Assign module, ‘Assign Arrival Time’, where it is assigned the attribute ‘Arrival Time’ with a value of TNOW (the current simulation time) and the attribute ‘type’ with a value of 1, 2, 3, 4, or 5 depending on the split/skill. Customer has a value of 3. Employee, Broker, Provider, and Provider B&E have values 1, 2, 4, and 5, respectively. It is also assigned the attribute ‘MyInterval’ with a value of Period.

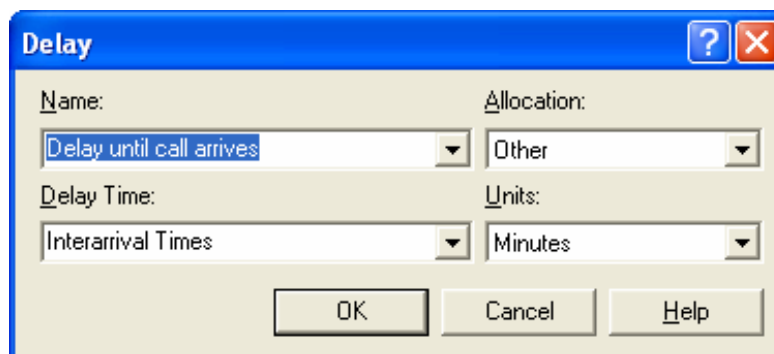


The 'Create' dialog box is used to configure a new control entity. It features a blue title bar with a question mark and a close button. The main area is divided into several sections: 'Name' and 'Entity Type' at the top, 'Time Between Arrivals' in the middle, and 'Entities per Arrival', 'Max Arrivals', and 'First Creation' at the bottom. The 'Name' field is set to 'Create Control Entity to Read Data' and 'Entity Type' is set to 'Control Entity'. The 'Time Between Arrivals' section has 'Type' set to 'Constant', 'Value' set to '30', and 'Units' set to 'Minutes'. The 'Entities per Arrival' is set to '1', 'Max Arrivals' is set to 'Infinite', and 'First Creation' is set to '0.01'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Name:		Entity Type:	
Create Control Entity to Read Data		Control Entity	
Time Between Arrivals			
Type:	Value:	Units:	
Constant	30	Minutes	
Entities per Arrival:	Max Arrivals:	First Creation:	
1	Infinite	0.01	
OK		Cancel	Help

Figure 6 Control Entity Create Module

The 'Arrival Time' attribute is used to keep track of the time that the current half-hour interval began. The MyInterval attribute is used for statistics collection to refer to the interval during which the entity actually arrived; the current period may be a later period than the arrival period. The Control Entity then enters a delay block where it is delayed for a period of time determined by the Expression 'Interarrival Times' (Figure 7).



The 'Delay' dialog box is used to configure a delay block. It features a blue title bar with a question mark and a close button. The main area is divided into several sections: 'Name' and 'Allocation' at the top, 'Delay Time' and 'Units' in the middle, and 'OK', 'Cancel', and 'Help' buttons at the bottom. The 'Name' field is set to 'Delay until call arrives' and 'Allocation' is set to 'Other'. The 'Delay Time' is set to 'Interarrival Times' and 'Units' is set to 'Minutes'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

Name:		Allocation:	
Delay until call arrives		Other	
Delay Time:		Units:	
Interarrival Times		Minutes	
OK		Cancel	Help

Figure 7 Call Interarrival Delay Module

Interarrival Times is a row of the Expression spreadsheet module and is the following expression:

$$(\text{ACD Calls} = 0) * 30 + (\text{ACD Calls} > 0) * \text{EXP}((30 / (\text{ACD Calls} + 0.000001))).$$

If the value of the attribute 'ACD Calls' is equal to zero, the first set of parentheses will return a value of 1; hence, the entity will be delayed 30 minutes, the full interval length. Otherwise, the first set of parentheses will return a value of 0 and the second set of parentheses will return a value of 1. In that case, the entity will be delayed an exponential amount of time with mean $30 / (\text{ACD Calls})$. Adding a small number is necessary for the compiler not to read a divide-by-zero when 'ACD Calls' is zero. Delaying a simple exponential length of time is not precisely correct due to the fact that some entities that should arrive in the interval in which the parent Control Entity was created will arrive in the following interval with a delay length based on the previous intervals call volume. However this approximation does not appear to distort the representation of the actual arrivals.

After delaying the length of the interarrival time, the Control Entity is sent to a Separate module (named 'Separate into duplicate entity for actual call arrival'). The duplicate entity can be thought of as the actual call that has arrived. It is the entity that will go on to the 'Process Calls' submodel. This entity has inherited all of the attributes of the parent Control Entity. It moves on to an Assign module where it is assigned its own attribute 'arrivalTime' with a value of TNOW. This attribute is used to reference the arrival of the actual call. The call entity is also assigned the attribute 'MyInterval' with a value of 'Period' since its arrival period may be different from its parent. Similarly, it is assigned the attribute 'MyDay' with a value of the variable Current Day. A couple of variables are also assigned here for statistics collection purposes; these assignments will be described later in Section 3.2.6.

After the Separate block, the Control Entity enters a Decide module (Figure 8). If the current time interval has ended, the entity is sent to a short Delay block. There is a dual purpose for the delay: to prevent simultaneous events, namely the assignment of the 'cust

calls per period' variable in both the 'Assign Call Arrival Time' and 'Reset cust calls variable' Assign modules; and to delay the recording of the interval SL until any entities that arrived in the interval being recorded have either completed their queuing or have waited longer than 30 seconds. Calls are then sent to two Decide modules. The first ('Split by num calls') is for statistics collection purposes and send calls to the next Decide module and hence to a Record Module only if the number of calls in the interval that just ended is greater than zero. The Record module (Figure 9) collects interval service level statistics by day (the second Decide module routes to the correct Record module based on the value of 'MyDay') and the Assign module resets two variables used to collect SL statistics back to zero. This statistics collection process will be described in more detail in Section 3.2.6.

The 'Decide' module configuration window shows the following settings:

- Name:** End of interval?
- Type:** 2-way by Condition
- If:** Expression
- Value:** (TNO'W - Arrival Time) >= 30

Figure 8 End of Interval Decide Module

The 'Record' module configuration window shows the following settings:

- Name:** Record Mon Cust Interval SL
- Type:** Expression
- Value:** cust calls answered in 30s(MyInterval)
- Record into Set:** ☒
- Tally Set Name:** Mon Cust Interval SL
- Set Index:** MyInterval

Figure 9 Record Module

The original Control Entity moves back to the ‘Delay until call arrives’ module after leaving the Separate module and repeats the process. Once the duplicate entity leaves the ‘Assign Call Arrival Time’ module, it completes the arrival process and moves on to the ‘Process Calls’ submodel.

3.2.2 Time Period Counter

The Time Period Counter (Figure 10) section of the model is completely separate from the rest of the model. The purpose is to keep track of the current day (1 = Monday, 2 = Tuesday, etc.) and the current time interval (1 = 8:00am-8:30am, 2 = 8:30am-9:00am, etc.). The counter entity is created only once per replication at the start of the run.

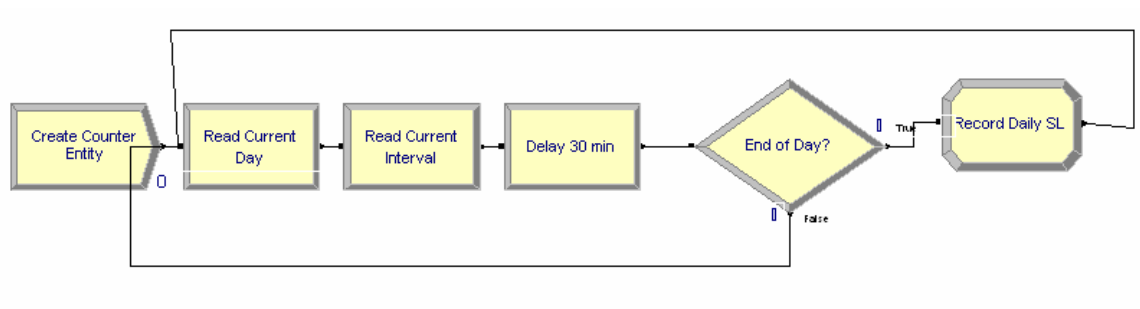


Figure 10 Time Period Counter Flowchart

When the counter entity passes through the ‘Read Current Day’ module, it assigns a variable ‘Current Day’ the value of the next input from the corresponding data (described in detail in Section 3.2.3). Passing through the ‘Read Current Interval’ module, it assigns a variable ‘Period’ the value of the next input from data. The entity is then delayed 30 minutes and sent to a Decide module (Figure 11). If it is the end of the day, i.e. if the current period is 21, then it is sent to a record module for daily statistics collection on SL. After the Record module it goes back to the ‘Read Current Day’ module to repeat the process until the end of the replication. If it is not the end of the day, it bypasses the Record module to go to the ‘Read Current Day’ module.

Decide

Name: Type:

If: Named: Is:

Value:

OK Cancel Help

Figure 11 End of Day Decide Module

3.2.3 Process of Answering Calls

When calls arrive in the Process Calls submodel, they are immediately routed by call type through a Decide module (Figure 12). So that queue times can be recorded for both ASA statistics and SL statistics, the seize and delay processes are separated. Because of the wrap-up process, the release of the resource is also separate.

Decide

Name: Type:

Conditions:

- Attribute, type, ==, 1
- Attribute, type, ==, 2
- Attribute, type, ==, 3
- Attribute, type, ==, 4
- <End of list>

Add... Edit... Delete

OK Cancel Help

Figure 12 Call Routing Module

After the decide module, the call enters the Seize block (Figure 13) where it will wait for an available agent. It seizes a resource from the appropriate resource set (in this case the set 'Cust 266', which is the set of agents that can answer customer calls) and gives the resource an attribute to designate which resource from the set has been used so it can release the correct resource later on. The resource selection rule is not tremendously important since there is no difference in skill level in the model between agents, but Preferred Order was chosen so that the part time agents would be used the least and so that in the model with cross-training the cross-trained claims agents would be used last. At the company, calls are routed to the agent that has been available the longest. It is possible to model this selection criteria, but not practical since agents are known to manipulate the system by going into non-call busy modes to shorten their availability time.

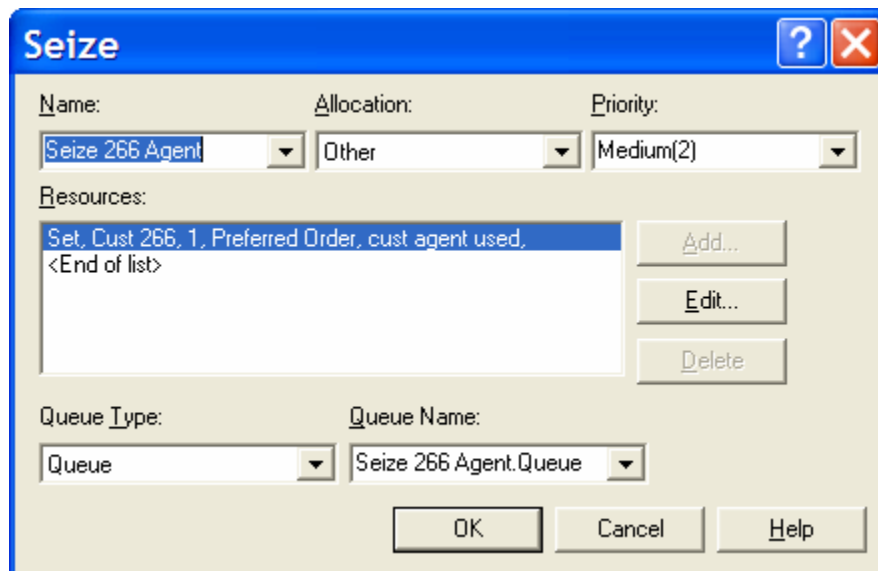


Figure 13 Queuing of Calls in Seize Agent Module

After seizing a resource, the call entity passes through a record block and an assign block for ASA and SL statistics collection. It then enters the delay block (Figure 14) that simulates the answering and completion of the actual call. The delay time is the expression 'ACD Times', which is the following expression, defined in the Expression spreadsheet module:

JOHN(1.354,.887,.58*(ACD Time),.87*(ACD Time),10)

The expression was found by assuming that talk times can be approximated by a JohnsonSB distribution (since no actual call time data was available- only averages over 30-minute intervals- an assumption as to the shape of the distribution had to be made). In the expression, 'ACD Time' refers to the attribute assigned to the Control Entity that gives the current interval's average talk time. The last parameter in the distribution is the random number stream used. To derive the Johnson distribution parameters, basic assumptions were made about the minimum, maximum, and mode. The program VIM (Version 2.0, Stephen Roberts, Lijun Wang) was used to determine the shape parameters that would result from this minimum, maximum, and mode combination and to graph the CDF and PDF (Figure 15). After several runs of the simulation (while using VIM to modify the mode or median) to improve the fit from this starting point, the minimum, maximum, and mode were modified to produce the above distribution.

The screenshot shows a 'Process' dialog box with the following fields and settings:

- Name:** Member Calls
- Type:** Standard
- Logic:**
 - Action:** Delay
- Delay Type:** Expression
- Units:** Minutes
- Allocation:** Value Added
- Expression:** ACD Times
- ☒ Report Statistics
- Buttons:** OK, Cancel, Help

Figure 14 Processing of Call Delay Module

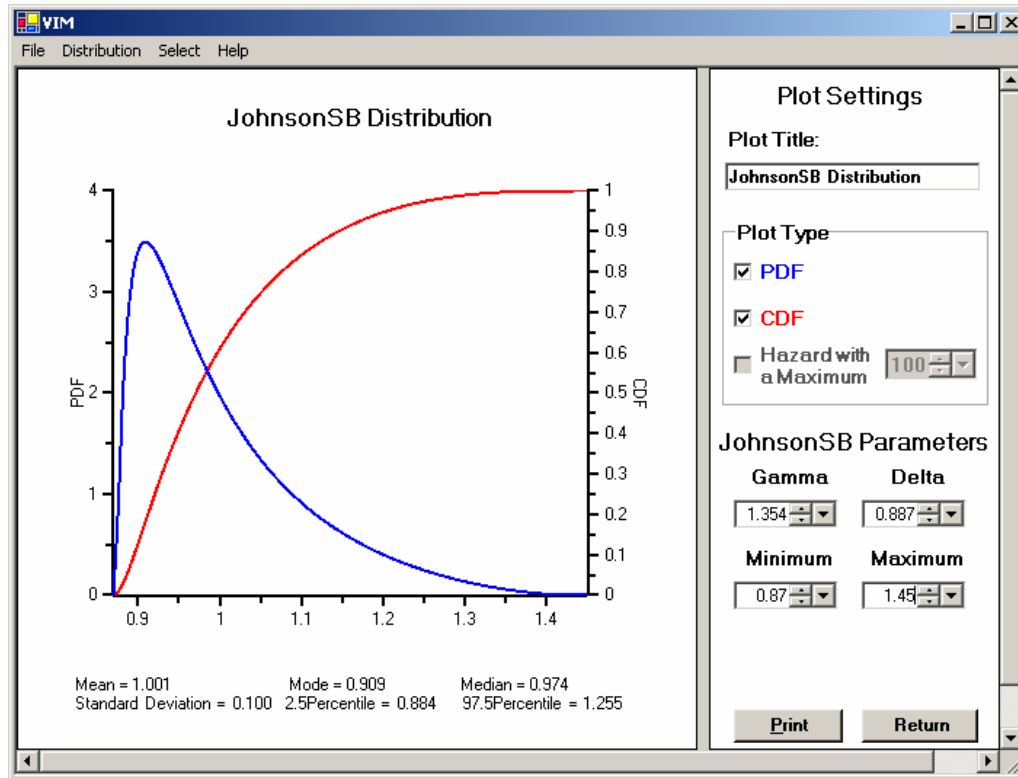


Figure 15 VIM PDF and CDF

Once the delay is complete and the phone call has ended, the call entity goes to a separate block where a single duplicate is created. It is duplicated so that the simulation does not model the call as being delayed again while the agent wraps up the phone call after it has ended. In this way, the simulation can model the agent being busy after the customer has hung up and left the system. Once the call has gone through the separate block it goes to a record module to collect cycle time statistics and is disposed. The duplicate entity goes to an assign module. Here it is assigned the attribute 'wrap start time' so that statistics can be collected on how long the agent is in 'wrap'. It also assigns the seized resource to the state 'Wrap', the purpose and details of which will be described more in section 3.2.6. It then enters a process block (Figure 16) that delays the entity and releases the resource. Here in choosing a resource to release, the attribute that was saved in the seize block is used to identify the correct resource used from the appropriate set. The delay time for the wrap-up process is again an Expression, 'ACW Times'. It is derived in the same manner

as 'ACD Times', the only difference being that it uses the attribute 'ACD Time' assigned in the arrival process submodel. Below is the Expression:

JOHN(1.354,.887,.58*(ACW Time),.87*(ACW Time),12)

Upon completion of the wrap process, the duplicate entity moves to a record block where the wrap delay time is recorded and is then disposed. This completes the flowchart of calls through the system.

