Managing the Driver Recruiting Pipeline in a Transportation Firm

Zahir Balaporia
Director, Advanced Planning and Decision Science
Schneider, Inc.
3101 S. Packerland Drive, Green Bay, WI 54313, USA
+1 (920) 592-2597
balaporiaz@schneider.com

Abstract

In the truckload transport business, the number of truck drivers is obviously a critical resource in order to achieve business goals. The industry has high turnover rates and hiring is a continual process. This project provided some key insights into the driver hiring process within the firm. The process was appropriately simplified and represented as a simple two stage stock adjustment model. One of the key insights was that while managers understood how hiring rates and termination rates resulted in growth or decline in the number of drivers, they were often unable to determine or agree on the number of open positions needed to meet the hiring goals that the budget called for. Part of the issue was that recruiting managers were unable to articulate to hiring managers what they "intuitively felt" was needed. The model and its accompanying explanation now provide the basis to approach driver hiring more prescriptively. It also provides the critical few metrics that need to be measured in order to make the decision on how many open positions are required to meet driver capacity targets.

Background

This project was undertaken as a semester long independent study towards completion of a System Dynamics degree from WPI. The author works at the firm represented below and is in the distance learning program. A WPI Adjunct Professor served as an advisor for the independent study.

About Schneider, Inc.

Schneider is the premier provider of truckload, logistics and intermodal services. Offering the broadest portfolio of services in the industry, Schneider solutions include Van Truckload, Dedicated, Regional, Bulk, Intermodal, Brokerage, Supply Chain Management, and Port Logistics services. Headquartered in Green Bay, Wis., Schneider has provided expert transportation and logistics solutions for 80 years. A $3.9 billion company, Schneider provides services throughout North America and China.

Executive Summary

This project provides some key insights into the driver capacity management process within Schneider's Intermodal line of business (the firm). In the truck related transport business, the number of truck drivers is a critical resource in order to achieve business goals. In this project, the hiring pipeline for drivers is simplified to only consider open positions that become drivers after a specified time that represents the time to hire and train. One of the key insights was that while managers understood the concepts of hiring rates and termination rates, and how the difference between the two resulted in growth or decline in the number of drivers, they were often unable to determine or agree on the number of open positions needed to meet the hiring goals that the budget called for. Part of the issue in convincing others was that recruiting managers were unable to articulate what they "intuitively felt" was needed. The model and its accompanying explanation now provide the basis to approach driver hiring more prescriptively. It also provides the critical few metrics that need to be measured in order to make the decision on how many open positions are required to meet driver capacity targets.

The model overview section in this paper provides a step-by-step explanation of the model to provide a base understanding of how the model maps to the actual business process. The output analysis section builds on this understanding by walking through a base scenario. It also provides sensitivity analysis for driver capacity related to five critical time constants in the process. A section on business applications outlines the interactions with functional managers and next steps on how this work is being used in business practice.

The project was originally focused on developing a system dynamics model to support the 3-year planning process for the firm. The model was designed to have six sectors, of which Driver Capacity was one. The project diverged from its original goal when the opportunity arose to take the Driver Capacity sector as a sub-model and apply it to a very real, critical and ongoing business opportunity. The model in this study was built with the dual purpose of helping with the near term opportunity and to be incorporated, with minimal changes, into the larger model.

Introduction

The model developed in this independent study is focused on the problem of hiring drivers in order to meet budgeted capacity levels at the right time. As one would expect, truck drivers are a critical resource in a truck transportation based business, and maintaining driver capacity is an industry wide challenge due to high turnover. The hiring process goes through various stages such as recruiting, interviewing, screening and training before an applicant can become a company driver. The model represents this process as a relatively simple two stage stock adjustment model and demonstrates the challenges of getting the right number of drivers hired to meet budgeted driver capacity levels.

This paper describes the process of deciding on the project scope, adjustments made along the way, the models developed and application at the company. A brief description of the business is also provided to provide context for readers unfamiliar with the industry and the firm.

Industry Background

At its core, the business is about managing networks and balancing flows. Economies of scale only help when the scale is balanced. For example, an increase in shipping volume from A to B could put the network out of balance if there isn't freight outbound from B to benefit from the new inflow of freight. Price is used to incentivize balance by charging higher prices for shipments outbound from markets that are net producers of freight (head haul markets), and providing discounts for shipping outbound from markets that are net consumers of freight (backhaul markets). Empty moves to reposition equipment from net consuming markets to net producing markets (backhaul lanes) are common. Profitability on a lane or for a customer must be evaluated within the context of the network. The business is capital intensive, competitive and very fragmented. Competition consists of a few very large fleets to many very small fleets, some as small as one truck independent contractors.

The firm represented in the model encompasses the continental United States with significant coverage in Canada and Mexico as well. This network is built around rail hubs, served by a fleet of over 1,000 drivers and trucks, and over 15,000 containers, which are used to move thousands of shipments per day. Dray shipment distances range from a few miles to about 250 miles, with over 80% of the drivers returning home each day. Driver turnover is a perennial issue in the industry, although it is comparatively lower in this firm due to a better quality of life for drivers who can be home with their families each day.

Scope

The original scope of this project was to build a model of the firm that would support the 3-year planning (3YP) process. While this was an ambitious goal, some initial work had already been done and an initial model to study the challenges in balancing service with operational costs was already built. The idea was to take the Service-Cost model and expand on it by incorporating some missing pieces that would cover the main asset stocks that make up the 3YP.

However, while working on the initial sector that focused on driver capacity, it became clear that there was an opportunity to build on this critical business process and derive value from just that piece of the model. We decided to put the 3YP model on the shelf temporarily and explore this sector more fully, but always with the intention of building it back into the 3YP model.

Driver Capacity Model Overview

The initial working version of this model was a single stock adjustment model that used a standard gap closing formulation between the actual stock of drivers and the budget. However, the behavior of the model was not acceptable as it did not use the hiring pipeline information appropriately. While exploring this issue, the project advisor suggested looking at Ch. 19 in Business Dynamics by John Sterman. The model in section 19.1, which refers to the labor supply chain, captured the essence of this problem almost perfectly, so the model for this project was adapted from it. A detailed description of the model is provided below.

