Operational Analysis of Global Staff Mobility

Zachary Mulder
Department of Industrial Engineering and Operations Research
University of California, Berkeley
May 15th, 2015

Introduction
Employee mobility is a key issues for many organizations when faced with the challenge of rotating staff through worldwide locations. Whether the organization is an airliner who is assigning home bases for crews, military units deciding stations for each troop, or international entities staffing offices throughout the world, designing a fair and equitable matching system is a common business problem. The client who engaged our team in this project falls into the last category, and has asked to not be named.

This project models a proposed relocation system and gives an operational analysis of the trends in such a system. Additionally, the simulation tool created allows the client to see how the trends in the system are affected by changing the structure of the relocation system or by changing a number of different input parameters. Extended work on this project also provides a prototype of a functional tool that could be used by the client to make location assignments as needed. Overall, the aim of this project was to give the client a picture of how the proposed system would perform and more importantly, to make them aware of any pitfalls or issues that they might have to address before adopting a new relocation scheme.

Model Design
The model created in this project comes from our own insights applied to specifications from the client. In the proposed system, employees, who are known as International Professionals (IPs), each have a functional group (FG) that corresponds to the duties they are capable of performing. This means that any IP within a given functional group is interchangeable, and serves as the basis of rotation. There are IPs who are designated as non-rotational; however, we assume that all IPs modeled in our system are eligible for relocation. Each location (also known as a station) can have multiple positions available (known as posts). Stations also have a hardship classification according to prevailing living conditions as established by the International Civil Service Commission (ICSC will be used to denote a station's hardship classification).

The client stipulated that there is an upper limit to the amount of time that an IP can spend at each station which is a function of the station's ICSC. This upper limit of time is known as the tour of duty (TOD) or the standard duration of assignment (SDA). Furthermore, there is a minimum amount of time that an IP can spend at each station, which is known as the time in post (TIP). TIP is approximately half of the SDA, and the values for each ICSC can be found in table 1.

Each IP also has an IP priority that signifies their order in the matching process. No specific structure relating to IP priority is assumed, except that a lower number corresponds to an IP with a higher priority (i.e. an IP with a priority of one has a higher priority that an IP with priority of five). Furthermore, each station also has a priority representing which stations are more important to the organization to staff.

Table 1. ICSC Classifications, IDs, TIPs, and SDAs

ICSC Classification	ID	Minimum Time (TIP)	Maximum Time (SDA)
H, A, and B Duty Stations	1	3 Years	6 Years
C Duty Stations	2	3 Years	5 Years
D Duty Stations	3	2 Years	4 Years
E Duty Stations	4	2 Years	3 Years
Non-family Duty Stations	5	1 Year	2 Years

Based on the above definitions, each matching cycle follows a general structure:

1. Retirement: IPs who are retiring are removed from the system and their posts are added to the list of posts available for IPs to fill (known as the post list).

2. SDA Voluntary Match: IPs at their upper limit of time at a station (SDA) have their posts added to the post list and are allowed to make a voluntary match of their own choice from the post list.

3. TIP Voluntary Match: IPs who are above their minimum required time (TIP) are allowed to make a voluntary match of their own choice from the remaining post list.

4. SDA Final Match: IPs who are at their upper time limit and have not accepted a voluntary position are assigned a single match. If they reject the match, they are released from the company and the post remains open.

5. New Hires: Each post that is still open is either filled by a new hirer with a certain probability or is left open to be on the post list for the next cycle.

Assumptions

Accepting a post is stochastic and stationary

In this model, it is assumed that the probability of an IP accepting a post is a Bernoulli trial with the probability of success as a function of both their current ICSC and the ICSC of the post. That is, the probability of accepting a post is a from/to matrix corresponding to the IP's current ICSC and the ICSC of the offered post. This probability is assumed to be stationary throughout the course of the simulation through all stages of matching for both IPs at their TIP and at their SDA.

Repeating Stations

An IP can remain at the station where they are currently located one time before having to move to a new station. This is a stipulation provided by the client, and allows IPs to have some long term stability on their living situation while still maintaining the spirit of the rotation system. As the current system allows IPs to remain in their current post for an indefinite amount of time, this limit is a compromise between forced relocation every time an IP reaches their SDA and the current system.

IP Priority

In this model, an IPs priority is assumed to signify how close the IP is to reaching their SDA. An IP who is at their TIP has the lowest priority while IPs at their SDA have the highest priority. This lubricates the system by giving IPs who need to relocate first choice of all available posts and naturally incentivizes them to select a new location. This is accomplished by increasing an IPs priority every time they are above their TIP but do not make a match. However, such a system can also be problematic as IPs may wait a few cycles before making a voluntary match in order to get a higher priority and have access to more desirable posts. Thus, at each cycle of the model, the user can input any priority value for an IP as long as it follows the scheme of one being the highest priority. That way, the client can test multiple priority systems and determine which they believe is best.