The screenshot shows a 'Process' dialog box with the following fields and options:

- Name:** wrap 3
- Type:** Standard
- Logic:**
 - Action:** Delay Release
 - Resources:**
 - Set, Cust 266, 1, Specific Member, cust agent used
 - <End of list>
- Delay Type:** Expression
- Units:** Minutes
- Allocation:** Value Added
- Expression:** ACW Times
- ☒ Report Statistics

Buttons at the bottom: OK, Cancel, Help.

Figure 16 Delay of Agent for Wrap Module

3.2.4 Abandons Arrival and Processing

As with Section 3.2.3, any screenshots in this section will be taken from the Customer 266 Split/Skill. The abandons arrival and processing is the same for all Split/Skills. Below (Figure 17) is a view of the entire abandons arrival and processing flowchart.




Figure 17 Abandons Arrival and Processing Flowchart

The Abandon entity is similar to the Control Entity in the call arrival process. It is created once every 30 minutes and immediately passes through two Read modules. The Read modules assign the Abandon entity the attributes 'Num Aban' and 'Aban Time'. It then passes through an Assign module where it is assigned the attribute 'MyInterval' with a value of 'Period'. From there it enters a Separate module where it creates a number of duplicates equal to the attribute 'Num Aban' (Figure 18). The Abandon control entity is the disposed. The duplicates go to an Assign module called 'Increment Cust Calls' where it increments a variable used to calculate SL. It passes next to an Assign module called 'Assign Cust SL Variable Abandon' that updates another SL variable and is then disposed. This process repeats once every interval during the simulation.

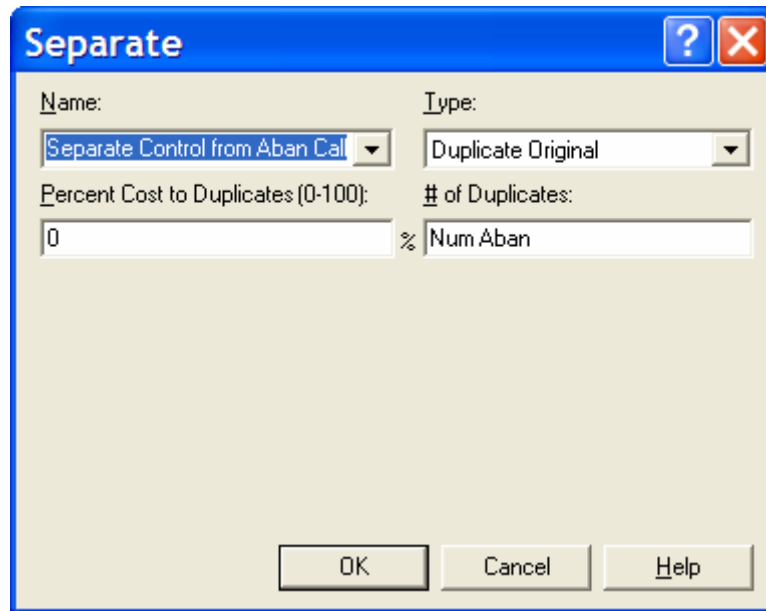


Figure 18 Abandons Separate Module

3.2.5 Employee Skill Sets, Schedules, and Failures

There is one resource for every agent and one schedule as well. For example, agent Jane Doe would be a single resource of maximum capacity one and corresponding individualized schedule ‘Jane Doe sched’. The schedule rule used for all resources is Ignore, meaning that when an agent is scheduled to go on break or off shift, he or she finishes the call they are currently processing before starting the change in the schedule. In addition, the time during which the capacity is zero starts when it is scheduled to start rather than when the resource actually begins the schedule change, as described in Kelton, et al (2002). This way the schedule is not shifted. For instance, if an agent works from 8:00am-4:30pm and finishes his last call at 4:35pm, he will start again at 8:00am and not 8:05am. The cost per hour is \$20.51, which is based on a loaded salary of \$40000, with 52 weeks per year and 7.5 hours per day of having a resource capacity equal to one. Each agent has StateSet ‘Agent States’, described in Section 3.2.6. Six of the full-time agents have one failure. This failure can only occur during the resource state Available (referring to “Idle”). The Uptime for the failure is 0.3 minutes and the

downtime ranges from 0 to 30 minutes over the periods. A list of the resource names and schedules used is provided in the appendix in Table A1. Some of the schedules are not the typical 8.5 hour day with a one-hour lunch. It was necessary to modify a couple of the schedules to better fit the data. Below Table A1 is a table of the resource skill assignments; this is in Table A2. The resources in bold have failures assigned to them.

There are five employee skills sets, one for each Split/Skill. They are ‘Cust 266’, ‘Prov 258’, ‘Prov BE 264’, ‘Brok 263’, and ‘Emp 262’. Each of these sets is of the type Resource set. To determine the composition of each of these sets, one month worth of agent login data was analyzed to determine what percentage of all agents logged in to each skill. It was found that each agent that logged into the system had logged into the Customer 266 skill. Approximately 96% of agents logged into the Provider 258 skill, 94% into the Provider B&E 264 skill, 37% into the Broker 263 skill, and 26% into the Employee 262 skill. However, during the validation process it was found that this distribution needed to be modified to fit the data. In the process, some part-time resources were added to the employee and broker skill sets only. The new percentages are:

- Customer: 94%
- Provider: 65%
- Provider B&E: 78%
- Broker: 48%
- Employee: 41%

3.2.6 Statistics Collection

There are two main types of statistics collected in this simulation: the ASA by interval and the SL by interval and day (e.g. Monday Period 1 SL). Each of these types of statistics is collected for each of the five skills. Other minor statistics are collected to assist in the model verification and validation process. These additional statistics include the number of calls by period for each skill, caller cycle time, and agent time spent in wrap. Resource states were also assigned using the Stateset ‘Agent States’ for

verification purposes only. Outside of resource states, statistics will be described in the order in which they appear in the model.

Service level statistics are the most complicated statistics collected and the process begins in the arrival submodel. After the Control Entity separates in the Separate module called 'Separate into duplicate entity for actual call arrival', it enters an Assign module where it updates a variable used in the SL statistic (Figure 19). In the case of the Customer 266 skill the variable is called 'cust calls per pd' and is a 21-dimensional array variable. This assignment increments by one the index of the variable corresponding to the period in which the call arrived, represented by the attribute 'MyInterval'.




Figure 19 A Service Level Variable Assignment

The 'Split by day' Decide module routes calls to Record modules only for Mondays, Tuesdays, and Fridays. These three days were chosen because, according to Cleveland and Mayben, Monday tends to be day of highest call volume in most inbound call centers, Friday tends to be the slowest day, and Tuesdays are generally similar to Wednesdays and Thursdays. As shown in Figure 20, the Record module logs the percent of calls answered within thirty seconds for the previous interval as given by the expression:

'cust calls answered in 30s(MyInterval)/cust calls per pd(MyInterval)'.

Figure 20 Recording of the ASA Statistic

It has already been described how the denominator of the above expression is arrived at. The numerator is a similar 21-dimensional variable that is updated in the ‘Process Calls’ submodel just after the caller seizes an agent. In the ‘Assign Cust SL Variable’ module, the index corresponding to ‘MyInterval’ of the variable ‘cust calls answered in 30s’ is updated to the following value:

$$((\text{tnow} - \text{arrivalTime}) \leq .5) * (\text{cust calls answered in 30s}(\text{MyInterval}) + 1) + ((\text{tnow} - \text{arrivalTime}) > 0.5) * (\text{cust calls answered in 30s}(\text{MyInterval}))$$

In words, the above expression will increment the variable by one if the caller has waited no more than 0.5 minutes; otherwise it will remain the value it is currently. After passing through the Record module, the call goes to an Assign module where these two variables are reset to 0 in preparation for the next day’s SL computation.

There is a tally statistic for each combination of skill, period, and day (Monday, Tuesday, or Friday). Sets are used to record into the tally statistic corresponding to the interval being recorded. For example, ‘Mon Cust Interval SL’ is a tally set composed of the 21 tally statistics representing the Monday Customer SL for each of the 21 intervals. Therefore, in Figure 20 ‘MyInterval’ provides the set index of ‘Mon Cust Interval SL’ in which to record the computed interval SL.

The SL variables described above are also updated in the abandons submodel. The ‘Increment Cust Calls’ Assign module increments by one the ‘cust calls per pd’

variable. The 'cust calls answered in 30s' variable is updated in a similar fashion as in the arrivals submodel. In this case the expression is:

$$((\text{Aban Time}) \leq .5) * (\text{cust calls answered in 30s}(\text{MyInterval}) + 1) + ((\text{Aban Time}) > 0.5) * (\text{cust calls answered in 30s}(\text{MyInterval}))$$

The result is the same as above, the only difference is that the entity is previously assigned a time to wait until abandoning. For simplicity's sake, this time is equal to average time to abandon from the data. If this time is not greater than thirty seconds, 'cust calls answered in 30s' is incremented by one.

The next statistic collected is the number of calls by skill. After being routed by skill type, the calls pass through a Record module called 'Record 266 calls' which counts by one and records into a Counter set. There are five sets, one for each skill, and they are composed of 21 Counter statistics corresponding to each period.

ASA is recorded just after the resource is seized by a call entity. Like SL, it is recorded into a tally set composed of 21 ASA tally statistics, each representing a period for that particular skill. For the Customer skill, the tallies are named '266 Period 1 ASA', '266 Period 2 ASA', etc. and the tally set is called '266 ASA'. The Customer Record module (Figure 20) uses 'MyInterval' as the set index and records into the tally set '266 ASA'. The value it records is the length of the time interval from 'arrivalTime' to TNOW, the current simulation time.

The next statistic recorded is the cycle time of the call. It records the value of the time interval from 'arrivalTime' to TNOW into the tally set 'Cycle Time' with a set index equal to the value of the attribute 'type'. The tally set is composed of five members: 'Emp Cycle Time', 'Brok Cycle Time', etc. After the duplicate entity created in 'Create wrap entity' leaves the wrap Process module, it enters a Record module that records the time interval from 'wrap start time' to TNOW in a tally statistic. In the case of skill 266, it is simply called 'Cust wrap time'.

Finally, the StateSet spreadsheet module is used to create a set of resource states agents enter into. The purpose is to be able to view Frequency statistics to see that the time spent in each state or failure is what is expected and seems to fit the real world proportions. The four states and their corresponding system state, when applicable, or failure name are given in Table 1.

Available	IDLE
ACD	BUSY
Wrap	
Break	Break 1

Table 1 StateSet 'Agent States'

3.3 Description of Model with Cross-Training

The model that includes cross-training has only a couple of key differences from the basic model. First, it has a new resource(s) whose primary function is to process claims. In addition, one or more of the current calls resources become cross-trained. This modification is made using the resource Sets. A submodel called 'Claims' has been added to process the claims. A flowchart view of the submodel is shown in Figure 21. There is a new set called 'Claims' that contains all of the cross-trained agents in order of claims agents then calls agents. The original resource sets for the call Split/Skills have been modified to contain the new claims agents at the end of the set (in this way, the Preferred Order method of resource seizing will choose the claims agents last, as it is done in the real-life situation).




Figure 21 Claims Submodel Flowchart

Claims are created once a day with a quantity of 1000. This creation process was chosen so that the number of entities in the system did not grow out of control but there were always claims waiting to be processed. It has been assumed that there is always claims work to be done. The claims Process module (Figure 22) uses a Priority of ‘Low’ for claims. This assignment assures that when call volumes are high enough to warrant the assistance of the claims representatives the claims are put aside until call handling improves. Additionally, the calls agents will only work on claims when call volumes are low enough that they are temporarily not needed. A triangular distribution for claims processing times was chosen for its simplicity. It is not known how long on average it takes to process a claim in the HMO department. Some of the employees guess that it is about six minutes. Since this thesis is a study in comparison, an accurate distribution is not needed. These are the only changes that were needed for the claims model. The analysis was performed by changing the number of claims agents and cross-trained calls agents available.

Process

Name: Type:

Logic

Action: Priority:

Resources:

Delay Type: Units: Allocation:

Minimum: Value (Most Likely): Maximum:

☒ Report Statistics

Figure 22 Claims Submodel Flowchart

The analysis was performed with the cross-training of between one and seven of both calls and claims agents. The following calls agents were added to the claims resource Set in the following order:

1. Resource 24
2. Resource 22
3. Resource 16
4. Resource 34
5. Resource 17
6. Resource 23
7. Resource 27

They were chosen for having regular schedules, no failures, a variety of lunch times, and for being trained in only skills 266, 258, and 264. However, it was necessary to choose three that were also trained in the Broker skill, 263. The schedules of the claims resources were as follows and were added in the following order:

1. 8:00am to 4:30pm with lunch at 11:00am
2. 8:30am to 5:00pm with lunch at 12:00pm
3. 9:00am to 5:30pm with lunch at 1:00pm
4. 9:30am to 6:00pm with lunch at 2:00pm
5. 8:00am to 4:30pm with lunch at 11:00am
6. 8:30am to 5:00pm with lunch at 12:00pm
7. 9:00am to 5:30pm with lunch at 1:00pm

These schedules were chosen to distribute as evenly as possible the shifts and lunch breaks, with a focus on the earlier shifts (in the actual environment there seems to be a preference for the earlier shifts).

3.4 Description of Model with Scheduling

The scheduling model is only different in its distribution of schedules. This distribution was found by starting with the basic HMO model and altering schedules until the ASA graph was reasonably smooth. The goal was to minimize the maximum average waiting time without changing the availability of resources (i.e. without adding any new resources or lengthening the current schedules of resources). Table 2 below contains the list of employees whose lunch times were changed and the time that it was changed to. The results of this change are shown in Section 4.4 where they can be compared to the basic model.

Resource 3	1:30
Resource 13	1:30
Resource 22	2:30
Resource 25	1:30
Resource 29	12:30
Resource 16	11:00
Resource 24	11:00
Resource 27	11:00

Table 2 New Lunch Schedules

4. Analysis

This chapter will describe the process through which the workings of the model were verified to accurately represent the real-working system as well as how the model was validated to fit the data obtained from the actual system. Various debugging capabilities in Arena were utilized to verify the model. The validation was done by changing various statistical distributions in the model, resource schedule, and resource set compositions and comparing the results to the ASA and SL from the data.