Driver Hiring & Termination

This process is modeled as a basic stock and flow structure shown in Figure 1. The hiring inflow is determined by an upstream process that will be discussed below. Drivers leaving (i.e. Terms) are based on the estimated normal tenure of drivers. The adjustment to tenure is a custom function that modifies normal tenure based on dray efficiency. If dray efficiency is low, drivers are less productive and earn less, resulting in a

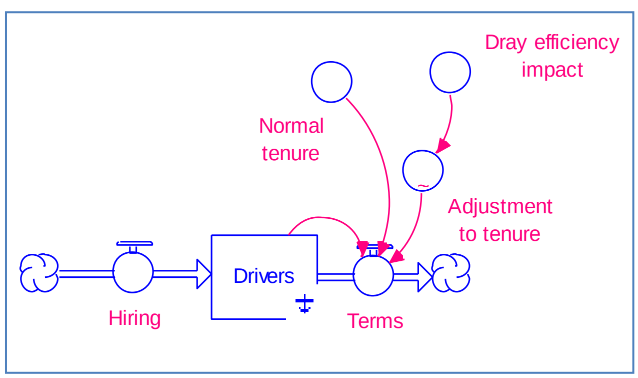

Figure 1 – Driver Stock Adjustment Structure

reduction in tenure and a higher term rate.

Open Positions for Hiring

The upstream process for hiring drivers is based on the number of open positions as shown in Figure 2.

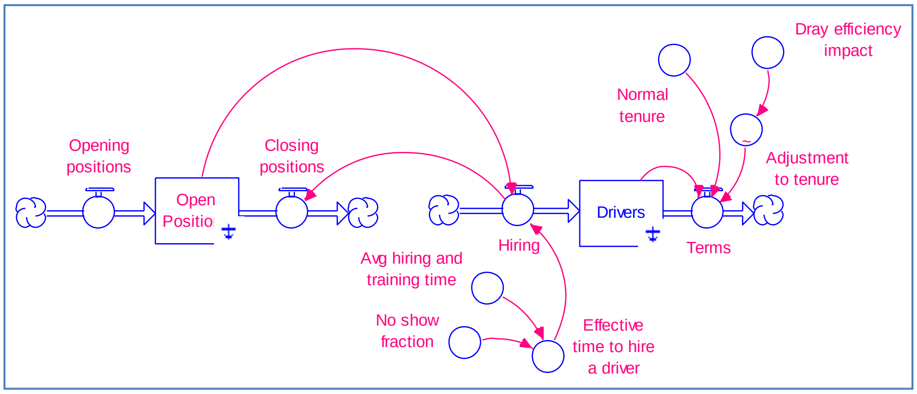

Figure 2 – Opening Positions for Hiring structure

Positions are opened in order to initiate any hiring, and positions are closed when drivers are hired. The stock of open positions determines the hiring rate which is based on the effective time to hire a driver. The effective time to hire a driver is based on the average hiring and training time and the no-show fraction. The no show fraction was initially modeled as an outflow that represented a reduction in open positions. This more complex formulation, when reduced to its essential impacts, was effectively an increase in the average time to hire and train. So the greatly simplified formulation is to divide the average time to hire and train by 1 minus the no show fraction to create the effective time to hire and train. Care is taken to not divide by zero.

Desired Hiring Rate

The stock of drivers is compared to the capacity budget to calculate the capacity gap (highlighted area in Figure 3 below). The gap estimation time is a decision delay that represents how management reacts to the capacity gap. Represented as a constant in this model, this time can vary based on the business climate. For example, if customer orders (not included in this model) were strong the gap estimation time would be shorter. In one sense it represents the speed at which the business would like to close the capacity cap.