Station Priority

Station priority is not altered throughout the simulation; however, like IP priority, it can be changed after each cycle based on the changing needs of the client. This would allow the client to test emergency situations and monitor how well their proposed relocation system could adapt to changes in staffing requirements.

IPs at SDA During Final Match
IPs who are at their SDA who do not voluntarily choose a new post in a previous stage are assigned a match and are not allowed to return to their previous post as often it already has been reallocated. If they do not accept their match, then they leave the company. Early iterations of this model examined what happens if IPs are allowed to not leave their current post meaning relocation is voluntary. The overwhelming result was IPs would continually transfer until they reached the highest ICSC and then would very rarely leave, often remaining at their current post until they retired. This result is exactly what the client is attempting to avoid, which is why IPs are forced to match after the voluntary selection stage.

Vacated TIP Posts
When an IP who is above their TIP but not at their SDA accepts a post, there current assignment is now vacated and must be added to the list of available posts. This model allows the post to be added in one of three places: immediately meaning another TIP may accept the vacated post, after all TIPS match meaning an IP at SDA who is in their final match stage may fill the open post, or after SDA final match meaning the post may be filled by a new hire. The second option may help reduce the number of failures, or the number of times when there is no match available for an IP in the third matching stage; however, with a sufficiently large functional groups failures are very rare.

Design Decisions
One of the key issues in this project was creating a system that was both equitable and well received by IPs. My main concern was that IPs would be frustrated if they did not have any control over their relocation or career path and would ultimately reject a system where they had very little control. Thus, this model seeks to give the IPs control over their own careers through the first two voluntary selection cycles, which have subtle features that mirror how this system would be implemented in the real world.

During the first two stages of matching, SDA voluntary match and TIP voluntary match, the stage begins with the IP with the highest priority being offered the station not with the highest station priority, but with the highest valued ICSC. The IP is then offered posts with ICSCs in decreasing value until they accept a match or the end of the list is reached, in which case the IP with the second highest priority is then offered all available posts by ICSC. This design mimics a real world system where the IPs would have access to a catalogue of available posts, and IPs with the highest priority get to pick their desired post first, similar to enrolling in classes at UC Berkeley through Telebears. This voluntary matching phase assumes that the primary factor in accepting a post is ICSC, which seems reasonable after many discussions with employees who work for the client as lower value ICSC stations are in remote locations.

However, to balance out the preference towards IPs in the first voluntary stages, the last stage, SDA final match, is favored towards the stations. In this stage, this post list is resorted by station priority, and the station with the highest priority is offered the IP with the highest priority who matches the post requirements. If the IP does not accept, the IP is removed from the system and IPs continue to be offered to the top post by decreasing IP priority until one accepts. This shift signifies that in the third stage, the system is attempting to fill the posts that are most important to the client first. IPs had their chance to choose their post in the first two stages, thus this stage removes any IPs who do not accept their match.

Inputs

The main input into the model is the current state of the system, known as the CAS. The CAS contains all information relating to each IP currently in the system. This could also be thought of as the initial conditions for testing. The information needed for each IP is listed in table 2. Additionally, the current open post list is needed which contains the same information as an IP with the IP ID set to 0, indicating that the post is open.

The other two main inputs are the probability of a post remaining open next cycle, which is used in the fifth stage of the matching cycle, new hires, and the from/to probability matrix representing the probability of an IP accepting an offered post as a Bernoulli trial.

Creating an Excel Model Using VBA

One of the main considerations in this project was what software to use to create the simulation tool. Generally, simulation requires software that is specifically designed for that purpose, such as SIGMA or Simio, since these programs already account for many important aspects of simulation such as resource contention and the flow of time.

However, with this project, we observed that even if the matching cycle was a continuous process throughout the year, we could model the system by having all matching take place at the same time. Whether the matching is performed every month or every year, if the time units in the simulation where merely cycles, then the time component of this model would drop out, and all that would be left is resource contention, or multiple IPs vying for the same post. Thus, as long as the matching was handled appropriately, any programming language could be used to create the model.

This opened up a number of possibilities for development, and ultimately we chose to use an Excel front end with a back end simulation using VBA. Excel is a familiar format to many employees, and also provides complex graphing capabilities already integrated into the program. Since the simulation returns the CAS after each cycle, this allows an employee who is already proficient in Excel to perform additional analysis with the given data. Additionally, we chose VBA as the back end simulation language because it integrates very easily with Excel. While VBA is not as rich nor as fast as other languages, we found that ultimately a cycle with 1000 IPs took only a few seconds, and the speed tradeoff for a more intuitive interface was appropriate. The simulation used approximately 2000 lines of code, and the high level functions comprising each cycle is shown in figure 1 on the next page.