Of central importance, this chapter will detail the results of the comparison study of cross-training and scheduling versus the current system in place. The interval ASA for the Customer skill is used as a basis of comparison. For the claims model, the difference in claims output is also examined. As described in the beginning of Chapter 3, the basis of the analysis is that the simulation is a trace simulation in that several inputs are read directly from the data, but that there is variability that creates a need for replications. This variability comes from the exponential arrival process and the JohnsonSB distribution of the talk times and wrap times (whose average values are read in from the data). Because of the variability, confidence intervals will be used to compare means between the Customer ASA results of the basic model and the Customer ASA results of the experimental model. To further examine the statistical difference between the basic Customer ASA and the ASA of the experimental model, JMP v. 5.1.1 (A BUSINESS UNIT OF SAS Copyright © 1989 - 2004 SAS Institute Inc.) was used to design an experiment to analyze the effect of cross-training on claims output and ASA.

While a couple of ASA graphs are provided within the body of the text, the majority of the ASA and SL graphs are provided in the Appendix. These graphs are not included for every experiment or for every statistic, but only for those that show a significant result. Note that in the ASA graphs, the y-axis units are seconds.

The number of replications used for running the models was determined with an ASA halfwidth of 5 seconds for the intervals 2:00pm and 2:30pm in mind (2:00pm and 2:30pm were chosen because they had the highest halfwidth in the range 8:00am to 5:00pm). Using the formula $n = n_0 h_0^2 / h^2$, available in Kelton, et al (2002), with $n_0 = 8$, $h_0 = 0.09$, and $h = .0833$, n , the number of replications needed, was found to be 9.3, which is rounded up to 10. After running the basic model with ten replications, the 95% confidence intervals were indeed reduced comfortably below five seconds.

For ease of notation, the intervals of the day will often be referenced by their period number, given in the table below (Table 3).

1	2	3	4	5	6	7	8	9	10	11
8:00am	8:30am	9:00am	9:30am	10:00am	10:30am	11:00am	11:30am	12:00pm	12:30pm	1:00pm
12	13	14	15	16	17	18	19	20	21	
1:30pm	2:00pm	2:30pm	3:00pm	3:30pm	4:00pm	4:30pm	5:00pm	5:30pm	6:00pm	

Table 3 Intervals of the Day

4.1 Verification

The first step in the verification was to check for accuracy of the flowchart of the model. This involved using animation to watch Control Entities enter the system and duplicate; call entities leave the Separate module and route correctly to the Call Process module; and duplicate entities enter and leave the wrap Process modules. To see that entities were being routed correctly, the Watch debugger was used to watch the ‘type’ attribute. In addition, the ‘ACD Calls’ attribute was watched to see that it matched not only the data being read (this was done for the other RecordSets) but also the incrementing of the ‘calls per pd’ variable. The built-in module counts that are animated were used to see that the number of entities duplicated matched the number being disposed after the call processing as well as to see that the number being duplicated agreed with the ‘ACD Calls’ attribute. The next step was to verify that the overall number of calls per period

for each skill equaled the sum of the ACD calls for the data. This was checked using the counter statistics collected on the number of calls per period.

The second part of the verification process was verifying that the statistics worked the way that they were expected to. The ASA statistics collection was easy to verify since it is a fairly simple and straightforward process. To see that it was working correctly, the Watch debugger and 'Break on Module' was used to examine the expression 'TNOW-arrivalTime' upon entering the Record module in which wait time is calculated and recorded in the ASA sets. By looking at call volume and agent utilization, the expression was examined for any inconsistencies such as zero wait times followed by one or two single extremely long wait times.

The service level statistics collection verification was much more complicated. The first step was to see that the minimum and maximum values for all periods were not less than zero and not greater than one, respectively. Initially the attributes 'MyInterval' and 'MyDay' were not used. It was discovered through watching the 'cust calls per pd' variable and the 'calls answered in 30s' variable that many recordings of SL were being recorded into the period after the one in which the calls arrived. In addition, some of the SL recordings the 6:30pm period were being recorded into the 8:00am period of the following day. Adding the aforementioned attributes solved many of the problems involved in the SL statistics collection. The Delay module 'Delay so all calls are answered b4 SL recorded' was added after it was discovered through watching the SL variables that even with the 'MyInterval' attribute, some of the calls would arrive at the end of a period and wait less than thirty seconds, but not be included in the SL calculation because the period would end and the calculation would be performed before the call was actually answered.

Much of the statistics verification changes were motivated by the output of the statistics. The minimum and maximum averages of the ASA and SL statistics as well as the overall minimum observations and maximum observations provided valuable feedback for possible problem areas. The cycle time statistics collected were also used to scan for any

glaring inconsistencies (one or more skills occasionally had extremely low or high cycle times). When these outputs appeared normal, a final run-through with the watch expressions was done to check for any additional problems, but none have been found. The model appears to be working as expected.

4.2 *Validation*

The model validation process involved changing the following general aspects of the model:

- ACD Time distribution
- ACW Time distribution
- Failures uptime and downtime
- Using exponential versus uniform interarrival times
- Resource schedules
- Resource set composition

For ease of fitting, it was assumed that the ACD and ACW times would have the same distribution, where both distributions would use the respective 'ACD Time' or 'ACW Time' attribute in the distribution. A beta distribution was used initially, but then was changed to the JohnsonSB distribution. The parameters were determined by starting with a mode of ' $0.9 * \text{mean}$ ', a min of ' $\text{mean}/1.2$ ', and a max of ' $\text{mean}/0.8$ ' and manipulating the median, mode, min, or max until a suitable fit of the ASA and SL was found for the Customer skill (Monday SL). Using exponential interarrival times seemed a better model of the real system and increased the ASA and decreased the SL.

The fitting was done initially using only the Customer ASA and Customer Monday SL because the two have the largest call volume of the five skills; it has been discovered that the other skills have a great effect on the Customer waiting times even though they go through separate process flows. Because of this discovery, it is possible that the fit achieved originally (which was not as close as was preferred) is not the best fit, but due to time restrictions the fit achieved after working with the other four skills is the one being used.

The service level data available is only from approximately three months of calls, so when broken down by interval and by day, each interval mean is calculated from one to eleven data points. Because of the restricted data availability, the data fitting was focused on the ASA. There is eleven months worth of ASA data available. Because finding the overall ASA for each interval involved some calculation (sum of ASA*ACD Calls divided by sum of ACD Calls), it was too cumbersome to calculate the ASA for each day and each interval. Therefore, the data fitting was concentrated on the ASA for each interval without day as a factor. To fit the data, the ACD/ACW Times distribution found from the first round of fitting was kept while changes were made to failures, schedules, and resource set composition. Confidence intervals reported from the Arena output were used to determine how close the fit was. The Appendix contains the graphs of the final fits of the ASA and the SL used in the basic model along with a table (Table A3) indicating the percent difference of the fit from the data for each interval.

There are several possible reasons for the simulation service level being much higher than the actual service level. One possibility is that in the simulation, some calls wait for up to an hour, thereby inflating the ASA. This long wait would of course not occur in real-life; callers would have abandoned long before. These long maximum queue times occur in the simulation because the resources are much more efficient, hence having as few as 44 in the simulation means that there is less ability to absorb high variation. When there are too many resources, queues rarely form and both ASA and SL are too good- the ASA drops well below 30 seconds and the SL rises well above 80%. Queues form in the live system because there is a lot of variation in agent activity; if a couple of agents take a few minutes to make a phone call or get a drink of water, there may be some queuing as a result. Sometimes managers take agents off of the phone when call volumes appear low to do non-call work; call volumes may then increase and queuing results. Another possible explanation for the difference in SL is that during a significant period of the three months that the data was taken from, the call center saw a higher call volume or was understaffed; this would make the SL lower on average for those three months but the yearly average could be much higher. The smoothness of the graphs from the simulation

in contrast with the shape of the SL from the data can be explained by the small number of data points available for each interval and each day. For a given day-interval combination, the simulation uses about 47 observations; the data has anywhere from one to eleven observations for each day-interval combination (by skill).

Due to great difference in data availability of the ASA over the SL, the ASA was fit to the data in all 21 periods for the Customer skill and for 15 to 18 periods for the other four skills. In addition, many of the intervals that were not fit within the confidence interval were off of the data mean by less than five seconds. The SL fits are in general good in shape but are mostly too high, with the exception of a couple of the Broker and Employee days.

4.3 Analysis of Cross-Training Model versus Basic Model

This section describes the analysis of the basic model compared to the cross-training model in terms of its effect on claims output as well as Customer ASA. Detailed analysis on ASA is performed by looking at 5 of the 21 intervals.

4.3.1 Claims Output without Cross-Training

The claims analysis was two-part: first, the model was run with claims reps that were not cross-trained to get an idea of how many claims could be processed with between one and seven claims agents. In the actual data, an average of 1,568 claims are processed per day, amounting to a total of 370,094 over the course of 236 days (the runtime of the simulation). See Table 4 below for the data provided. The second phase was to run the model with the claims agents being cross-trained and the call agents being cross-trained. Runs were made with one of each group being cross-trained up to seven of each group.

Claims Volume Data	3-Jan	3-Feb	3-Mar	3-Apr	3-May	3-Jun	3-Jul	3-Aug
Days	31	28	31	30	31	30	31	31
Monthly total	44516.1	41085.2	51951.9	45541.6	47834.6	44634.6	59816.4	46124
Average Per Day	1436.0	1467.3	1675.9	1518.1	1543.1	1487.8	1929.6	1487.9

Table 4 Claims Volume Data by Month

Table 5 shows the output from the simulation for the total claims volume with one to seven full-time claims agents added without cross-training. It includes the cost for 236 of employing the claims representatives at the yearly salary of \$40,000 (the salary for a cross-trained representative is approximately \$49200). Because the average difference between each additional claims resource has a mean of 18,909.23 with a relatively small standard deviation of 46.7, projections can be made for total claims volume over the 236 day period by adding the mean to the previous output. This projection is done for eight to twenty claims resources, as shown below. At twenty claims resources, the total projected claims output exceeds the data average 236-day output.

Claims Reps	Number of claims out	Halfwidth	Cost
1	18893.3	57.91	36302.70
2	37798.5	28.97	72605.40
3	56704.1	32.59	108908.10
4	75683.5	80.27	145210.80
5	94525.3	102.27	181513.50
6	113462.8	121.15	217816.20
7	132348.7	120.62	254118.90
8	151258	N/A	290421.60
9	170167	N/A	326724.30
10	189076	N/A	363027.00
11	207986	N/A	399329.70
12	226895	N/A	435632.40
13	245804	N/A	471935.10
14	264713	N/A	508237.80
15	283623	N/A	544540.50
16	302532	N/A	580843.20
17	321441	N/A	617145.90
18	340350	N/A	653448.60
19	359260	N/A	689751.30
20	378169	N/A	726054.00

Table 5 Claims Volume Output Totals without Cross-Training

4.3.2 Claims Output with Cross-Training

The simulation was then run with cross-training added, where the claims representatives were trained in skills 266, 258, and 264 (Customer, Provider, and Provider B&E, respectively). Table 6 shows the simulation output for each run with cross-training from one each of calls and claims agents to seven each. From eight to thirteen agents each, the number of claims out is projected by adding the average difference of adding one agent to each group of 30053.5 claims (with a standard deviation of 744.14) to the previous total. The cost of employing for 236 days the respective number of cross-trained calls and claims agents is also given. The calculation is based upon an estimated fully-loaded salary of \$49,200.

Number of Each Group Cross-trained	Number claims out	Halfwidth	Cost
1	30667	581.65	53011.50
2	61384	1224.11	106023.00
3	92019	761.19	159034.50
4	121810	1234.89	212046.00
5	151404	1685.14	265057.50
6	182083	988.52	318069.00
7	210988	764.66	371080.50
8	241042	N/A	424092.00
9	271095	N/A	477103.50
10	301149	N/A	530115.00
11	331202	N/A	583126.50
12	361256	N/A	636138.00
13	391309	N/A	689149.50

Table 6 Claims Volume Output Totals with Cross-Training

Table 6 shows that the total number of cross-trained calls agents and claims agents is 13 each to surpass the total average output from the actual data. This result involves employing seven fewer claims agents total while cross-training 13 existing claims agents and 13 existing calls agents. The percent difference in cost is $(726054.00 - 689149.50) / 726054.00 = .0508$, or a 5% decrease in cost from the basic model. Given the approximate 13,000 increase (an increase of 3.4%) in claims volume that would accompany going from 20 non-cross-trained claims agents to 13 cross-trained, the

decrease in cost is notable. Additionally, at this point the decrease in resource needs for the calls agents as a result of cross-training has not been factored into the cost.

4.3.3 JMP Analysis of Cross-Training Effects on Customer ASA and Claims Output

The Appendix includes graphs of the Customer ASA for the following scenarios: one cross-training each of claims and calls agents, two cross-training each of claims and calls agents, and so on up to seven each (Figures A20-A26, respectively). Given that the company goal is to have an ASA of 30 seconds or less for each interval, it would appear that adding two cross-trained agents to each group would suffice to meet that goal. In terms of percentages, that would involve cross-training 4.167% of the calls staff and the equivalent number of the claims staff. Cross-training any more than that would likely result in overstaffing. However, there is no significant improvement in the SL from cross-training two agents per group, as can be seen in Figure A27. With four cross-trained agents each, the Customer Monday SL shows some improvement in intervals 11:00am-12:30pm and 1:30pm-3:30pm; a graph of the SL values is shown in Figure A28.

To determine the true effects of cross-training, Arena's Process Analyzer (PAN) was used to run eight scenarios with ten replications each: the basic model, cross-training one of each of claims and calls, cross-training two of each, and so on up to seven. Figure 23 below shows the results of the Customer ASA by scenario for 16 of the 21 periods from PAN. It shows a general trend downward for all intervals as more agents are cross-trained. To determine the significance of this trend, JMP was used to analyze the PAN results of the seven cross-training scenarios in addition to a Taguchi 9-run design. The design had three factors of three levels each: Number of Cross-Trained Claims Agents (1, 3, or 5), Number of Cross-Trained Calls Agents (1, 3, or 5), and Total Number of Calls FTEs (42, 43, or 44). The numbers for the three levels of each were based on the preliminary results in the eight scenarios as well as the results of quick experiment in reducing the number of calls agents. In this brief experiment, the simulation was run with two agents of each group cross-trained and the removal of one regular calls agent. It

was run again with four agents of each group cross-trained and the removal of two calls FTEs. The results indicated that removing agents had a great effect on ASA, so removing more than two would bring the ASA well above the 30-60 second range. The Taguchi design is shown below in Table 7. The models for the Taguchi design were all run in PAN with ten replications each. For the Total FTEs, the number 44 refers to the original 44 full-time resources and excludes the additional part-time resources.

Cross-Trained Claims	Cross-Trained Calls	Total Calls FTE
1	1	42
1	3	43
1	5	44
3	1	43
3	3	44
3	5	42
5	1	44
5	3	42
5	5	43

Table 7 Taguchi Design

The responses examined were (Customer) Period 3 ASA, Period 7 ASA, Period 10 ASA, Period 13 ASA, Period 17 ASA, and Number of Claims Out. The periods were chosen in an attempt to have a representation from across the day; period 13 was specifically chosen because it is the period of the day that resulted in the highest ASA. Additionally, the confidence intervals from PAN were graphed for periods 3, 7, 10, 13, and 17 for the first eight scenarios (Figures A32-A36, respectively). It is clear that after a certain point, simply increasing cross-training without changing FTEs has diminished or even null improvement of ASA. The JMP analysis clarifies the effect of total FTE.