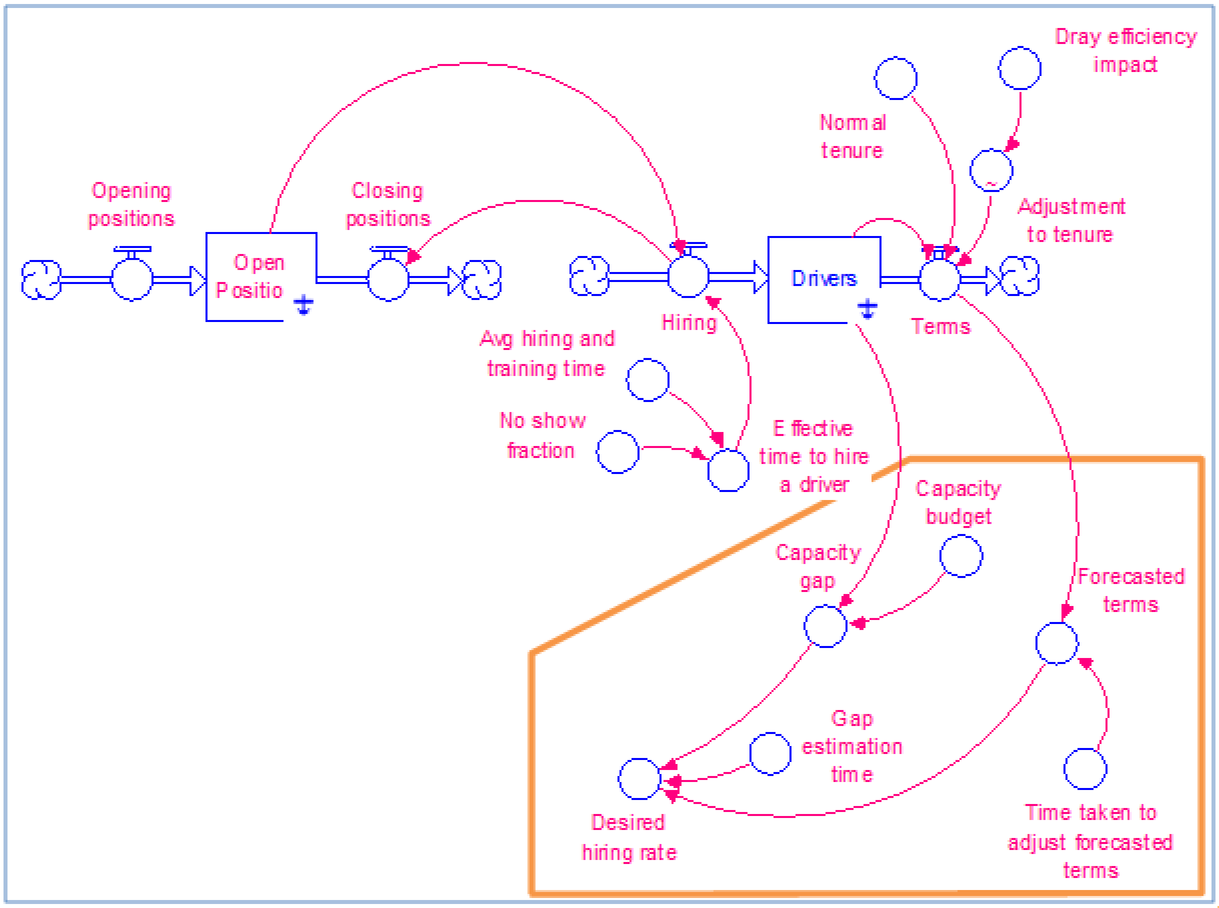

Figure 3 - Calculating the desired hiring rate

The other component of desired hiring rate is accounting for driver terms. Historical terms are forecasted based on a simple moving average where the time taken to adjust forecasted terms determines how quickly the forecast is adjusted.

From Desired Hiring Rate to Open Positions

The desired hiring rate needs to be converted into open positions. The first step is to calculate the desired open positions by multiplying the desired hiring rate by the effective time to hire a driver. This concept is best illustrated with a simple example. If the desired hiring rate is 10 drivers/week and the effective time to hire a driver is 4 weeks, then we need 40 open positions to generate the required 10 drivers/week. It is important to note that in this model desired open positions cannot be negative.

The adjustment for open positions takes the desired open positions and subtracts the already open positions to determine the net open positions that need to be created. This is divided by the time to

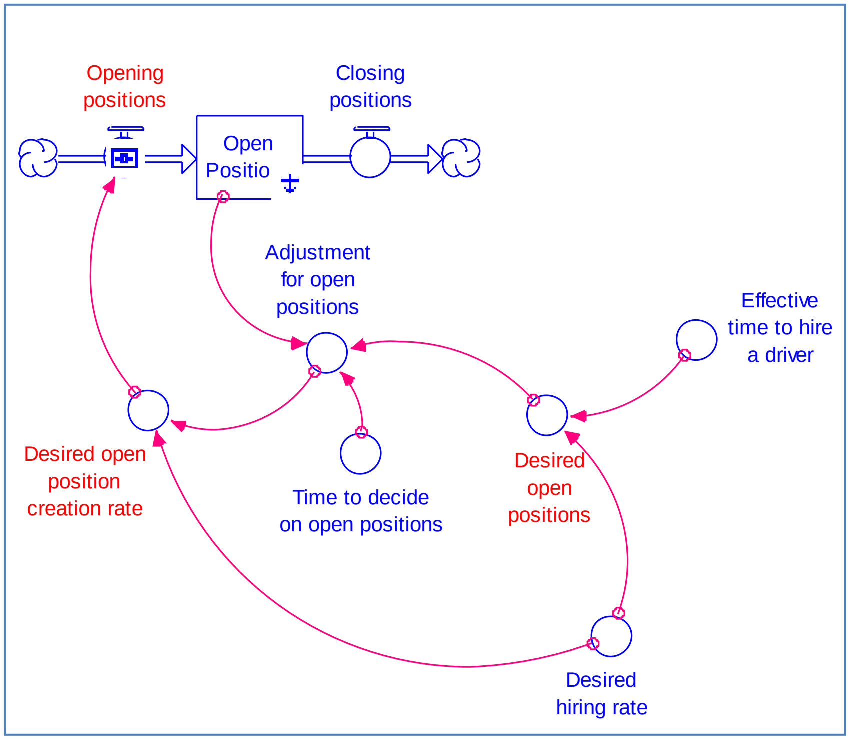

Figure 4 - Calculating the Opening Positions Rate

decide on open positions. This time constant represents the time it takes for the business to commit to the change in open positions, if a change is required.

The desired open position creation rate is simply the desired hiring rate plus the adjustment for open positions. The opening positions rate is this desired rate, except that it cannot represent a negative number as this model does not allow for a separate process for closing positions. In the actual business process this is obviously possible if driver capacity is significantly higher than the budget. However, that is a rare case given the reality of extremely high driver turnover rates and it is often dealt with by letting turnover solve for excess capacity, i.e. an increase in the time to decide on open positions and the gap estimation time.

The complete model is shown below in Figure 5.