Table 2. IP Information

This table contains all the necessary information for each IP which is input as the initial condition. Note, the open post list also requires the same information with IP ID set to zero to indicate the post is open

IP ID	IP Priority	FG	Post FG	Station
A unique number corresponding to an IP	The IP's priority with one being the highest	The IP's functional group, assumed to be the same for all IPs	The functional group required by the post, assumed to be equal to FG	The current station of the IP (or the station of the post).
Cycle Retire	**Cycle Hit TIP**	**Cycle Hit SDA**	**Prev Station**	**ICSC**
The cycle number that the IP will retire	The cycle the IP is at their TIP	The cycle the IP is at their SDA	The IPs previous station	The ICSC classification of the station

```
'1. Create all needed lists by removing IPs who are retiring, IPs at their SDA, and IPs above their TIP.
Call setLists

'2. Give IPs at SDA first pick
Call SDAFirstPick

'3. Next give any IPs above TIP (who is not at SDA) a chance to pick
Call TIPSecondPick

'4. Force IPs to match ONLY IF SPECIFIED.
'Otherwise, allow the IP at SDA to stay in their post if it is not taken
If (forceMatch) Then
    Call SDAThirdAssign
End If

'5. Check all remaining open posts and determine which are filled by new hires
Call fillOpenPosts

'6-8 Are specific analysis conducted each cycle
'6. Output probability matrix or flow rate of transfers
Call calcTransfers

'7. Output count of number of times IPs at SDA did not have a match
Call calcSDANoMatch

'8. Output z-index showing average IP time in a post per ICSC
Call calcZIndex
```

Figure 1. VBA Code Outline

The high level functions called during each cycle. Note that steps 6 – 8 are added statistical outputs demonstrating how the code can be easily adapted to changing future needs.

Interface

As stated above, one of the key reasons we used Excel was to create a friendly an intuitive interface for the user. A screenshot of the interface along with the simulation commands are shown in figure 2. There are four options for the user to choose after each cycle is completed: cycle, which progresses the simulation forward one cycle, batch cycle which performs multiple cycles in rapid succession, reset which resets the spreadsheet back to the initial conditions specified at the beginning of the run (i.e. deletes all cycles except the first), and settings which allows the user to change the input settings. Note that updating settings will only have an effect if done before the first cycle is simulated (i.e. immediately after resetting the simulation or immediately after opening the tool before cycle is clicked).

Figure 2. Simulation Tool Screenshot

The initial conditions entered into the system is the information found in cycle 1. The open post list is entered on the worksheet "OpenPost." The user can perform four actions: cycle, batch cycle, reset, or input settings. Simulation results are then output for each cycle on the main worksheet.

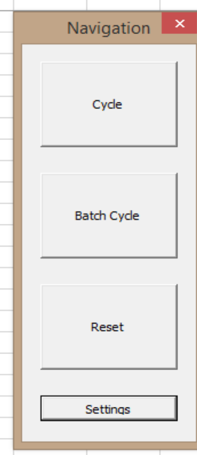

Cycle	IP ID	IP Priority	FG	Post (FG)	Station	Station Priority	Cycle Retire	Cycle hit TIP	Cycle hit TOD	Prev Station	ICSC	Cycle Hired
1	1	3	1	1	1	1	6	1	4	2	1	0
	2	4	1	1	3	5	11	2	5	5	2	0
	3	5	1	1	4	3	2	1	2	4	5	0
	4	5	1	1	2	2	7	0	2	4	3	0
	5	1	1	1	5	4	18	0	2	4	4	0
	6	3	1	1	5	4	10	1	3	5	4	0
	7	3	1	1	1	1	6	1	4	4	1	0
	8	3	1	1	3	5	11	1	4	4	2	0
	9	2	1	1	2	2	14	1	3	1	3	0
	10	5	1	1	2	2	8	2	4	1	3	0
	11	3	1	1	2	2	12	0	2	2	3	0
	12	2	1	1	2	2	3	0	2	3	3	0
	13	1	1	1	2	2	10	0	2	1	3	0
	14	5	1	1	5	4	5	0	2	1	4	0
	15	5	1	1	5	4	2	1	3	1	4	0
	16	5	1	1	5	4	6	1	3	4	4	0
	17	3	1	1	5	4	10	1	3	3	4	0
	18	1	1	1	5	4	3	0	2	5	4	0
	19	1	1	1	3	5	16	1	4	3	2	0
	20	4	1	1	4	3	3	1	2	2	5	0
2	1	3	1	1	1	1	6	1	4	2	1	0
	2	4	1	1	3	5	11	2	5	5	2	0
	3	5	1	1	4	3	2	1	2	4	5	0
	4	5	1	1	2	2	7	0	2	4	3	0
	5	1	1	1	5	4	18	0	2	4	4	0
	6	3	1	1	5	4	10	1	3	5	4	0
	7	3	1	1	1	1	6	1	4	4	1	0