Customer ASA by Scenario




Figure 23 Customer ASA by Scenario

In the end, the data used for the statistical analysis in JMP consisted of the average over the ten replications for each response for the nine-run Taguchi design plus the seven previously run scenarios with 48 FTEs on calls as well as two additional scenarios- two each of calls and claims cross-trained with 47 FTEs and 4 each with 46 FTEs (the experiments described earlier that were used to determine the range on the FTEs for the Taguchi design). The data analyzed is shown in Table 8. All main effects and interactions were included in the analysis. After performing the stepwise analysis, the three-factor interaction was found to be insignificant and was removed. Also, one or two or the two-factor interactions were removed in some of the models since they were significant in some responses but not in others. In all cases, the F ratio for the two-factor interactions was near 2, which is much less than the F ratio for the main effects.

Claims CT	Calls CT	FTE Calls	Period 3	Period 7	Period 10	Period 13	Period 17	Claims Out
5	3	42	0.261	0.261	0.29	0.331	0.12	118102.4
1	1	42	0.491	0.626	0.723	0.932	0.437	28211.1
1	3	43	0.489	0.482	0.491	0.781	0.351	55698.4
1	5	44	0.395	0.336	0.32	0.681	0.207	84884.7
3	1	43	0.277	0.235	0.312	0.441	0.172	61260.8
3	5	42	0.344	0.387	0.457	0.547	0.209	108069.2
5	1	44	0.15	0.119	0.132	0.248	0.045	98169.5
5	5	43	0.275	0.218	0.2	0.321	0.089	147018.3
3	3	44	0.236	0.192	0.22	0.387	0.106	91843.9
1	1	44	0.281	0.302	0.284	0.644	0.189	30667.6
2	2	44	0.231	0.234	0.282	0.477	0.146	61778.2
3	3	44	0.236	0.192	0.22	0.387	0.106	91843.9
4	4	44	0.246	0.162	0.188	0.377	0.091	122144.1
5	5	44	0.211	0.159	0.17	0.315	0.073	151884.8
6	6	44	0.193	0.144	0.173	0.271	0.067	181581.3
7	7	44	0.195	0.136	0.165	0.247	0.062	210070.6
2	2	43	0.357	0.325	0.356	0.481	0.202	58845.6
4	4	42	0.334	0.293	0.337	0.469	0.147	113467.5

Table 8 PAN Data used in JMP Analysis

The Normal Quantile Plots (Figure 24) show all 18 PAN data points to be within the range of normality for each of the six outputs. The 'Fit Model' tool was used with a Stepwise approach based on the F-ratios with a goal of minimizing the AIC criterion to find the best Least Squares model to run. Figures 25-36 show the results of the statistical analysis for Period 3, Period 7, Period 10, Period 13, and Period 17 ASA and Number of Claims Out, respectively.