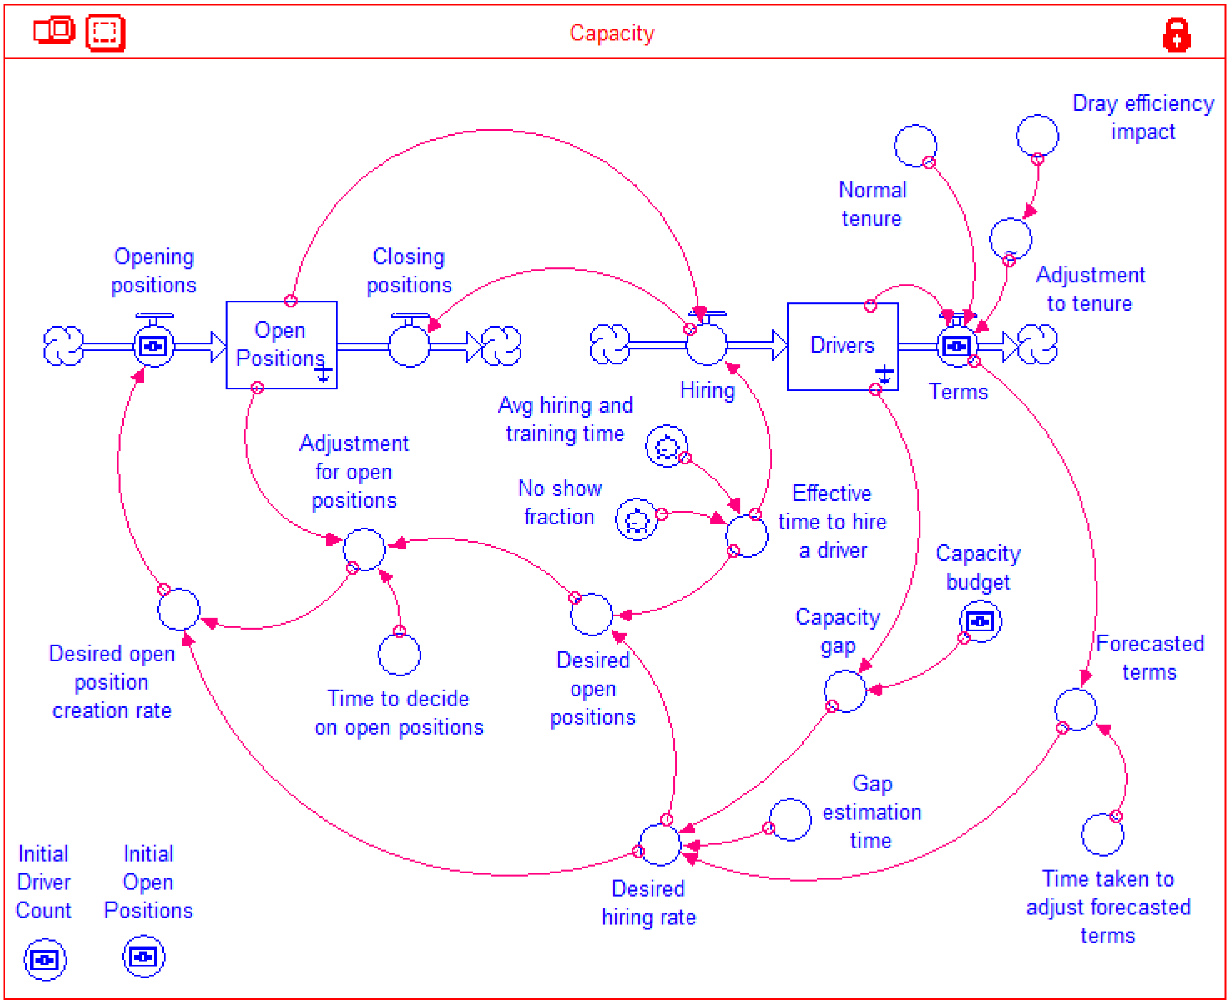

Figure 5 - Complete Model Diagram

Output Analysis

The base scenario uses an initial driver count of 1,250 drivers. The budgeted driver count starts at 1,250 and steps up to 1,350 in week 10. The normal tenure is set at 125 weeks which equates to approximately 10 drivers terminating per week. There is no adjustment to tenure (i.e. = 1). The average hiring and training time is set to 4 weeks with a no-show fraction of 0.06, which results in an effective time to hire a driver of approximate 4.26 weeks. Time taken to adjust forecasted terms is set to 1, and forecasted terms are modeled as a first order smooth with an initial value equal to terms. The gap estimation time is set to 2 weeks and the time to decide on open positions is set to 3 weeks.

In order for the model to start in steady state, the stock of open positions has to be set to a value so that the hiring rate is equal to the rate of terms i.e. 10 drivers/week. Since hiring is equal to open positions divided by the effective time to hire a driver, therefore open positions is set to effective time to hire a driver multiplied by terms i.e. 42.6.

The model starts in steady state. Then in week 10, the capacity budget steps up to 1,350 and creates a capacity gap of 100 drivers. With the gap estimation time of 2 weeks, the desired hiring rate due to the capacity gap is 50 drivers/week. The forecasted terms are 10 drivers/week, so the total desired hiring rate is 60 drivers/week. As a result, the desired open positions jumps to 255 drivers. With ~43 open positions and a time to decide on open positions of 3 weeks the adjustment for open positions is 71 drivers/week. This results in a desired open position creation rate of 131 drivers/week. The graph in Figure 6 shows the result of this initial step up in capacity budget at t=10.

Figure 6 - Initial Response to Step Up in Capacity Budget

With the increase in open positions, the hiring rate increases and the number of drivers start to increase. As soon as the number of drivers starts to increase, the capacity gap starts to decrease, which reduces the desired hiring rate and desired open positions. The gap between open positions and desired open positions falls until at ~t=12.15 weeks and the adjustment for open positions is 0. This is shown in Figure 7 below (t=12.13 weeks and adjustment for open positions is still 1, but approaching 0).

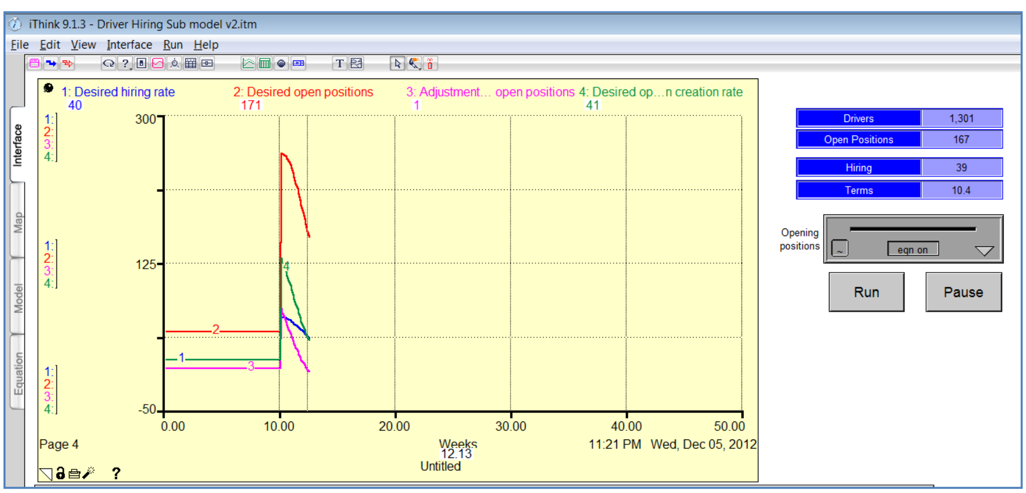

Figure 7 - Adjustment for Open Positions Approaching Zero

The number of drivers continues to grow because the hiring rate is above the term rate. But the number of open positions peaks at 168 and starts to drop as the desired open positions is less than the open positions, and the adjustment for open positions is now negative. At ~t=14.6 weeks, the number of drivers is equal to the capacity budget, but then continues to over shoot because the hiring rate is still higher than the term rate. This is shown in Figure 8 below.

