September, 2011

Statistical Correlation Between The LinkUp Job Search Engine and Bureau of Labor Statistics Jobs Reports

Qi Zhang

Bowen Yu

Junlin Hu

Study Advisor:

Laurie Derechin *(Executive Director, MCFAM)*

Minnesota Center For Financial

& Actuarial Mathematics

School of Mathematics

College of Science & Engineering

UNIVERSITY OF MINNESOTA

Step 1 - Count the number of jobs on Linkup

The graph below depicts the monthly statistics of jobs on Linkup for the 3 years between January 2008 and January 2011. We used Matlab to compute this number. We believe that a qualified opening position must satisfy two constraints:

1. The job listing first emerged on Linkup in a given month

2. The job listing has a duration longer than one day

We decided to skip the option of a 30-day average because a simple average of the daily statistics will underestimate the total number of the possible future hires. Jobs can be filled during the current month while the length of the period of each job to be filled is changing according to the economic cycle and is distinctive between different job types.

Therefore, what we took into account are the job listings that emerged on the official corporate websites of companies indexed by LinkUp during each month because we believe this statistic will prove to be a good indicator of the monthly change in the national employment.

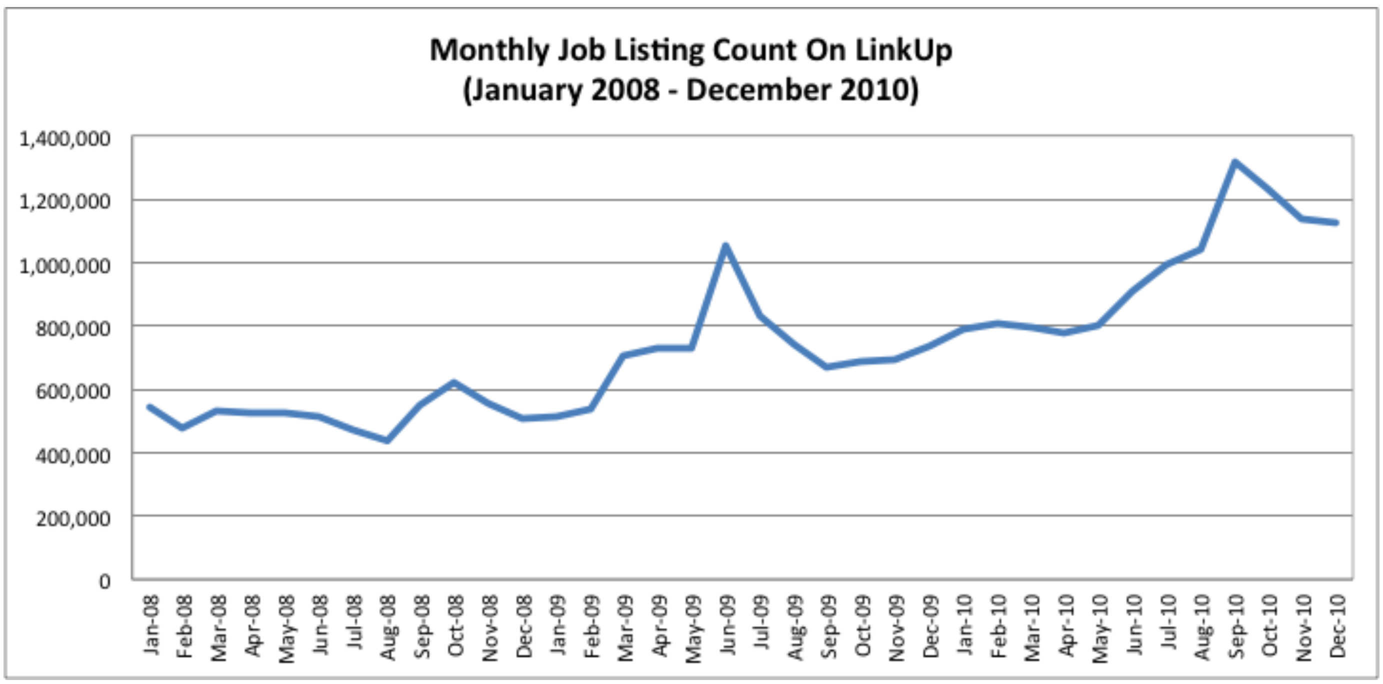

We noticed immediately what appeared to be unusual spikes in the monthly job count in June 2009 and September 2010. Job count data for these two months will be covered later in this report.