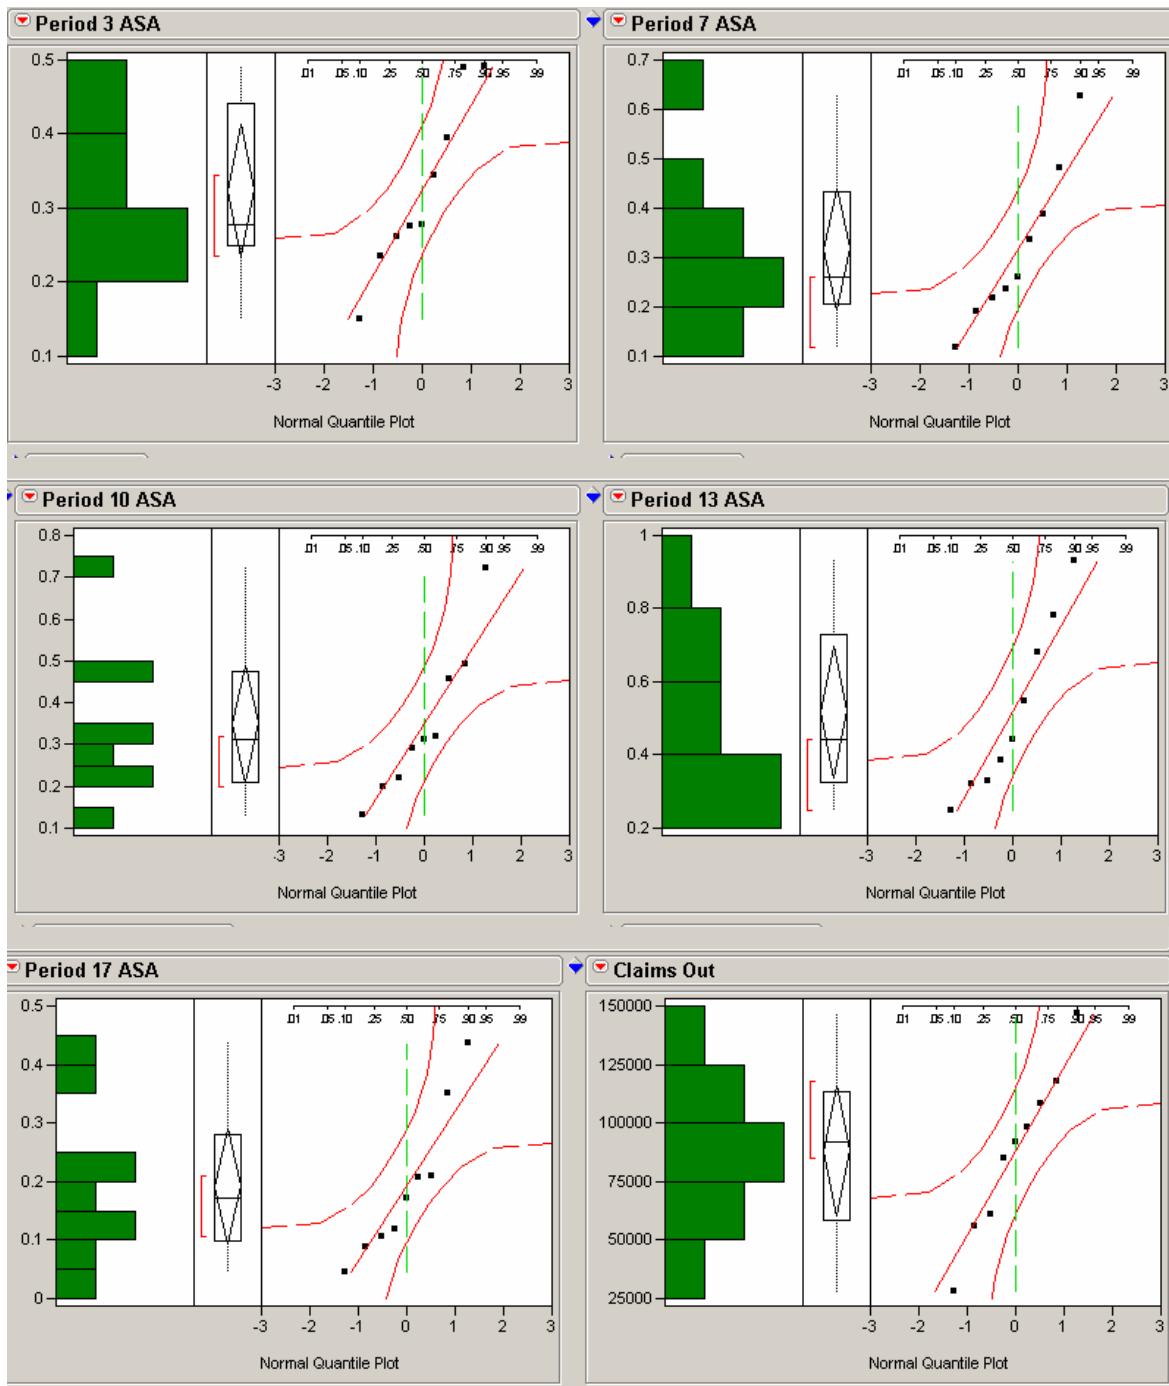


Figure 24 Normal Quantile Plots of Responses

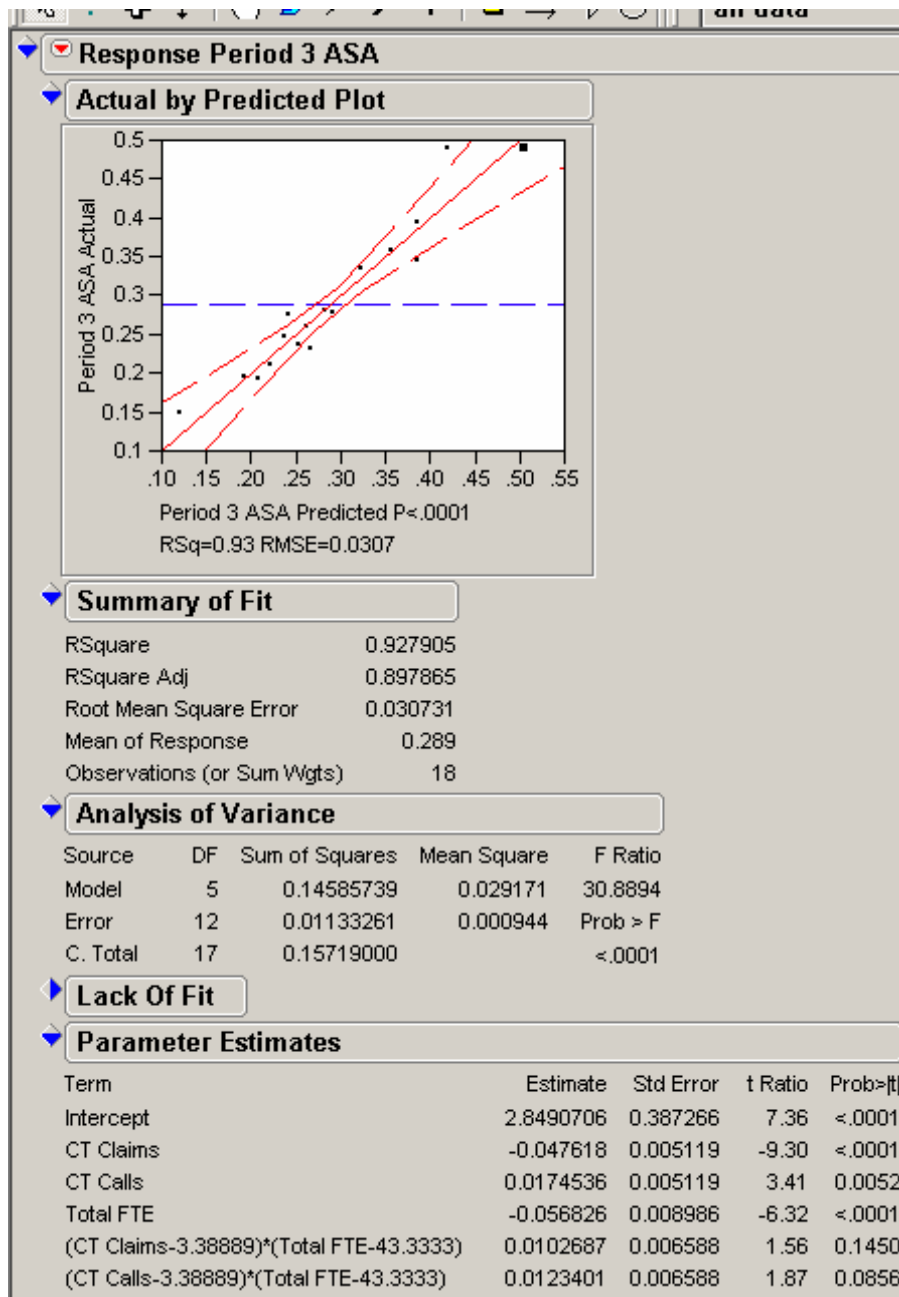


Figure 25 Period 3 ASA Least Squares Regression Model Results

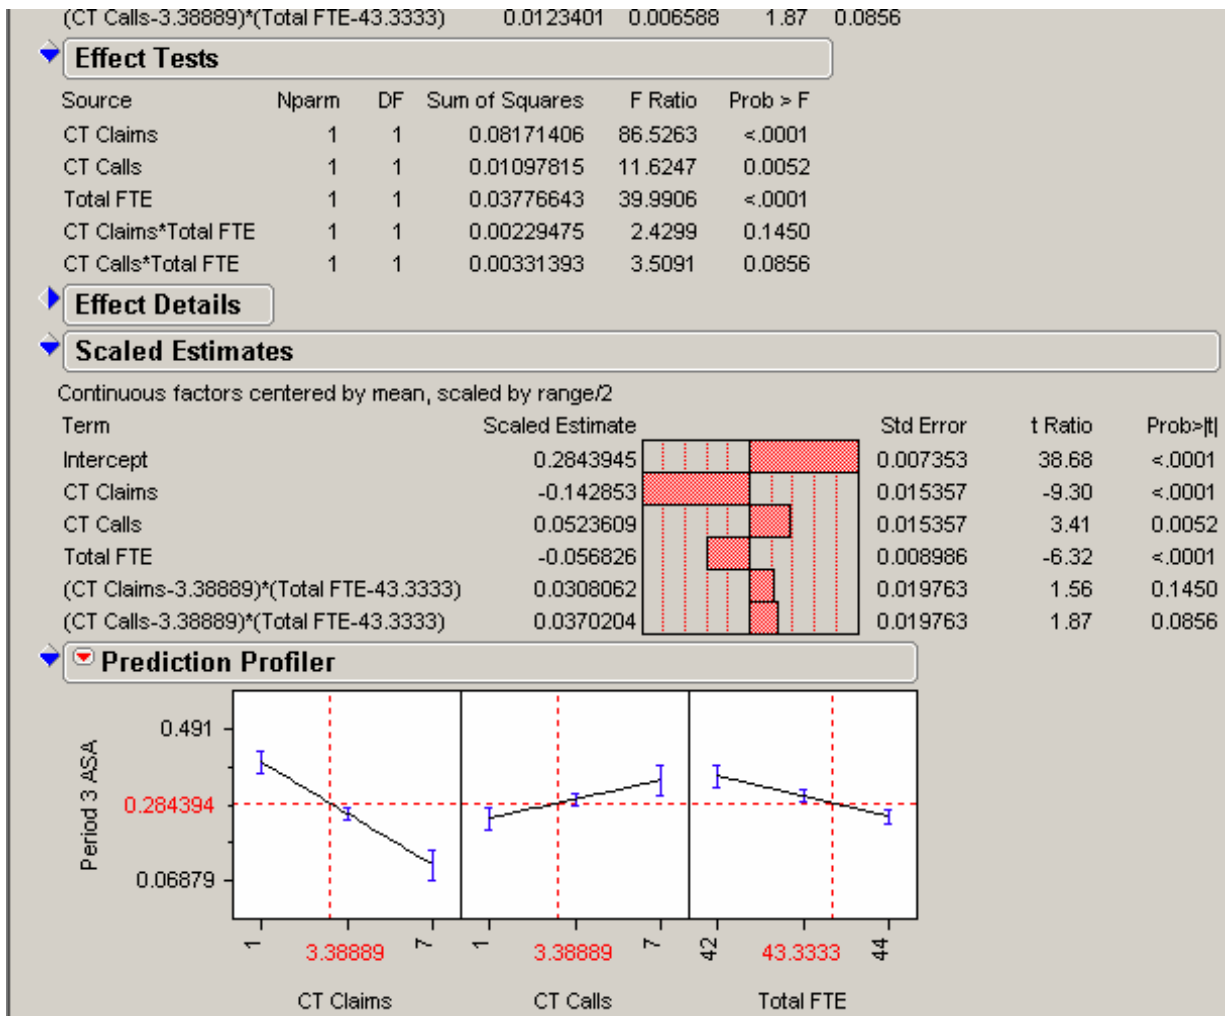


Figure 26 Period 3 ASA Least Squares Regression Model Results




Figure 27 Period 7 ASA Least Squares Regression Model Results

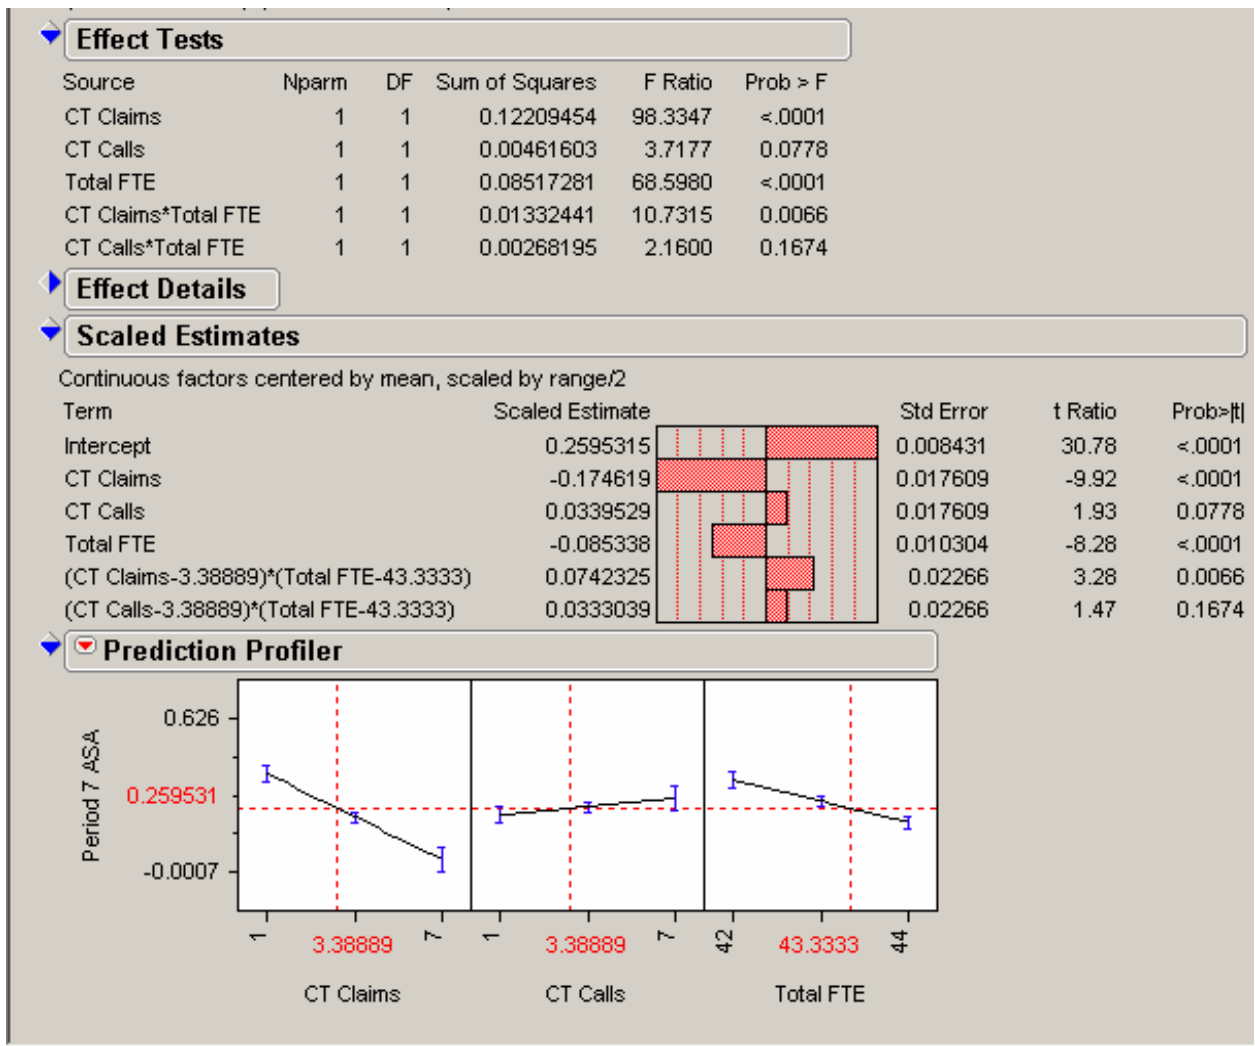


Figure 28 Period 7 ASA Least Squares Regression Model Results




Figure 29 Period 10 ASA Least Squares Regression Model Results

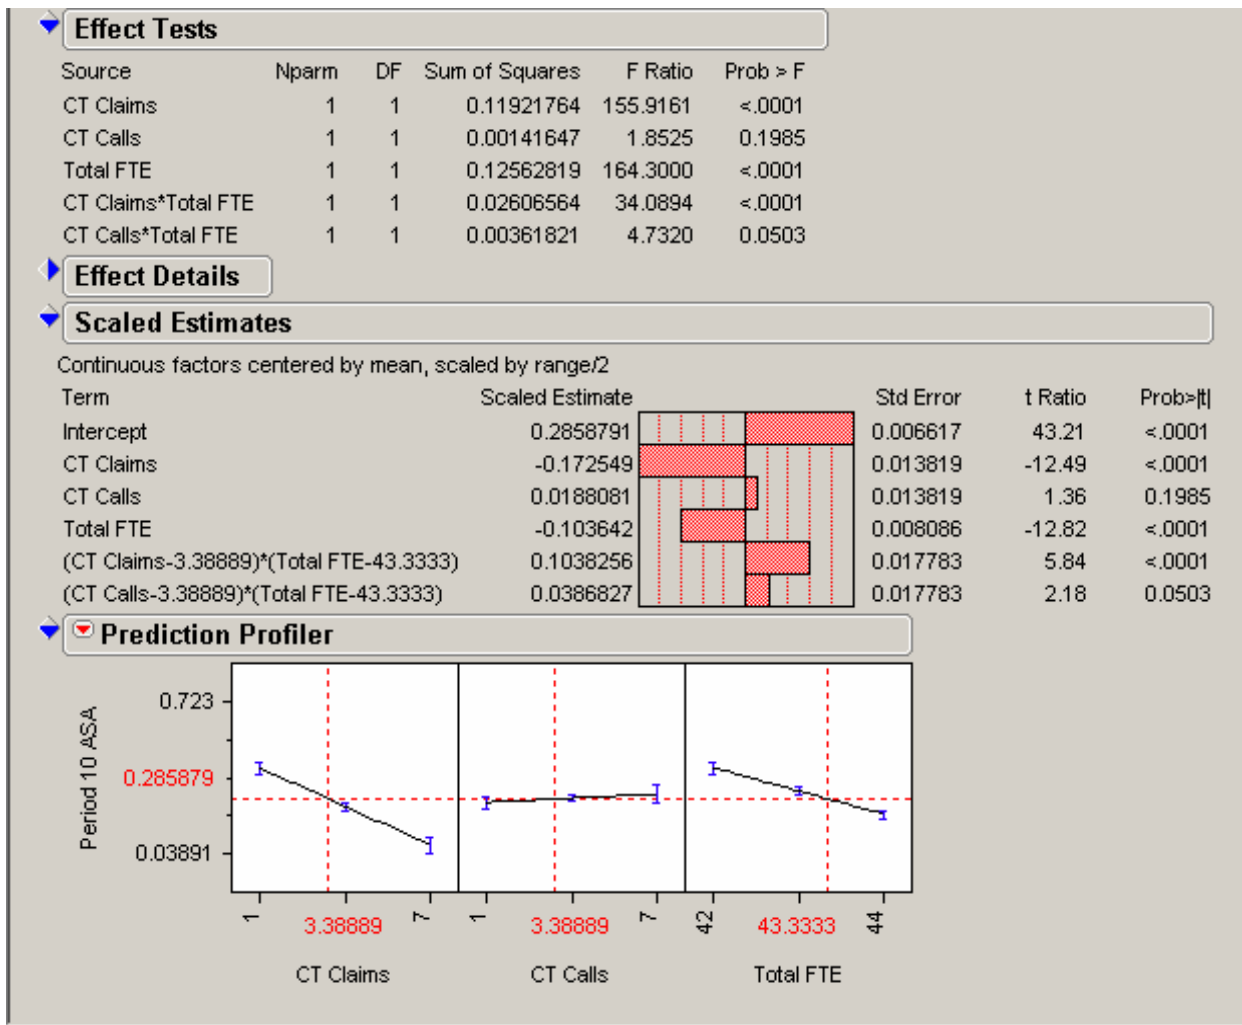


Figure 30 Period 10 ASA Least Squares Regression Model Results

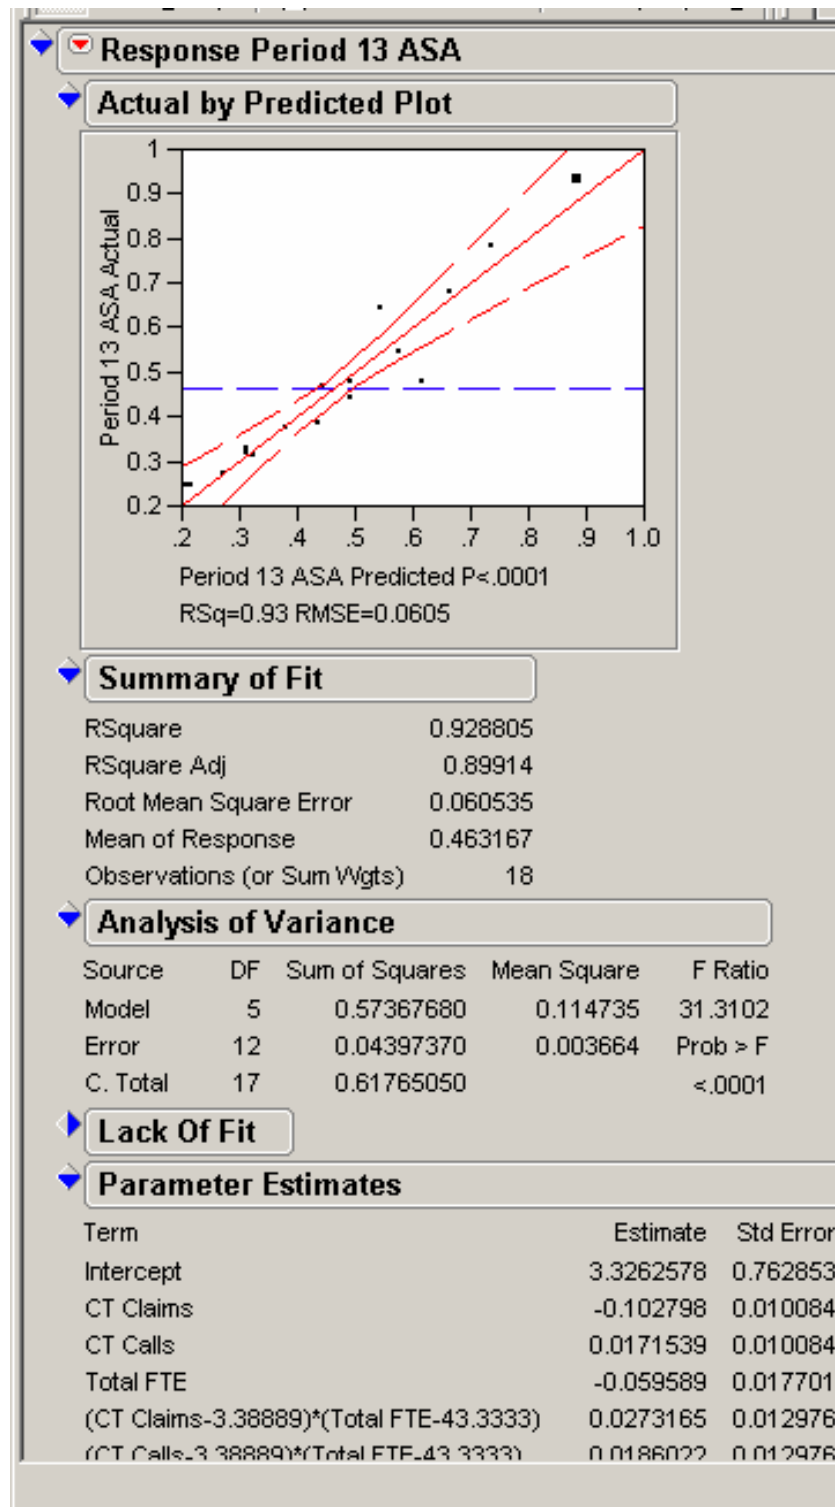


Figure 31 Period 13 ASA Least Squares Regression Model Results

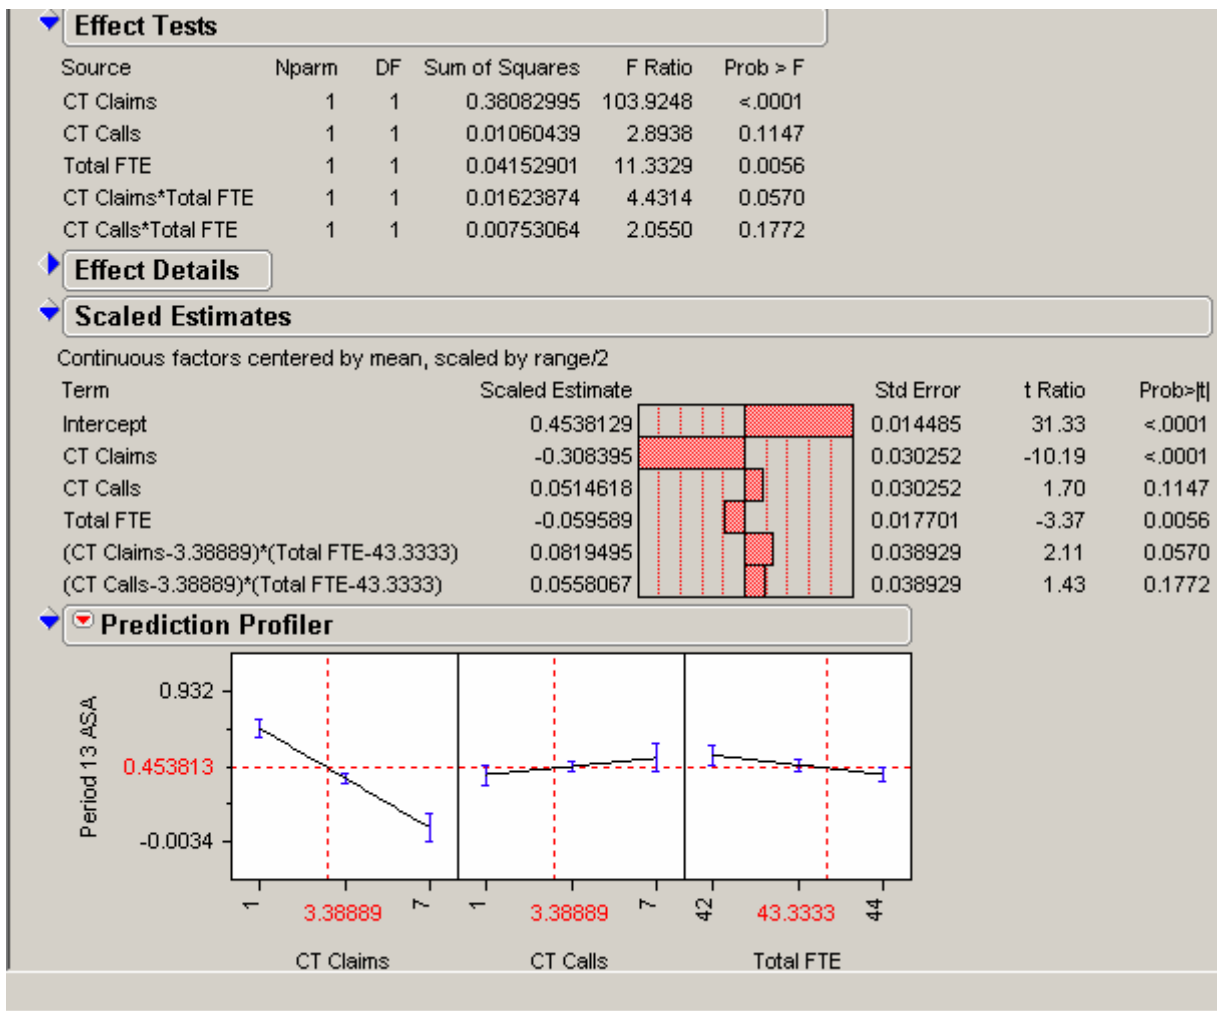


Figure 32 Period 13 ASA Least Squares Regression Model Results




Figure 33 Period 17 ASA Least Squares Regression Model Results

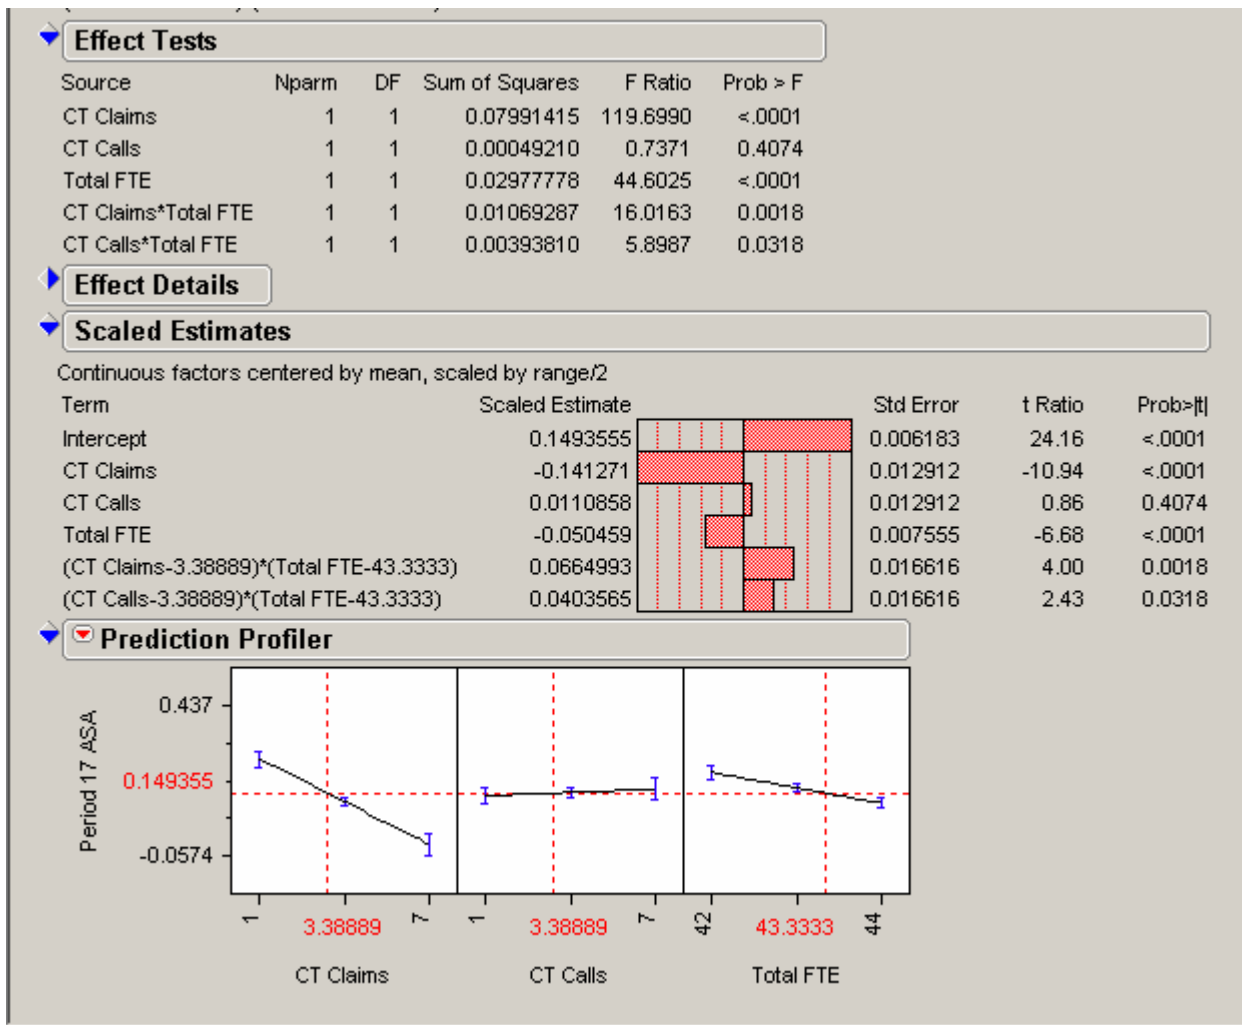


Figure 34 Period 17 ASA Least Squares Regression Model Results

Response Claims Out					
Summary of Fit					
RSquare			0.999853		
RSquare Adj			0.999792		
Root Mean Square Error			713.7746		
Mean of Response			100863.4		
Observations (or Sum Wgts)			18		
Analysis of Variance					
Source	DF	Sum of Squares	Mean Square	F Ratio	
Model	5	4.16623e10	8.33245e9	16355	
Error	12	6113689.48	509474.12		
C. Total	17	4.16684e10			
Lack Of Fit					
Parameter Estimates					
Term		Estimate	Std Error	t Ratio	Prob> t
Intercept		-158360.9	9038.083	-17.52	<.0001
CT Claims		16581.972	117.9712	140.56	<.0001
CT Calls		12930.843	118.5847	109.04	<.0001
Total FTE		3673.7328	209.6503	17.52	<.0001
(CT Claims-3.38889)*(CT Calls-3.38889)		-81.03736	53.65298	-1.51	0.1568
(CT Calls-3.38889)*(Total FTE-43.3333)		824.1716	146.1076	5.64	0.0001
Effect Tests					
Source	Nparm	DF	Sum of Squares	F Ratio	Prob > F
CT Claims	1	1	1.00657e10	19756.98	<.0001
CT Calls	1	1	6057852275	11890.4	<.0001
Total FTE	1	1	156439427	307.0606	<.0001
CT Claims*CT Calls	1	1	1162264.67	2.2813	0.1568
CT Calls*Total FTE	1	1	16211080.1	31.8192	0.0001

Figure 35 Number of Claims Out Least Squares Regression Model Results

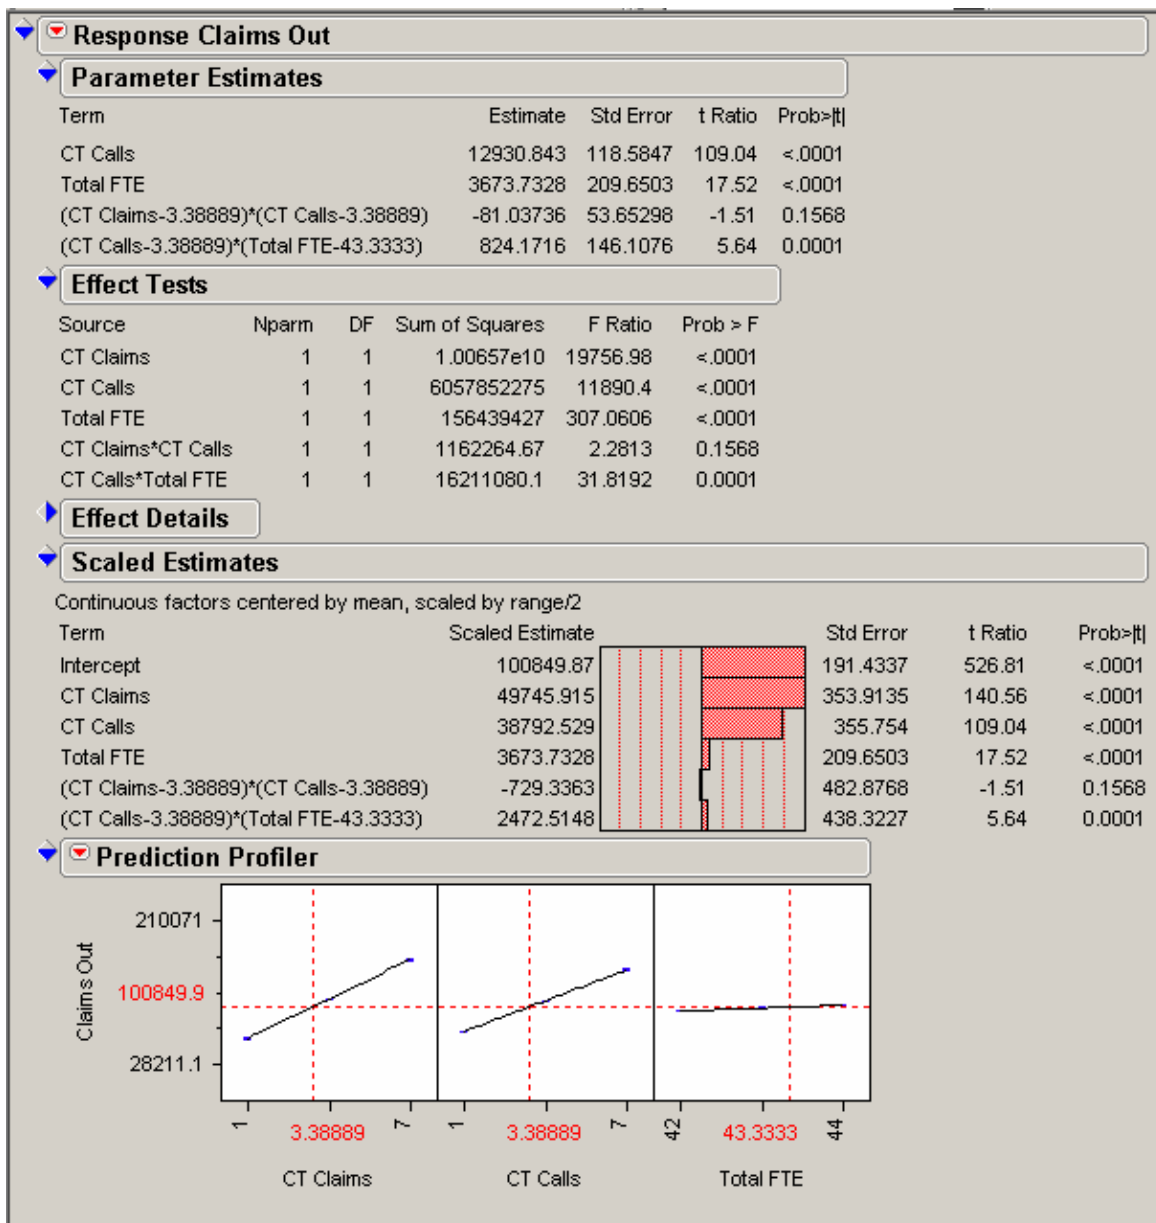


Figure 36 Number of Claims Out Least Squares Regression Model Results

First, note the F ratio for the ANOVA on each of the six responses: the least significant of the six is Period 3 ASA with an F ratio of 30.89. It also has the lowest adjusted R-Square value of 0.89, which is also good. The Claims Out model appears to be completely linear- its R-Square value is almost exactly 1. It is clear that the respective model is a good fit for each of the responses.