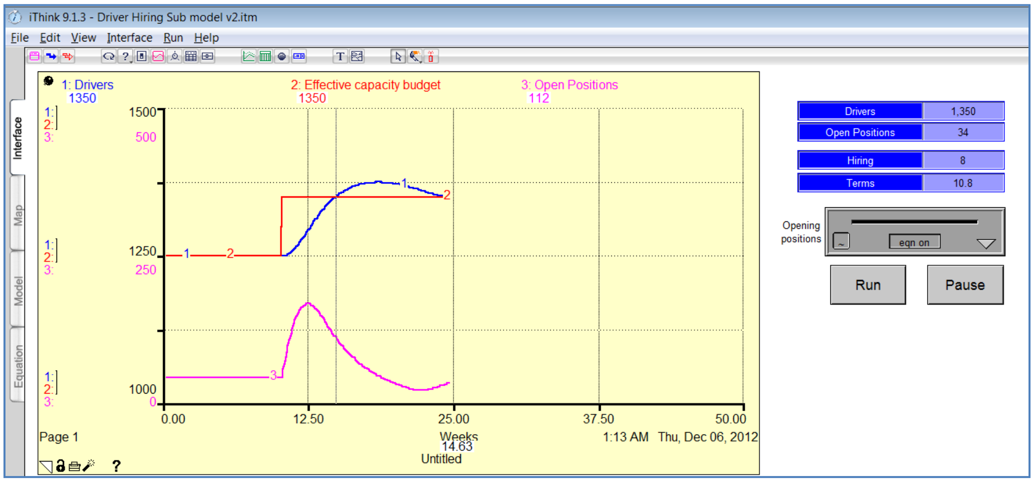

Figure 8 – Drivers and Open Positions

The overshoot is caused by a relatively aggressive gap adjustment time and the additional delay in opening positions. Open positions continue to drop while drivers exceed the budget. At ~t=18.25 weeks, drivers peak when open positions result in a hiring rate below the terms rate. Since the rate of opening positions cannot go negative, the only way that the number of open positions can drop is if the hiring rate is greater than the rate of opening positions. At the other end of the pipeline, due to the increase in drivers and driver tenure held constant, terms per week have increased by a small fraction. Drivers eventually drop to the budget level at ~t=25 weeks. However open positions have started to increase at ~t=21.5 weeks even though there are drivers in excess of the budget because, with the addition of forecasted terms, the desired hiring rate is positive again. This is shown in Figure 9 below.

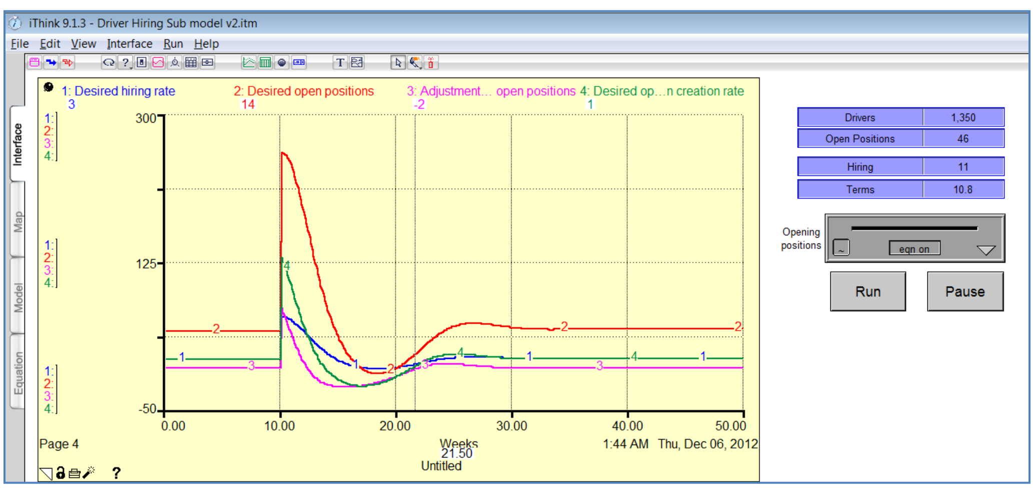

Figure 9 – Start Opening Positions Again

The figure above also shows the behavior of the model over the remaining time periods. The damped oscillation continues until ~t=30.5 weeks where the model settles into its new equilibrium with drivers at budget and open positions at 46, which results in a hiring rate that is equal to the term rate (Figure 10).

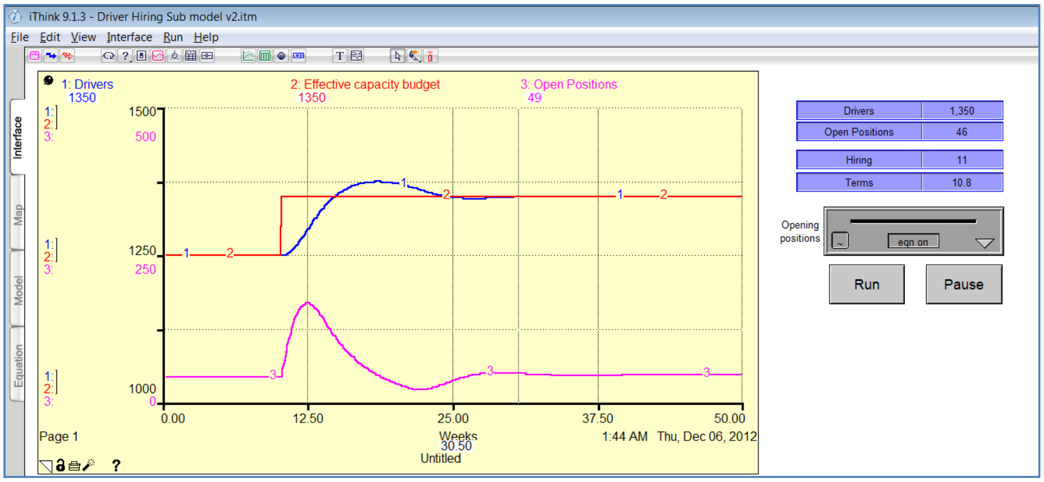

Figure 10 – Drivers Reach New Equilibrium at Budget

Sensitivity Analysis

In order to understand the sensitivity to some of the time constants in the model, the following graphs show the behavior over time of the stock of drivers when each constant is varied independently, holding all the others constant, at the values in the base case discussed above. Four values are used which produce four curves. The values are specified in the title of each sensitivity scenario and are the same for all scenarios.