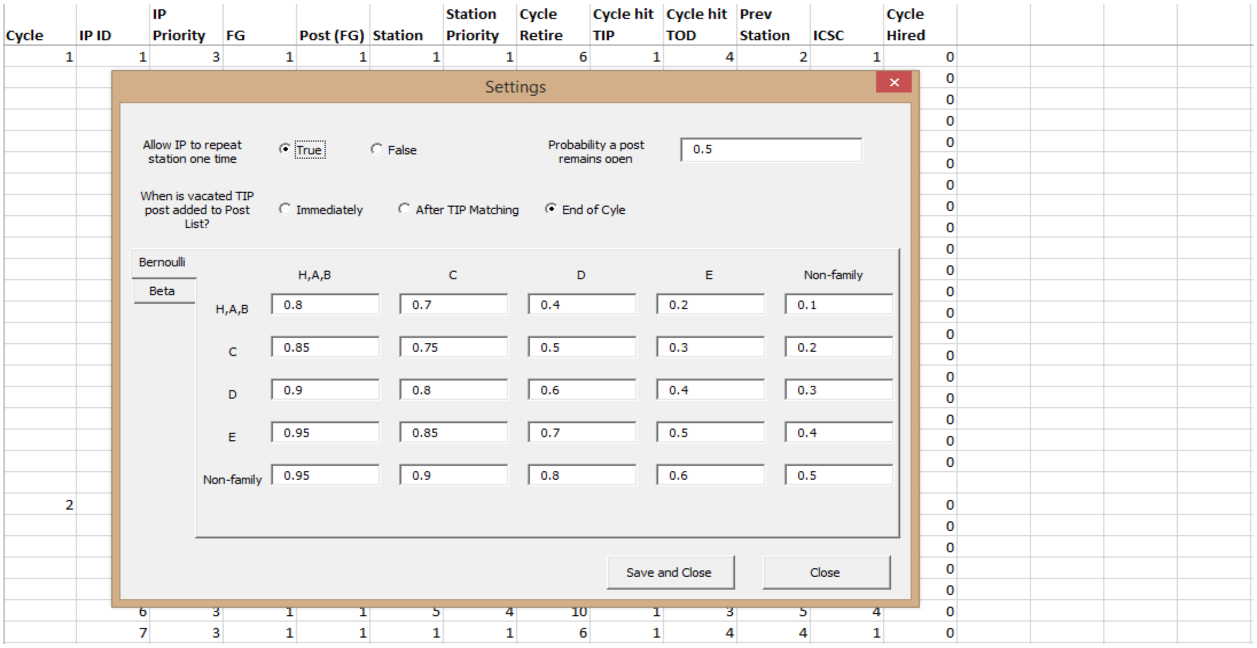

Figure 3. Settings Pane

This window allows the user to input many different settings used in the simulation. Note that using the Beta distribution page will add randomness to the probability of an IP accepting a position

In the settings pane, a number of different options can be modified. The first choice is if IPs are allowed to repeat their current station one time before being forced to move to a new station. Next is the probability that a post remains open until the next cycle. This represents how likely it is that a new IP can be hired to accommodate any remaining posts that are not filled at the end of each cycle. Additionally, there is an option of specifying when vacated posts by IPs at their TIP are added to the list. The options were detailed previously in the report. Finally, there is an input area for the probability that an IP accepts a post based on the ICSC of their current post and the post being offered. An image of the settings pane can be found in figure 3.

There are two options for this section: a Bernoulli distribution or a Beta distribution. If the Bernoulli option is used, then the input value represents the probability that the post will be selected. If the Beta distribution is used, then each entry is a dropdown box with three options: attractive, neutral, or unattractive. Each of these three options corresponds to sampling from a different distribution to define the probability of acceptance. Selecting attractive samples from a right triangle distribution between zero and one with a mode of one, selecting neutral samples from a uniform random variable between zero and one, and unattractive samples from a right triangle distribution between zero and one with a mode of zero. Essentially, attractive transfers are more likely to have a high probability of acceptance, while unattractive transfers are more likely to have a low probability of acceptance. This randomness in probability mirrors the client's uncertainty in which pairings will be attractive, and allows them to run multiple simulations with varying values to determine how sensitive their system is to transfer probabilities.