In all of the ASA models, the Prediction Profiler shows a clear negative correlation between the number of cross-trained claims agents and the ASA. This result implies that cross-training claims agents reduces ASA. The prediction profiler also indicates a negative correlation between the number of calls FTEs, which is logical since a higher number of resources reduces the ASA. The correlation between ASA and cross-trained calls agents appears to be weakly positive. For the Claims Out model, the correlation between number of claims out and number of FTEs is, not surprisingly, weak; the correlation with respect to cross-trained claims agents and cross-trained calls agents is strongly positive, as indicated by the prediction profiler. The model indicates, therefore, that claims volume out is heavily impacted in the positive direction by the amount of cross-training present.

Examining the parameter estimates and effect tests paints the complete picture. In the ASA responses, there is a slight interaction between calls FTEs and cross-trained claims agents as well as with cross-trained calls agents. The number of claims agents cross-trained has a highly significant effect which seems to increase as the ASA increases: from period 3 to period 13, the ASA increases in the basic model (Figure A1) and the parameter estimate on CT Claims increases. From period 13 to period 17 the parameter estimate decreases, as does the ASA. This result confirms the purpose of looking to cross-training in the first place for lowering ASA: during periods of poor ASA, bringing in extra help from claims agents clearly helps to relieve the queuing. During periods where there is less queuing, the claims agents are rarely needed and therefore have less of an impact on the system. The effect of cross-trained calls agents on ASA is less clear, though there is some indication that the effect is higher during slower periods. When there is no queuing, cross-trained calls agents will work on claims more often and are more likely to be busy with a claims when new calls come in, possibly resulting in some queuing and hence, the former results.

The parameter estimates were put into Microsoft Excel to generate a table of the value of each response for a given set of parameter values. Shown below, a partial representation of the table shows several options that result in a maximum ASA across the five intervals

of less than 30 seconds, an increase in claims output, and a decrease in cost. The cost column is calculated by finding the change in wages of cross-training existing agents as well as subtracting the wages of any unneeded agents (i.e. if Total FTE is 38, then the salaries of 6 agents would be subtracted) over the 236-day period. According to the parameter estimates, of the values chosen for the inputs the greatest cost savings (\$200,895-a savings of 8.4%) come from cross-training five claims and three calls agents, eliminating the need for 6 of the 44 calls agents (reducing the total calls staff by 12.5%). In this case, an additional 10,176 claims could be processed- an increase of 10.8%. If the company focus were on claims output, it is better to cross-train more of the calls agents; in the table below, this highest additional claims output occurred when four claims and 5 calls agents were cross-trained, resulting in an additional 44,386 claims processed at a savings of \$58,675 (a savings of 2.5%).