Gap Estimation Time (1, 4, 7 and 10 weeks)

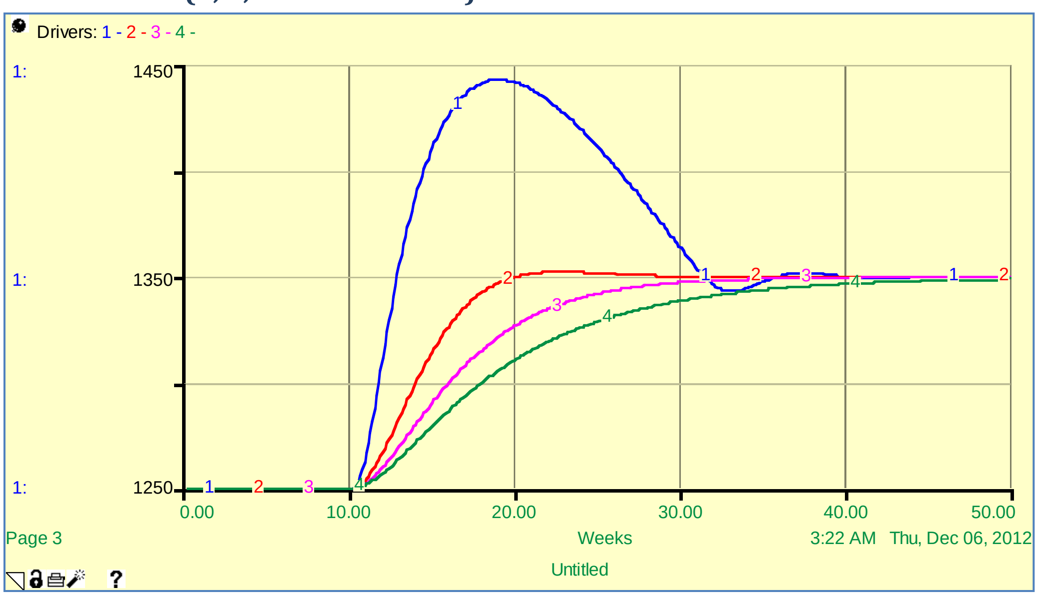

Figure 11 – Sensitivity of Drivers Stock to Gap Estimation Time

The sensitivity of the stock of drivers to gap estimation time is high. The lower the gap estimation time, the greater the oscillation and magnitude of the overshoot. The longer times eliminate any overshoot at the expense of taking longer to close the gap. So, longer times here create a damping effect.

Time to Decide on Open Positions (1, 4, 7 and 10 weeks)

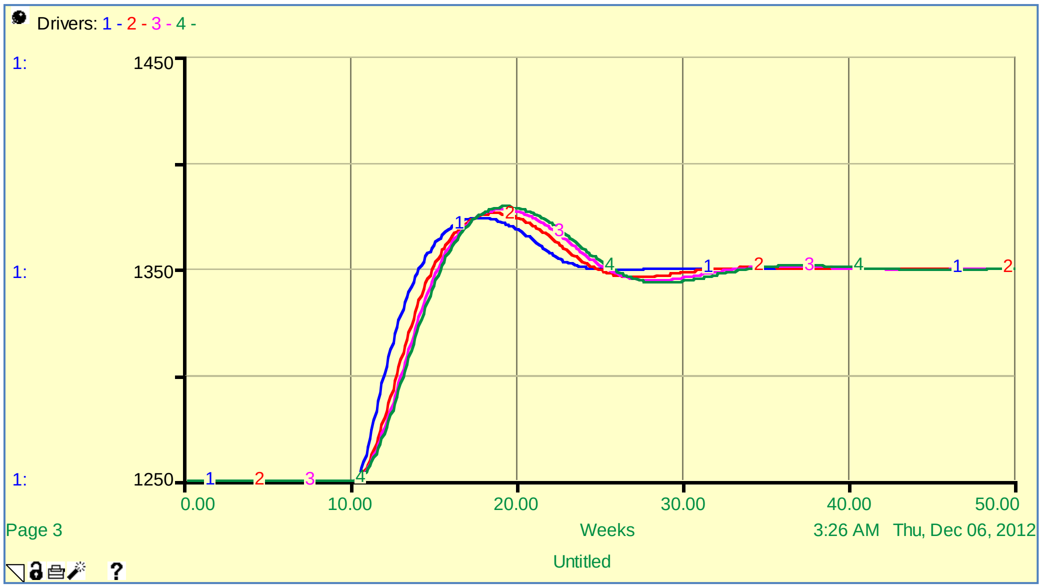

Figure 12 – Sensitivity of Drivers Stock to Time to Decide on Open Positions

The sensitivity to time to decide on open positions is lower. While it is consistent that the longer the time, the longer it takes to reach budgeted capacity, it is interesting to note that the longer times in this case create a slight increase in the overshoot and oscillation. So, longer times in this case have an amplifying effect.