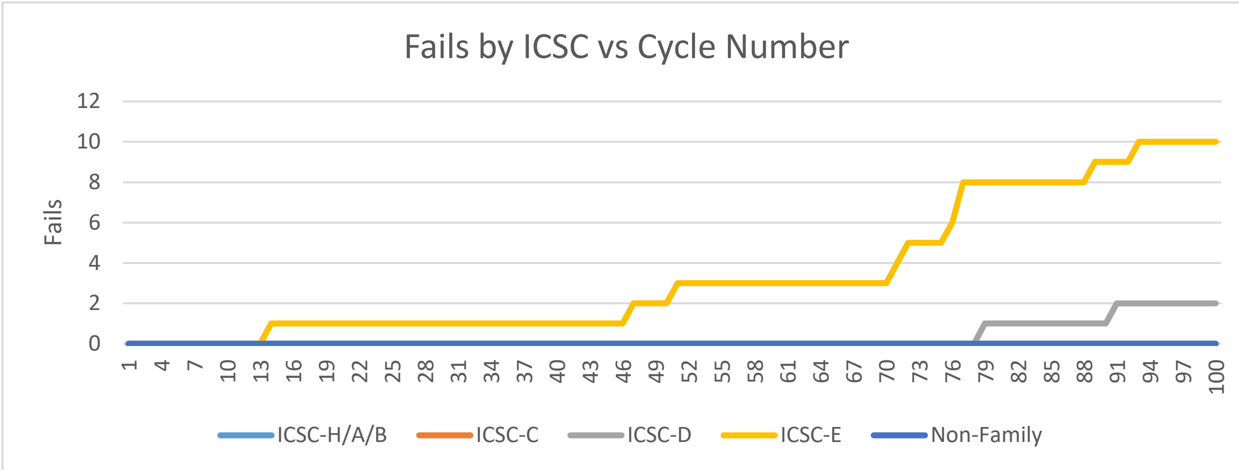

Figure 4. Cumulative Failures by Cycle
A failure is defined as an IP leaving the system when no match is offered to them during the last matching stage.

Outputs

There are three main outputs generated from the simulation which are used as metrics to judge system performance. The first counts the cumulative number of fails that occur throughout the span of the simulation and plots the value versus time. This metric indicates when the system breaks down and IPs who are at their SDA are not offered any posts during the third matching stage, SDA final match, because there are no posts available that they are eligible to fill. The number of fails decreases rapidly as functional group size and the number of stations increase which is reasonable since creating more total posts means it is less likely for an IP to not find at least one match. An example graph is shown in figure 4.

Another output generated by the simulation is a graph representing the transfer rates from each ICSC to another. This graph indicates how IPs are generally moving through the system starting with where new IPs are hired traced through to retirement. An example output is shown in figure 5.

Figure 5. Transfer Rates
Notice that often IPs transfer to a higher valued ICSC. This is due to the voluntary matching phases favoring the IPs preference, which is assumed to be ICSC, over the client preference, which is assumed to be station priority.