CT Claims	CT Calls	Total FTE	Period 3 ASA	Period 7 ASA	Period 10 ASA	Period 13 ASA	Period 17 ASA	Additional Claims Output	Change in Cost
4	1	42	19.2	19.6	23.2	28.0	12.2	2246.7	-47471.40
4	2	42	19.3	19.3	22.6	27.5	11.4	14031.9	-39117.00
4	3	42	19.3	19.1	21.9	27.1	10.5	25817.0	-30762.60
4	4	41	22.0	22.7	25.7	28.5	11.4	33424.8	-67029.90
4	5	41	21.3	21.9	24.3	26.9	9.7	44385.7	-58675.50
4	5	40	23.1	25.0	28.0	27.7	10.6	39384.2	-103297.20
4	5	39	25.0	28.1	31.7	28.5	11.6	34382.7	-147918.90
5	2	40	22.5	21.2	24.2	24.1	10.8	6826.8	-120006.00
5	2	39	25.9	24.8	28.2	26.6	12.8	4297.9	-164627.70
5	3	40	21.1	19.6	22.0	21.5	8.3	16882.5	-111651.60
5	3	39	23.8	22.6	25.2	22.8	9.5	13529.4	-156273.30
5	3	38	26.5	25.6	28.4	24.2	10.7	10176.3	-200895.00
5	4	40	19.6	18.1	19.8	18.8	5.9	26938.1	-103297.20
5	4	39	21.6	20.4	22.2	19.0	6.2	22760.9	-147918.90
5	4	38	23.6	22.7	24.6	19.3	6.6	18583.6	-192540.60
5	5	40	18.2	16.6	17.6	16.1	3.4	36993.8	-94942.80
5	5	39	19.5	18.2	19.3	15.2	3.0	31992.3	-139564.50
5	5	38	20.7	19.9	20.9	14.4	2.6	26990.9	-184186.20
5	2	41	19.0	17.5	20.3	21.7	8.8	9355.7	-75384.30
5	3	41	18.4	16.7	18.9	20.1	7.1	20235.6	-67029.90
5	4	41	17.7	15.8	17.4	18.5	5.5	31115.4	-58675.50
5	5	41	17.0	14.9	16.0	16.9	3.8	41995.3	-50321.10
5	2	42	15.6	13.9	16.3	19.2	6.8	11884.6	-30762.60

Table 9 Use of Parameter Estimates

Table 10 below shows the actual parameter estimates from JMP. Recall that in the basic model, the ASA generally increases from periods 1-13 and begins to decrease after period 13. The absolute value of the CT Claims estimates follows this general trend, indicating that as more resources are needed, cross-trained claims agents play a more important role in the ASA. Total FTE follows a similar trend except that its absolute peak is during period 10.

When comparing CT Claims with CT Calls, it can be seen that CT Claims has much more influence on the ASA over CT Calls. From period 1 to 17 respectively, it is about 3, 5, 9, 6, and 13 times as influential as CT Calls. Likewise, Total FTE is about 3, 8, 17, 3, and 14 times, respectively, more influential than CT Calls. CT Claims and Total FTE have a comparable influence: Total FTE is anywhere from 0.6 to 1.8 times more influential than CT Claims. However, Total FTE clearly has the least influence on Claims Output. CT Claims is 4.5 times more important than Total FTE on claims output; similarly, CT Calls is 3.5 times more influential. CT Claims is only slightly more influential than CT Calls, but this is because of the design of the model- there has not been a separation between cross-trained and regular claims. Obviously, cross-training claims agents takes them away from their claims work at times so it would affect claims output negatively. This can be seen when the parameter estimates are used with zero CT Calls to see the claims output from the cross-trained claims agents: the claims output actually decreases from the original claims model with claims agents that are not cross-trained.


	Intercept	CT Claims	CT Calls	Total FTE	CT Claims *Total FTE	CT Calls *Total FTE	CT Claims *CT Calls
Period 3	2.849071	-0.047618	0.0174536	-0.056826	0.0102687	0.0123401	
Period 7	4.116419	-0.058206	0.0113176	-0.085338	0.0247442	0.0111013	
Period 10	4.950794	-0.057516	0.0062694	-0.103642	0.0346085	0.0128942	
Period 13	3.326258	-0.102798	0.0171539	-0.059589	0.0273165	0.0186022	
Period 17	2.482976	-0.04709	0.0036953	-0.050459	0.0221664	0.0134522	
Claims Out	-158361	16581.97	12930.8	3673.7		824.2	-81.03

Table 10 Parameter Estimates

The simulation outputs suggest that minimal cross-training results in a significant improvement of customer service in most periods of the day, especially those periods of higher call volume and lowered agent availability, i.e. the late morning and early afternoon hours. This result agrees with the original motivation for investigating cross-training since the primary purpose is to aid in these poorer-service intervals; the way it is implemented (i.e. claims agents only get on the phones when there is more queuing than call agents can easily control) implies that there would naturally be little improvement in the intervals that already have a low occurrence of queuing. In addition, the cost of cross-training would involve not only the cost of paying higher salaries to cross-trained associates, but also the cost of the increased training required. However, with cross-training comes service improvement and hence, a lower need for call center agents overall, bringing savings from the salaries of the agents that are no longer needed.

4.4 *Analysis of Scheduling Model versus Basic Model*

The lunch schedules of the employees were modified in an attempt to smooth out the ASA curve for each skill. A smoother curve is desired because during the bulk of the day, the staffing level does not change (even half of the time spent at lunch is paid time off); when ASA is too high, the call center is experiencing understaffing and when ASA is too low, the call center is experiencing overstaffing. Ideally, say Cleveland and Mayben, a call center wants its performance to be as consistent as possible to optimize the staffing level. The results of the schedule changes are shown below in Figure 37.






Figure 37 ASA Results of Scheduling versus No Scheduling

Clearly, simply shifting the schedules in the model resulted in a much smoother ASA curve. It is likely that an even smoother fit could be achieved with more resources. In the actual call center, there are many more agents and hence many more options for schedule changes. In the model, the smoothness of the ASA curve was limited by the number of schedules available to change since the effect one change may have on one skill may be opposite the effect the change would have on another skill.

Because the schedules in the model were created based on fitting the data and were significantly altered from the original sample schedule obtained from the company, this result addresses more than just the issue of proper scheduling. It is suspected that the majority of the increase in midday ASA levels is due not to lunch breaks but to post-lunch slowness, managers assigning non-call work, or a myriad of other schedule adherence issues. It is likely that most agents really were scheduled to take lunches

earlier in the day, so this result suggests that an equally important issue to address is schedule adherence.

Had the company not only based their schedules around call volumes, but also enforced schedule adherence, the resulting ASA would likely have been much more consistent throughout the day. Also notice that by simply shifting the schedules, the highest mean ASA was only 43 seconds, as compared with the original high of 58 seconds. The cost of improving the schedule and maintaining a good schedule as well as maintaining schedule adherence is difficult to define, but would likely involve the use of a specially designed call center scheduling software as well as the employment of personnel to design the schedule and use the scheduling software to monitor adherence.

4.5 Analysis of Cross-Training and Scheduling Together

The scheduling changes that were made were added to the model with cross-training. The simulation was again run with one to seven of each group cross-trained. Not surprisingly, there was almost no change in the claims output, as shown in Table 11. The effect of having both cross-training and improved scheduling seems to be a general decrease in ASA and a smoothing of the ASA curve. Since this analysis is concerned with the shape of the curve, the models were simply run in Arena and the outputs were graphed directly.

Number of Each Group Cross-trained	Number claims out	Halfwidth
1	30294	772.76
2	61486	534.78
3	91576	1553.09
4	120716	3586.88
5	151914	1669.48
6	181372	2418.06
7	210775	1443.43

Table 11 Claims Output from Model with Scheduling and Cross-Training

When one agent from each group is cross-trained, there is clearly a decrease in the ASA for the later half of the day (Figure 38). The case is similar for the case of cross-training two agents each, only the overall ASA is lowered still further (Figure 39). Cross-training just one agent from each group brings the ASA down comfortably below the 30-second mark in every interval.

With three agents each, the ASA confidence interval is lower than the original fit in intervals 9:30am-10:30am, 1:30pm-4:00pm, and 4:30pm-5:00pm: there does not appear to be much of a change from two to three. However, cross-training four in each group adds period 1:00pm-1:30pm and 4:00pm-4:30pm to the former list. The graph shows a much more consistent ASA with a maximum of just 18 seconds. When a fifth agent from each group is cross-trained, the ASA CI is lower in the 9:00am-10:30am and 1:00pm-5:30pm ranges and the maximum drops to less than 15 seconds. Adding a sixth shows a strong decrease in ASA during periods 8:30am-12:00pm and 1:00pm-5:30pm with a maximum of 15 seconds; there appears to be a slight drop in overall ASA from five to six, but the difference is minimal. Adding a seventh brings little improvement: the maximum average drops to 14 seconds.




Figure 38 Customer ASA with Cross-Training One Calls and One Claims Agent Along with Improved Scheduling




Figure 39 Customer ASA with Cross-Training Two Calls and Two Claims Agents Along with Improved Scheduling




Figure 40 Customer ASA with Cross-Training Three Calls and Three Claims Agents Along with Improved Scheduling




Figure 41 Customer ASA with Cross-Training Four Calls and Four Claims Agents Along with Improved Scheduling




Figure 42 Customer ASA with Cross-Training Five Calls and Five Claims Agents Along with Improved Scheduling




Figure 43 Customer ASA with Cross-Training Six Calls and Six Claims Agents Along with Improved Scheduling




Figure 44 Customer ASA with Cross-Training Seven Calls and Seven Claims Agents Along with Improved Scheduling

5. Conclusion

This thesis has shown with statistical significance that cross-training improves ASA and claims output. It is also clear that manipulating agent schedules can have a great impact on ASA during the busiest periods. Combined with schedule adherence, this can reduce overstaffing during the lighter call volume periods and reduce understaffing during the periods that are hit the hardest. When cross-training is combined with improved scheduling, the result is an increase in claims output, an overall decrease in ASA and reduction in staffing needs, and a smoothing of the ASA curve that eliminates strong peaks in ASA and levels out staffing needs throughout the day.

The simulation has been valuable in providing a flexible environment in which to model the call center and gather information about the behavior of the call center. It has shown that there is a direct and linear relationship between the number of claims agents- both cross-trained and not- and the number of claims output, which eliminates the need for complex calculations in determining how to staff the claims department. While the verification of the model was relatively simple, the validation was difficult enough to recommend using simulation only in complex call centers. The difficulty came in balancing ASA with SL as well as balancing the five skills across all of the intervals. Because of the failures and random distributions of the talk times and wrap times, balancing these responses was difficult and tedious.

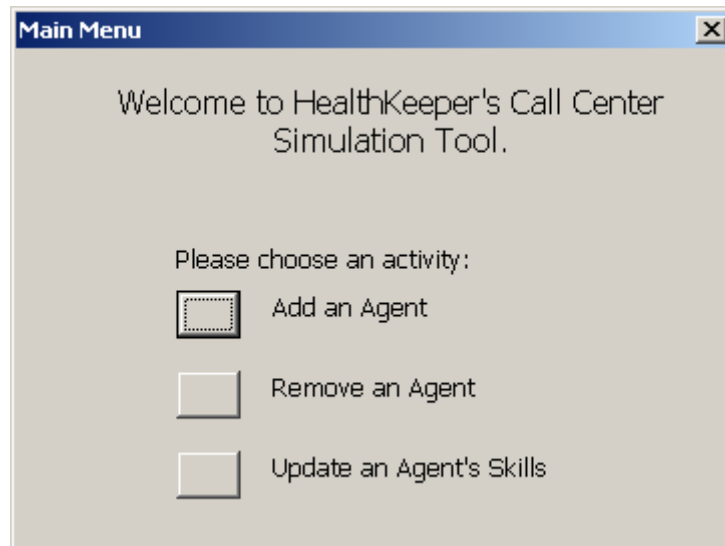
The analysis has shown generally what was expected from the simulation. To improve upon the cross-training analysis, it is recommended that all replications be used as data points and higher order factors and interactions be considered in the regression. However, the models used do provide strong results and give responses that make intuitive sense. Further analysis on the improved scheduling could be done by manipulating the model such that Arena's OptQuest could search for the schedule that would minimize ASA or SL. This change to the model may be made by introducing

variables as the schedule durations that OptQuest would change the values of in searching for the optimum.

Originally, a sub-goal of this thesis was to create a simulation model that could be used by the call center to manage their scheduling and cross-training needs with a salary input to assist in finding the best combination of good service and lower costs. As a first step in this process, VBA in Arena was used to create user forms that would add resources to the model with custom schedules and skill assignments (Figures 45 and 46). The forms were helpful when building the model to easily add agents, but the forms were not completed in such a way that the company could use the simulation. Additional work needs to be done to add forms that update the data used to run the simulation and add custom reports on cost and ASA information.

Further suggestions for future work include the areas of employee satisfaction and the ideal percentage of staff to cross-train on both teams. Job satisfaction plays a large role in productivity; it is quite possible that cross-trained employees would have higher job satisfaction due to the added variety of job functions and higher pay. This increased job satisfaction may increase productivity and further lessen the staffing needs of the call center. Moreover, with increase job satisfaction may come decreased attrition and hence a reduction in cost from the expenses involved in termination of employment, hiring, and training.

This thesis has investigated to some extent the ideal percentage of staff to cross-train, but because there are three goals involved in determining this proportion it has not been investigated in detail. In determining the optimal distribution of cross-trained and non-cross-trained associates the call wait time objectives, claims output targets, and cost goals must all be resolved. Future work could be done in this area using OptQuest to help call centers better determine their staffing arrangements.



Main Menu

Welcome to HealthKeeper's Call Center Simulation Tool.

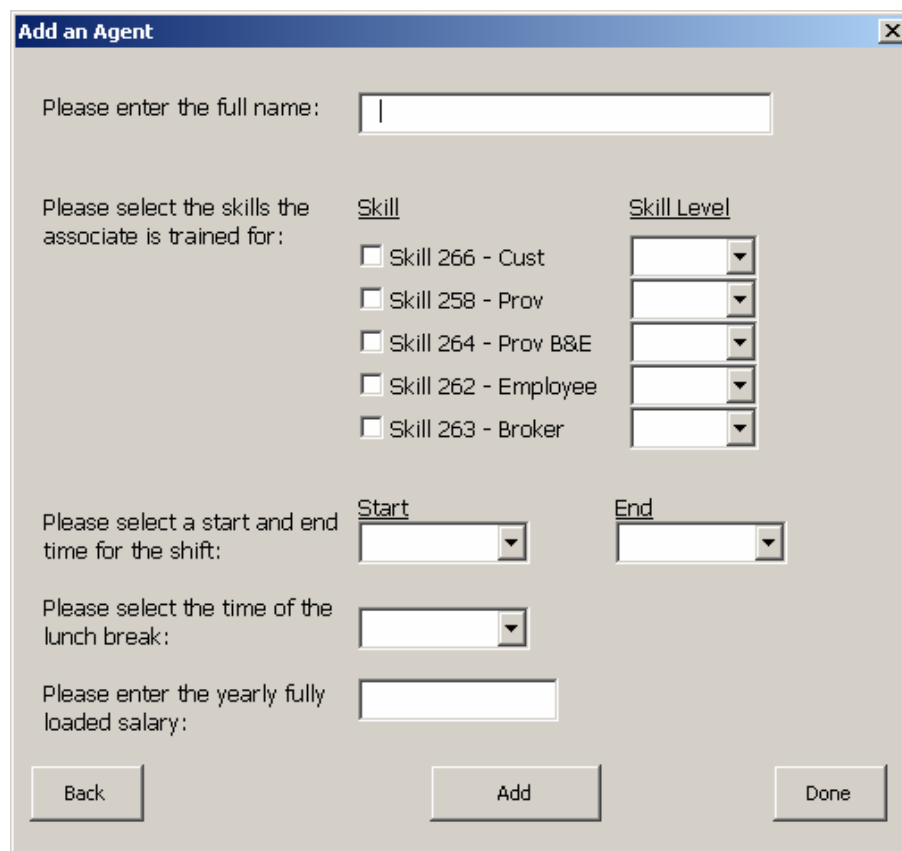
Please choose an activity:

☒ Add an Agent

☐ Remove an Agent

☐ Update an Agent's Skills

Figure 45 Main Menu Form



Add an Agent

Please enter the full name:

Please select the skills the associate is trained for:

Skill	Skill Level
<input type="checkbox"/> Skill 266 - Cust	<input type="text"/>
<input type="checkbox"/> Skill 258 - Prov	<input type="text"/>
<input type="checkbox"/> Skill 264 - Prov B&E	<input type="text"/>
<input type="checkbox"/> Skill 262 - Employee	<input type="text"/>
<input type="checkbox"/> Skill 263 - Broker	<input type="text"/>