Time Taken to Adjust Forecasted Terms (1, 4, 7 and 10 weeks)

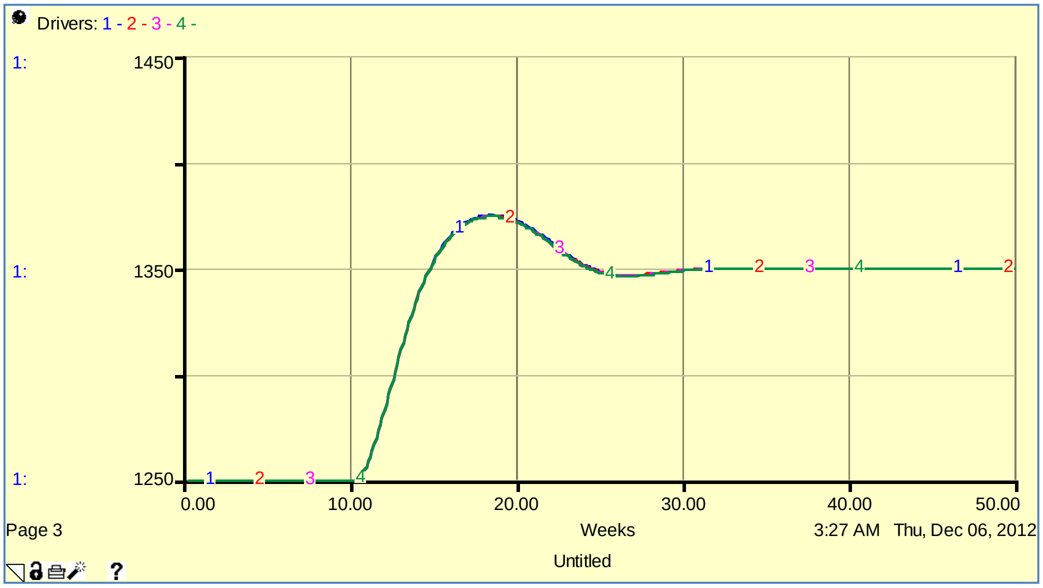

Figure 13 - Sensitivity of Drivers Stock to Time Taken to Adjust Forecasted Terms

The time taken to adjust forecasted terms has very little impact on the rate at which the drivers stock adjusts to budgeted levels. This is due to the fact that in the base scenario, the term rate does not change significantly, so longer delays in adjusting forecasted terms have no significant impact on forecasted terms.

Avg Hiring and Training Time (1, 4, 7 and 10 weeks)

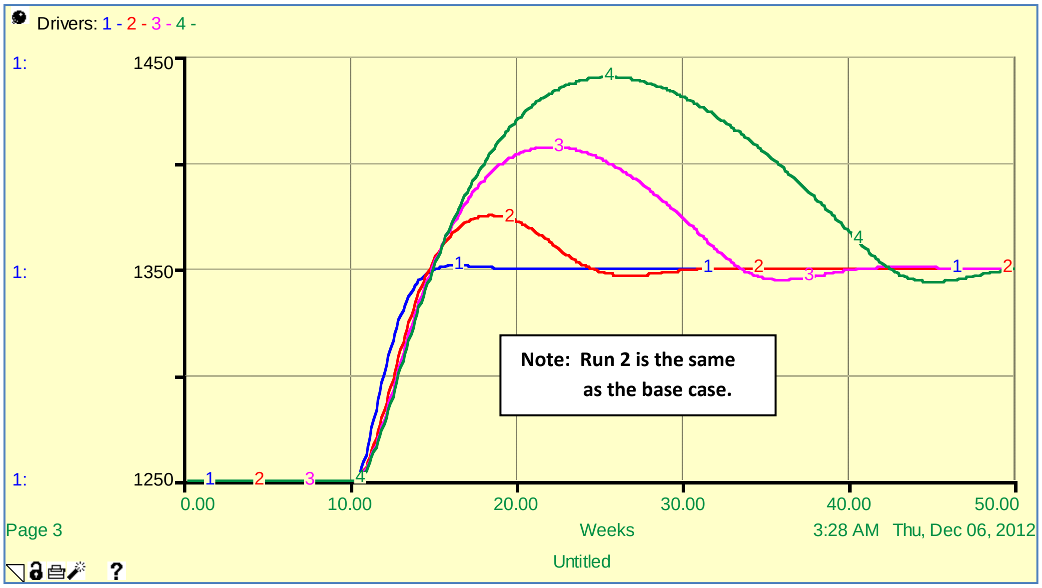

Figure 14 - Sensitivity of Drivers Stock to Average Hiring and Training Time

The average hiring and training time clearly has a significant impact. As expected, the shorter the time, the faster the stock reaches steady state. The longer the time, the greater the overshoot and therefore the longer time to reach steady state. Interestingly the longer hiring and training times do not delay the time to reaching the budget which is close to t=15 weeks in all four cases. This is because the gap adjustment process opens significantly more positions to compensate for the hiring and training time, which in turn results in the large overshoot.

Business Opportunity

This model addresses a critical process within the company. The number of drivers is one of the key asset stocks required to generate revenue. Furthermore, one of the constant conversations within the company, and the industry at large, is turnover rates for drivers and the possibilities of long term driver shortages due to changing demographics. While the larger driver turnover problem is beyond the scope of this effort, the fundamental impact of this larger issue is captured in the concept of effective time to hire. In general, the concept of hiring rates and term rates are well appreciated within the firm, and the fact that growth only occurs when the hiring rate exceeds the termination rate is well understood. Weekly status meetings show charts of hires last week versus terms last week, and the weekly trend of each rate over the last many weeks are reviewed. However, when the desired number of drivers is greater than the actual, it always seems to take longer to close the gap than people expect. There seemed to be an opportunity to improve our mental models of how this system behaved. So this model was used with the team responsible for driver capacity to explore our collective understanding of how this system works.

To do this, an early version of the model was used to socialize the concepts that the model addressed and to provide a brief introduction to system dynamics models and modeling. The team seemed to get comfortable with the idea of flows representing hiring and termination rates, and the stocks representing the number of drivers. The initial meetings were more helpful in cleaning up some terminology so that the model mapped to the business process more naturally. For example, the initial model was built using the concept of a "job requisition" or "req". However, a "req" represented a type of job, whereas "positions" were the actual number of openings within a "req". So the model terminology was changed to "positions" which drivers are hired against, and then closed as drivers are hired.