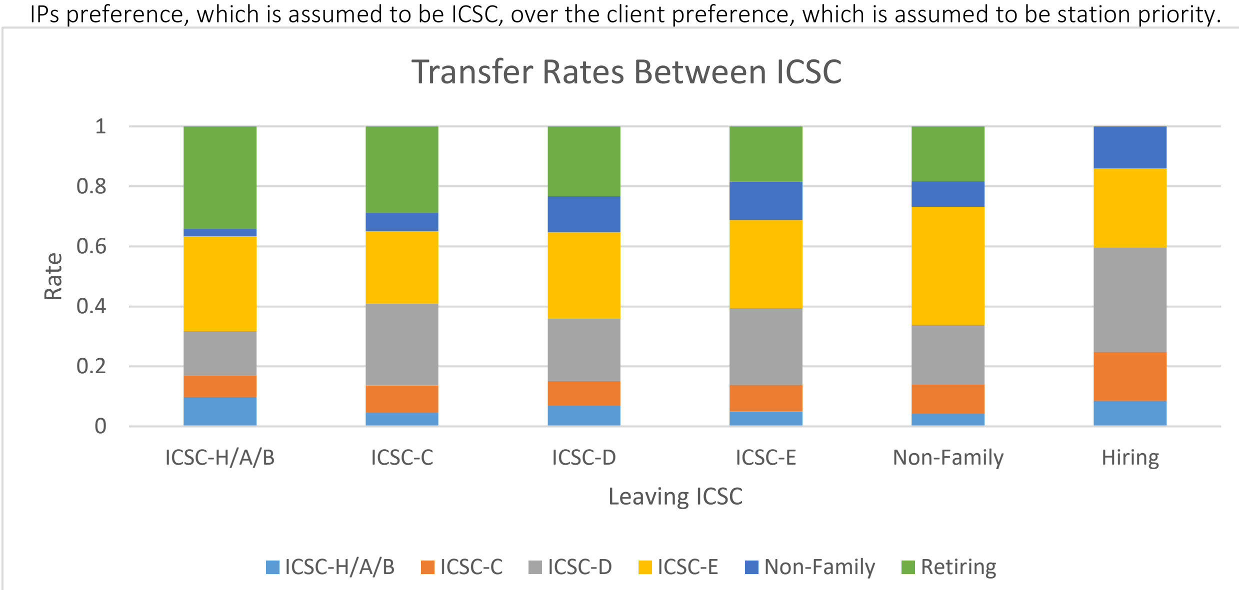

The last output generated by the simulation is a metric that represents on average how much time an IP has till their SDA broken down by ICSC classification. This metric classifies on average how long IPs are staying at a given ICSC station and also depicts the frequency and rotational pattern of IP relocation. This metric was coined the Z-index by Professor Lee Schruben, and follows the following formula:

$$Z_{i,t} = \frac{1}{n_{i,t}} \sum \frac{t - IP_{SDA}}{TOD_i} \quad \forall\, IP \in \{IP_{ICSC} = i\}$$

Where $Z_{i,t}$ is the Z-index for ICSC i at cycle t, and $n_{i,t}$ is the number of IPs belonging to ICSC i at cycle t. The Z-index is essentially an average of how many cycles each IP is over their SDA. If an IP is above their SDA, then they will contribute a positive value, whereas if an IP is below their SDA, they will have a negative value. Since we assume that IPs must transfer when they hit their SDA, all values will be less than zero. Thus, a value of negative one, the minimum, would represent all IPs in that ICSC just began a new post, whereas a value of zero, the maximum, would represent that all IPs in an ICSC are at their SDA. Examples and analysis of the Z-index are shown in figure 6.

Figure 6. Z-Index

Sample values of the Z-index over a simulation run. Top left is the Z-Index for the lowest ICSC class with an SDA of 2 years, which clearly has a strong periodicity. The top right graph also includes the Z-Index for the highest ICSC class with an SDA of 6 years. There is much lower variance and much less pronounced periodicity to the values. Finally bottom left and bottom right compare the Z-index for the lowest ICSC when IPs never accept a match and IPs always accept a match. In the former scenario, after two years, an IP is always replaced by a new hirer, whereas in the latter scenario, the system generally has a smother flow of IPs in and out of posts.