Please select a start and end time for the shift:

Start: End:

Please select the time of the lunch break:

Please enter the yearly fully loaded salary:

Back Add Done

Figure 46 Add Agent Form

Reference List

1. Andrews, Bruce and Henry Parsons. 1993. Establishing Telephone-Agent Staffing Levels through Economic Optimization. *Interfaces* 23, (2): 14-20.
2. Askin, O. Zeynep and Patrick T.Harker. 2001. Modeling a Phone Center: Analysis of a Multichannel, Multiresource Processor Shared Loss System. *Management Science* 47, (2): 324-336.
3. Askin, O. Zeynep and Patrick T.Harker. 2003. Capacity sizing in the presence of a common shared resource: Dimensioning an inbound call center. *European Journal of Operational Research* 147: 464-483.
4. Avramadis, Athanassios N., Alexandre Deslauriers, and Pierre L'Ecuyer. 2003. Modeling Daily Arrivals to a Telephone Call Center.
5. Bapat, Vivek and Jr Eddie B.Pruitt. 1998. Using Simulation in Call Centers. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D.J.Medeiros, E.F.Watson, J.S.Carson, and M.S.Manivannan, 1395-1399.
6. Brown, Lawrence, Noah Gans, Avishai Mandelbaum, and et al. 2002. Statistical Analysis of a Telephone Call Center: A Queueing-Science Perspective.
7. Cleveland, Brad and Julia Mayben. 1997. *Call Center Management on Fast Forward: Succeeding in Today's Dynamic Inbound Environment*. First ed. Annapolis: Call Center Press.
8. Cooper, Robert B. 2004. "Queuing Theory." Available on-line via <http://www.cs.usm.maine.edu/~pfiorini/erlang-B-C-models.pdf> [accessed May 6,2004].
9. Kalimo, Raija, Keri Lindstrom, and Michael J.Smith. 1997. Psychosocial Approach in Occupational Health. In *Handbook of Human Factors and Ergonomics*, edited by Gavriel Salvendy 2nd ed. (New York: John Wiley and Sons, Inc.).
10. Kelton, W. David, Randall P.Sadowski, and Deborah A.Sadowski. 2002. *Simulation with Arena*. New York: The McGraw-Hill Companies.
11. Klungle, Roger. 1997. The Role of Simulation in Call Center Management. In *Proceedings of the 1997 MSUG Conference*.
12. Klungle, Roger. 1999. Simulation of a Claims Call Center: A Success and a Failure. In *Proceedings of the 1999 Winter Simulation Conference*, ed.

- P.A Farrington, H.B.Newbhard, D.T.Sturrock, and G.W.Evans, 1648-1653.
13. Koole, Ger and Erik van der Sluis. 2002. Optimal shift scheduling with a global service level constraint. *IIE Transactions* 35: 1049-1055.
 14. Leshner, Martin and Anne Browne. 1993. Increasing efficiency through cross-training. *Best's Review*, 1993, 39-40.
 15. Mehrotra, Vijay and David Profozich. 1997. Simulation: the best way to design your call center. *Telemarketing & Call Center Solutions*, 1997, 28-34.
 16. Pichitlamken, Jutta, Alexandre Deslauriers, Pierre L'Ecuyer, and Athanassios N.Avramidis. 2003. Modeling and Simulation of a Telephone Call Center. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S.Chick, P.J.Sanchez, D.Ferrin, and D.J.Morrice.
 17. Swezey, Robert W. and Robert E.Llaneras. 1997. Models in Training and Instruction. In *Handbook of Human Factors and Ergonomics*, edited by Gavriel Salvendy 2nd ed. (New York: John Wiley and Sons, Inc.).
 18. Tanir, Oryal and Richard J.Booth. 1999. Call Center Simulation in Bell Canada. In *Proceedings of the 1999 Winter Simulation Conference*, 1640-1647.
 19. Wickens, Christopher D. 1992. *Engineering Psychology and Human Performance*. New York City: Harper-Collins.

Appendix

The appendix contains tables and figures that provide information about the basic model, support the validation, and support the analysis. The first set of tables give the resource schedule and skill assignment information. The following set of figures, Figures A1-A19, and Table A3 show the fit of the basic model to the data. The next set of figures, Figures A20-A31, show the graphs of the ASA with cross-training in different scenarios compared with the basic model ASA. Figures A32-A36 graph the confidence intervals over various cross-training scenarios of periods 3, 7, 10, 13, and 17.

Name	Shift					Lunch								Exceptions
	8:00-4:30	8:30-5:00	9:00-5:30	9:30-6:00	9:30-6:30	11:00 AM	11:30 AM	12:00 PM	12:30 PM	1:00 PM	1:30 PM	2:00 PM	2:30 PM	
Resource 1	1								1					
Resource 2			1							1				
Resource 3	1									1				
Resource 4	1												1	
Resource 5		1							1					
Resource 6	1							1						1.5 hour lunch
Resource 7		1							1					
Resource 8			1								1			
Resource 9			1											no lunch
Resource 10			1								1			
Resource 11					1								1	
Resource 12					1							1		
Resource 13	1					1								
Resource 14	1											1		
Resource 15	1													no lunch
Resource 16	1								1					
Resource 17		1								1				
Resource 18		1								1				
Resource 19		1								1				
Resource 20		1												9:30-11:00 "lunch"
Resource 21		1								1				8:30-6:00
Resource 22		1										1		
Resource 23			1							1				
Resource 24	1										1			
Resource 25					1							1		
Resource 26	1								1					1.5 hour lunch
Resource 27	1										1			
Resource 28		1						1						

Resource 29		1											1	
Resource 30		1												no lunch
Resource 31			1						1					
Resource 32			1									1		
Resource 33			1					1						
Resource 34			1					1						
Resource 35				1								1		
Resource 36					1					1				
Resource 37				1								1		
Resource 38					1						1			
Resource 39	1													no lunch
Resource 40				1							1			
Resource 41				1					1					
Resource 42					1						1			
Resource 43					1			1						
Resource 44					1			1						
Emp Agent				1			1							

Part Time Resources

	4:30-6:00	5:00-6:30	12:30-3:00	11:00-2:30	5:00-6:00	4:30-5:30	1:30-3:00	8:30-1:00
PT end day 1	1							
PT end day 2	1							
PT end day 3	1							
PT end day 4		1						
PT end day 5	1							
PT end day 6	1							
PT end day 7					1			
Brok Emp lunch			1					
Emp lunch 2				1				
Brok end day						1		
Brok lunch 2							1	
Brok PT Agent								1

Table A1: Resources Schedule

Name	Skills				
	266	258	264	263	262
Resource 1	1	1	1		
Resource 2	1	1	1	1	1
Resource 3	1	1	1		1
Resource 4	1	1	1		1
Resource 5	1	1	1		
Resource 6	1	1	1		1
Resource 7	1			1	1
Resource 8	1	1	1		
Resource 9	1		1	1	
Resource 10	1		1		
Resource 11	1	1	1	1	1
Resource 12	1	1	1		
Resource 13	1	1			1
Resource 14	1	1	1	1	
Resource 15	1	1	1	1	1
Resource 16	1	1	1		
Resource 17	1	1	1	1	
Resource 18	1	1	1		
Resource 19	1		1		
Resource 20	1	1	1	1	1
Resource 21	1	1	1		1
Resource 22	1	1	1		
Resource 23	1	1	1	1	
Resource 24	1	1	1		
Resource 25	1	1	1	1	1
Resource 26	1	1	1		
Resource 27	1	1	1		
Resource 28	1	1		1	1
Resource 29	1	1	1	1	1
Resource 30	1	1		1	
Resource 31	1		1	1	1
Resource 32	1	1	1	1	1
Resource 33	1		1	1	
Resource 34	1	1	1	1	
Resource 35	1		1		
Resource 36	1				
Resource 37	1		1		
Resource 38	1	1		1	
Resource 39	1	1		1	1
Resource 40	1		1		
Resource 41	1		1	1	1
Resource 42	1		1	1	1
Resource 43	1		1		

Resource 44	1		1	1	
Emp Agent					1
PT end day 1	1	1	1	1	1
PT end day 2	1	1	1	1	
PT end day 3	1	1	1		
PT end day 4	1	1	1		
PT end day 5	1	1	1	1	
PT end day 6	1	1	1		
PT end day 7	1	1	1		1
Brok Emp lunch				1	1
Emp lunch 2					1
Brok end day				1	
Brok lunch 2				1	
Brok PT Agent				1	

Table A2: Resource Skill Assignments




Figure A1: Customer ASA Basic Model Fit with Simulation Confidence Interval




Figure A2: Provider ASA Basic Model Fit with Simulation Confidence Interval




Figure A3: Provider B&E ASA Basic Model Fit




Figure A4: Broker ASA Basic Model Fit




Figure A5: Employee ASA Basic Model Fit

Interval	Provider				
	Customer	Provider	B&E	Broker	Employee
8:00	22%	14%	19%	20%	8%
8:30	6%	28%	20%	29%	38%
9:00	4%	20%	3%	13%	18%
9:30	13%	23%	26%	5%	7%
10:00	9%	6%	5%	11%	39%
10:30	1%	15%	15%	7%	21%
11:00	5%	3%	12%	1%	1%
11:30	2%	6%	4%	2%	7%
12:00	7%	20%	13%	5%	3%
12:30	11%	2%	5%	18%	13%
13:00	3%	2%	8%	11%	5%
13:30	9%	5%	8%	10%	1%
14:00	4%	16%	12%	2%	9%
14:30	1%	8%	5%	2%	3%
15:00	3%	15%	12%	5%	1%
15:30	4%	3%	23%	1%	12%
16:00	5%	25%	15%	10%	4%
16:30	8%	43%	43%	21%	6%
17:00	22%	38%	66%	5%	40%
17:30	36%	66%	43%	9%	50%
18:00	51%	22%	80%	NA	29%

Table A3: Percent Difference of Simulation Interval ASA from Data Interval ASA




Figure A6: Customer Monday SL Basic Model Fit




Figure A7: Customer Tuesday SL Basic Model Fit




Figure A8: Customer Friday SL Basic Model Fit




Figure A9: Provider Monday SL Basic Model Fit




Figure A9: Provider Tuesday Basic Model Fit




Figure A10: Provider Friday SL Basic Model Fit




Figure A11: Provider B&E Monday SL Basic Model Fit




Figure A12: Provider B&E Tuesday SL Basic Model Fit




Figure A13: Provider B&E Friday SL Basic Model Fit




Figure A14: Broker Monday SL Basic Model Fit




Figure A15: Broker Tuesday SL Basic Model Fit




Figure A16: Broker Friday SL Basic Model Fit




Figure A17: Employee Monday SL Basic Model Fit




Figure A18: Employee Tuesday SL Basic Model Fit




Figure A19: Employee Friday SL Basic Model Fit




Figure A20: Customer ASA with Cross-Training One Claims and One Calls Agent




Figure A21: Customer ASA with Cross-Training Two Claims and Two Calls Agents




Figure A22: Customer ASA with Cross-Training Three Claims and Three Calls Agents




Figure A23: Customer ASA with Cross-Training Four Claims and Four Calls Agents




Figure A24: Customer ASA with Cross-Training Five Claims and Five Calls Agents




Figure A25: Customer ASA with Cross-Training Six Claims and Six Calls Agents




Figure A26: Customer ASA with Cross-Training Seven Claims and Seven Calls Agents




Figure A28: Customer Monday SL with Cross-Training Two Claims and Two Calls Agents




Figure A29: Customer Monday SL with Cross-Training Four Claims and Four Calls Agents







Figure A30: Cross-Training with Two Additional Agents for Each Group; Provider, Provider B&E, Broker, and Employee Skills

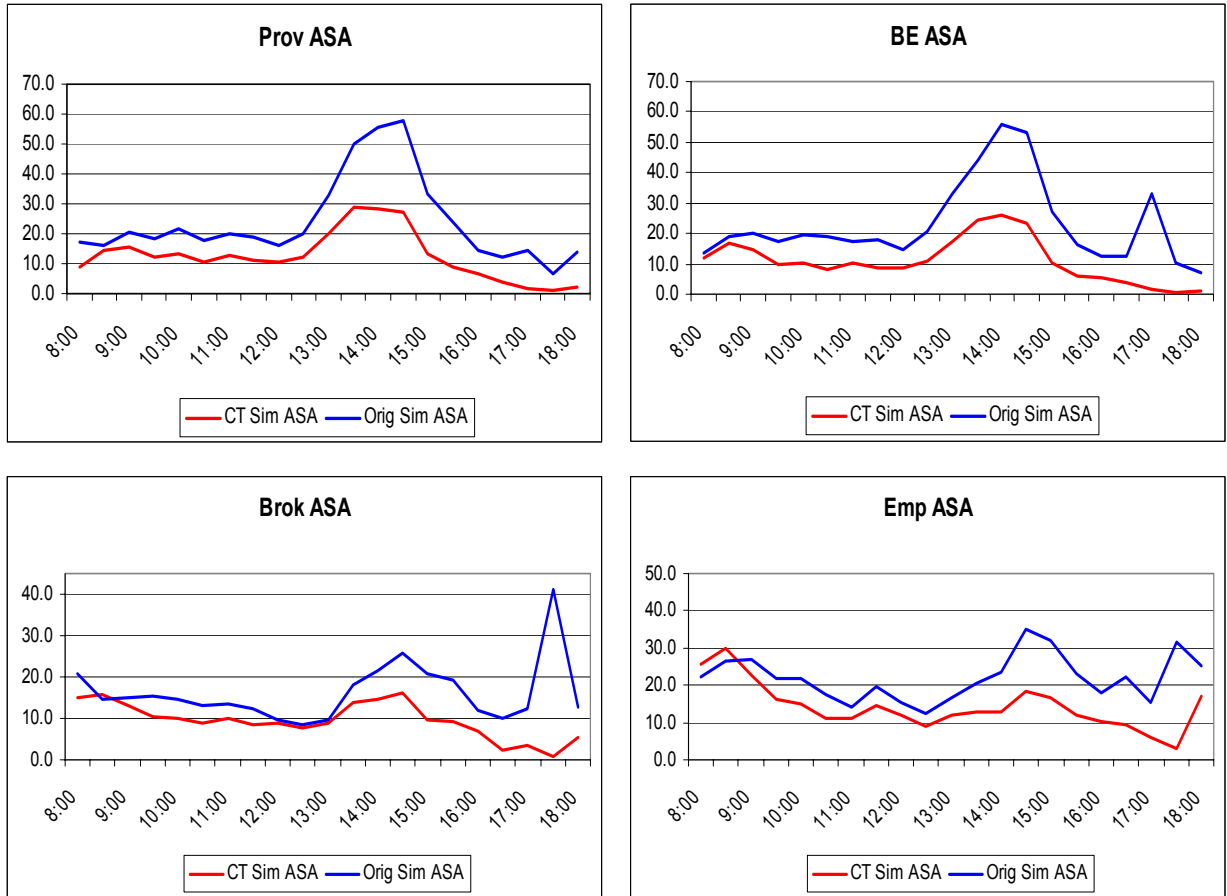


Figure A31: Cross-Training with Four Additional Agents for Each Group; Provider, Provider B&E, Broker, and Employee Skills




Figure A32: Period 3 PAN Confidence Intervals by Scenario




Figure A33: Period 7 PAN Confidence Intervals by Scenario




Figure A34: Period 10 PAN Confidence Intervals by Scenario




Figure A35: Period 13 PAN Confidence Intervals by Scenario




Figure A36: Period 17 PAN Confidence Intervals by Scenario