But beyond terminology, even in this very initial meeting, the conversation got very interesting as we looked at a model run very similar to the base case. We noticed the number of positions the model suggested opening based on the formulation that accounts for the lead time to hire and train. When I suggested that the reason we have a tendency to lag in our hiring needs is because we don't open enough positions, both managers responded with "that's what we've been saying, but they won't listen to us." Further probing led me to understand the issue as an incomplete understanding of how open positions, combined with the time to hire and train, resulted in a hiring rate. The simplest example was one market where it took a long time to hire e.g. 10 weeks. If turnover was just 1 driver per week, we would need 10 open positions on average to maintain current capacity. However, when the recruiting team worked with operations to discuss hiring needs, operations would push back on opening 10 positions because "we only need 1." A simple suggestion to avoid this problem was to not inform operations of the number of positions being opened as it did not impact their work processes directly in any way. Given that this first meeting had generated good input and interest, a number of follow-up meetings were scheduled. The model is used today in cases where there is disagreement between the recruiting manager and the hiring manager on the number of open positions needed.

Summary

The original scope for this independent study project was to build a model to support 3-year planning (3YP) for the firm. However, when the opportunity to take the Driver Hiring sector of the 3YP model and apply it to an existing issue presented itself, the direction of this study was modified to focus on the driver hiring process. However, the desire to build the 3YP model still exists, and this driver hiring model has already been incorporated into the 3YP model.

The model used in this study is an adaptation of the labor supply chain model in Business Dynamics by John Sterman. The model in Section 19.1 was used with minimal modification. It represents the hiring process as a second order pipeline adjustment model, and maps well to the current business processes used to manage driver capacity. The model shows the behavior over time of the stock of drivers based on parameters selected for various time delays in the decision making process, and the process used to determine the number of initial and weekly positions to open in order to hire drivers and achieve the budget for the required number of drivers.

Initial model runs showed a much larger number of open positions than the actual business process tended to produce, but any lower number of open positions would cause us to be always short of our capacity budget. When this was shared with the managers responsible for recruiting, they responded enthusiastically that they always felt they needed to open more positions but the hiring managers would argue for fewer positions, and the recruiting managers had no convincing argument to respond with. The gap in understanding seemed to come down to the fact that open positions were being viewed as imminent hires rather than averaged across the time to hire and train. The managers responsible for the recruiting and hiring pipeline now have a better understanding of the dynamics of their process.

Appendix A

Complete Model Diagram

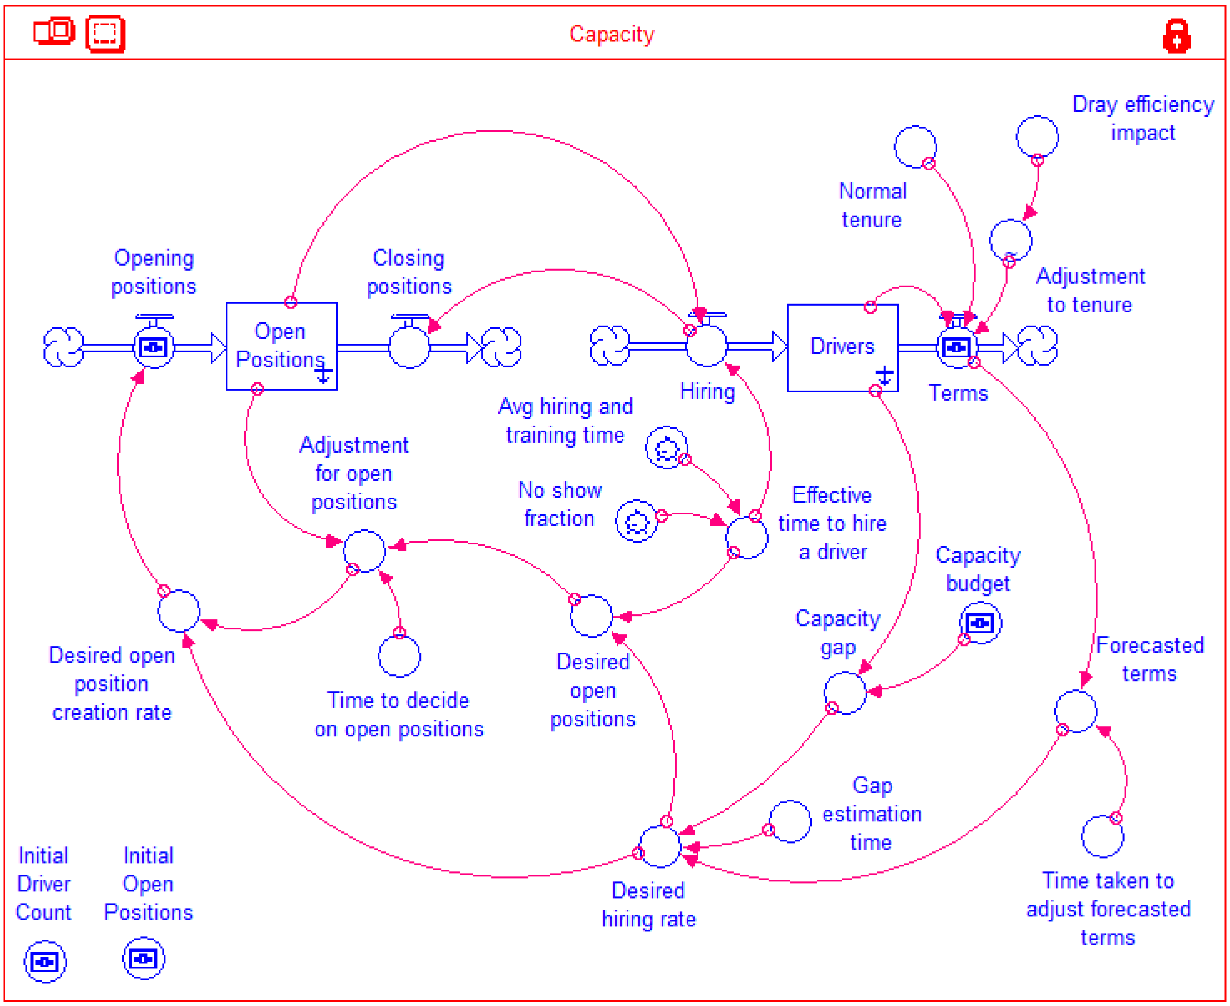

Appendix B

User Interface in iThink