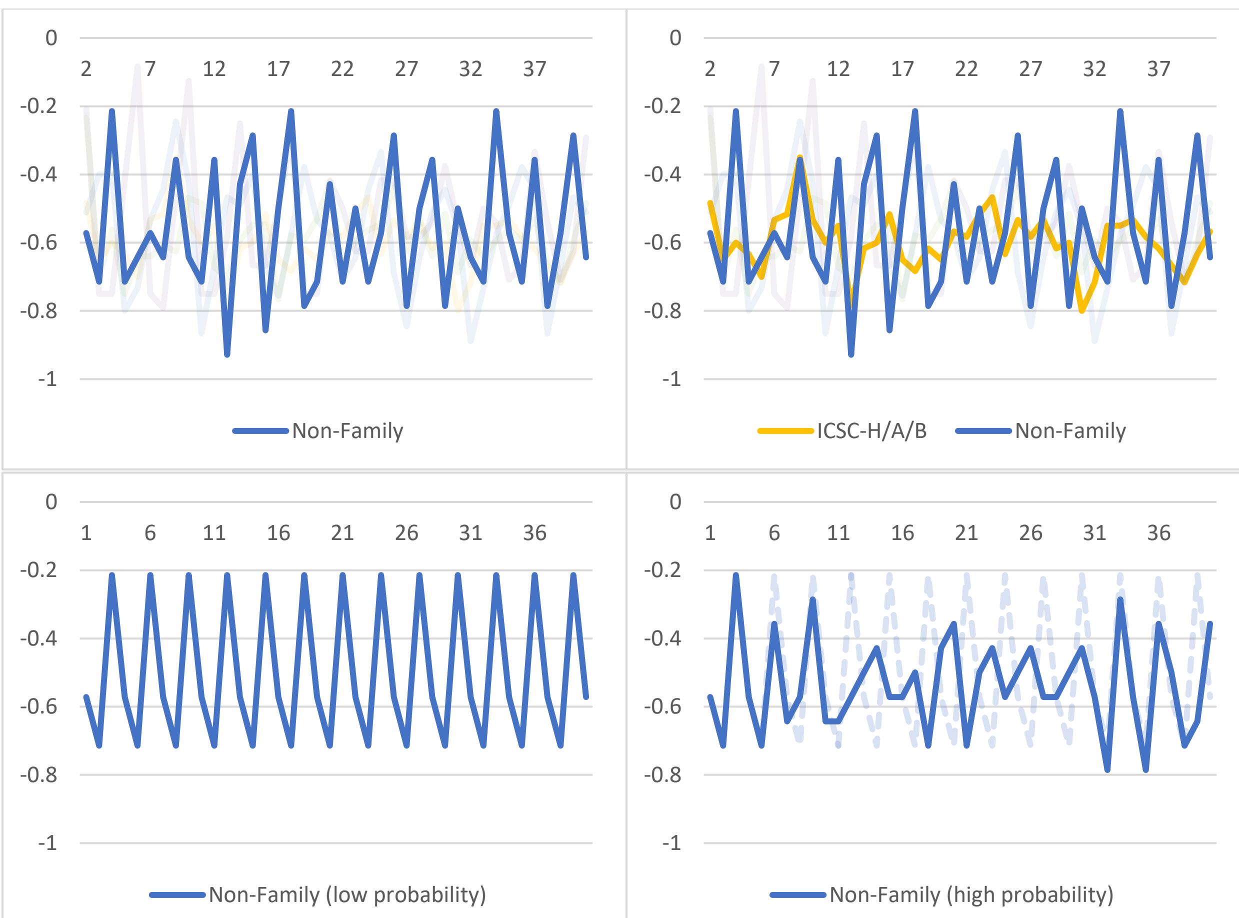

Additional Study

One of the main facets of simulation is that often more questions arise from the study than are answered. This was especially true in this project as our results show that the system performs much more smoothly when IPs are willing to transfer. However, this raises the question of how can the proposed system be implemented in a way that makes IPs more willing to transfer, especially in the third matching stage, SDA final match, when the system decides a match for the IP.

Many systems have been considered, and we decided to study the effects of an open market where IPs have a certain number of points each cycle, and they allocate points to their most desirable posts. The system then assigns matches to globally maximize the points allocated to each match. While the results from this study are intuitive, the main objective is to educate the client on potential pitfalls of adopting such a system and provide a prototype tool that could be used in assigning matches.

To define the mathematical model, let i be the index representing the current post an IP is assigned to and let j be the index to represent what post the IP will be transferring to. Let J be the set of all posts, and note that both i and j belong to the set J. Furthermore, let c_{ij} denote the points the IP in post i allocated to post j (where c_{ij} = -M if the IP in post i did not allocate points to post j), and let p_i be a binary input denoting if an IP is current filling post i. Finally, let the binary variable X_{ij} denote if the IP in post i is matched with post j. This yields the following formulation:

$$\begin{aligned} \max \sum_{i \in J} \sum_{j \in J} & X_{ij} c_{ij} \\ s.t. \sum_{j \in J} & X_{ij} \le p_i \quad \forall\, i \in J \\ \sum_{i \in J} & X_{ij} \le 1 \quad \forall\, j \in J \\ & X_{ij} = 0 \quad \forall\, i = j, i, j \in J \end{aligned}$$

The optimization problem essentially maximizes the number of points each IP uses in their assignment subject to the constraints that a post can only have an IP leave if the post currently has an IP stationed at the post, an IP can only be assigned one transfer, and an IP cannot be assigned to their current post.

During an initial experiment, each IP was allowed to allocate a maximum of 12 points with varying allocation strategies. The general strategy of each IP is they would allocate an assigned maximum number of points to their first preference, with their remaining points scaled accordingly. Thus IPs would allocate anywhere from five to nine points to their most preferred location, while their third location preference would receive anywhere from one to three points based on how they allocated points to their top preference. The result of this experiment is shown in figure 7 on the next page.

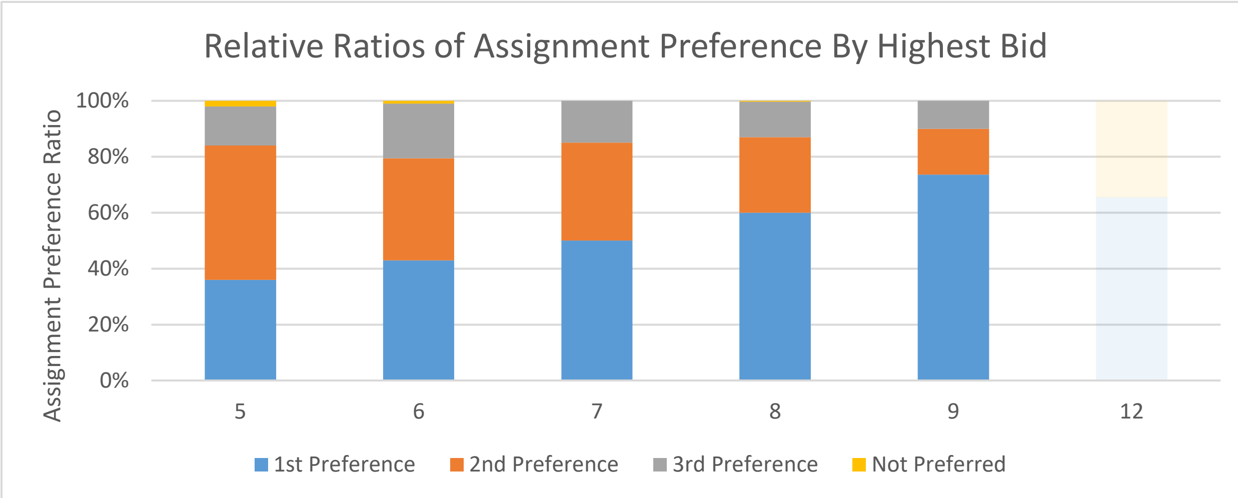

Figure 7. Ratio of Assignment Preference When Highest Bid is 5 to 9 Points

Unsurprisingly, IPs who allocated almost all of their points into their first choice were much more likely to receive it as a match. This is fairly intuitive as the difference in point allocation for somebody who allocates five points to their first choice and four points to their second choice is much smaller than those who wager nine points to their first choice and two points to their second choice. Thus, the system naturally favors those who allocate more points to their first preference. The main concern that accompanies this result is that eventually IPs will realize the system can be exploited, and the system could easily devolve into one where all IPs allocate all of their points to their first pick. The results from simulating this scenario is given in figure 8 juxtaposed with the results from the previous experiment, and clearly indicates that the system breaks down when all IPs allocate 12 points to their first preference with 35% of IPs given a random match and 65% of IPs given their first choice.

Again, these results were expected from the start of the experiment; however, the goal of this analysis was to educate the client on how a point allocation system may work and caution them on some of the pitfalls they may need to address if they choose to allocate such a system. Per a suggestion from Professor Anil Aswani, there may be alternative matching schemes that are “game proof,” such as deferred acceptance models, that would be more appropriate for this system as they are easily solvable and are immune to any individual strategy. These suggestions have been passed on to the client.

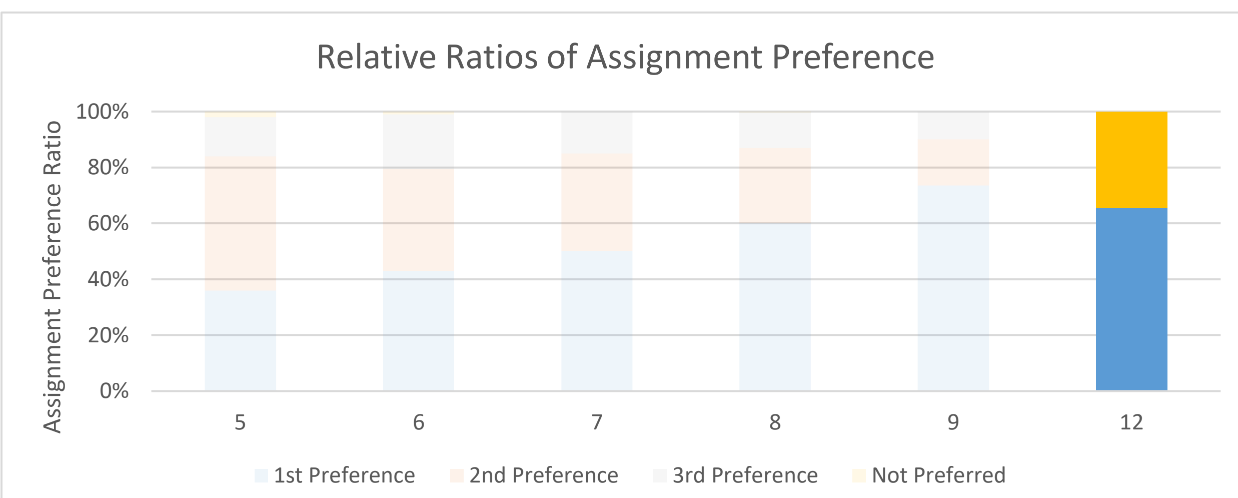

Conclusion

During this project several simulation tools were developed for our client that not only provide insights into how the proposed system behaves, but also raises many questions that may need to be answered before a relocation system can be implemented. One of the most prominent results from this study was the importance of IPs being willing to relocate. Otherwise, IPs will continually leave the system as they reach their SDA which may bring large hiring and training costs on the employer. Thus, effective incentive systems and adopting implementations that give IPs more control over their career will ultimately be very beneficial to the relocation program. Finally, the client should put some consideration into what is considered success in the relocation scheme. Whether it be minimizing cost or maximizing IP morale, the behavior of the optimal system will change drastically based on what constitutes success in the eyes of the client. Overall, this project has been an interesting study in the real world dynamics of complex systems, and the simulation tools we’ve created will help the client to better understand their proposed relocation scheme while also providing key pitfalls that need to be avoided once the system is implemented.