Step 2 - Count the number of companies on Linkup

The 2nd step in our analysis was to count the number of companies in the LinkUp index that delivered at least one job opening during a given month. This number is only a fraction of the entire set of companies that LinkUp indexes because in any given month, more than 50% of the companies being scraped for jobs did not have any job listings on their website. We requested data for the entire list of companies with working spiders, but LinkUp does not have that data as the company is not able, at the present time, to distinguish between a company with no job listings and a broken spider, so we agreed that the company count as indicated below was the most suitable available data for our analysis. We did note at the end of the study that correlations might be higher if we were able to use company count data that included all working spiders. LinkUp is working on establishing such a data point.

Step 3 - Calculating the LinkUp Index

After determining the number of job listings and the number of companies, we calculated the LinkUp Index as follows:

LinkUp Index = (number of jobs per month)/(number of companies per month)

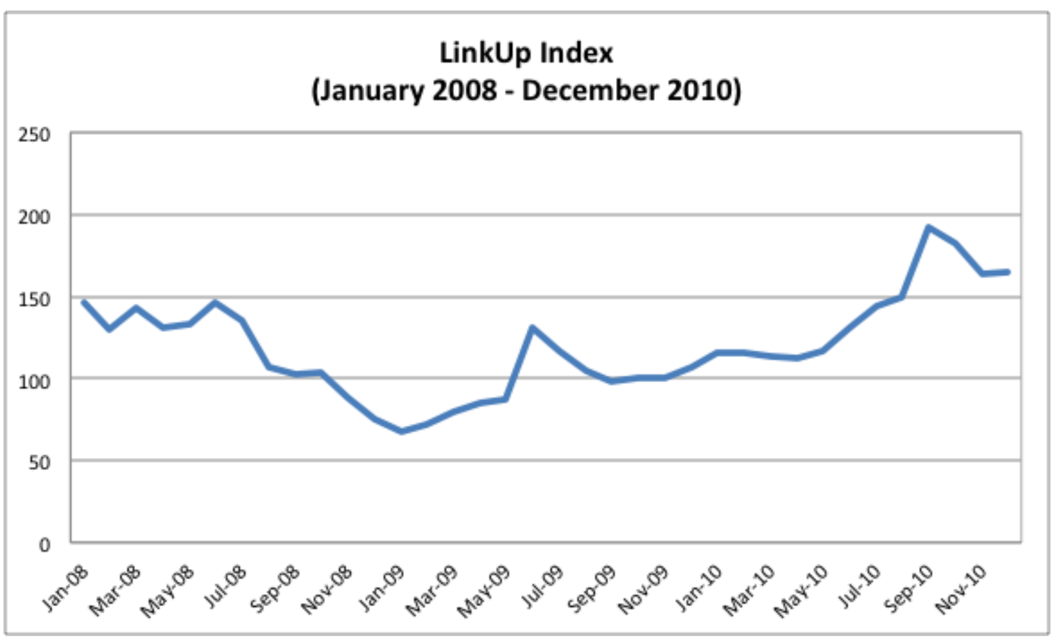

Step 4 - Establish baseline data (U.S. Department of Labor Nonfarm Payroll Report)

Each month the Current Employment Statistics (CES) program within the Department of Labor's Bureau of Labor Statistics, surveys about 140,000 businesses and government agencies, representing approximately 440,000 individual work sites, in order to provide detailed industry data on employment, hours, and earnings of workers on nonfarm payrolls. This nonfarm payroll report (NFP), together with the monthly Household Survey report, form the basis for the monthly jobs numbers issued by BLS.

For purposes of this analysis, we will consider NFP as the primary baseline data rather than the Household Survey and the national unemployment rate. While we would certainly expect to see some correlation between the LinkUp index and the U.S. unemployment rate, the national unemployment rate is more an indicator of labor supply, impacted by factors such as new entrants into the job market and discouraged workers that have left the labor market. The NFP and the LinkUp index are more closely aligned as indicators of labor demand.

Before we calculate the correlation coefficients between NFP and the LinkUp Index, we had to adjust the NFP data to remove the job counts associated with the 2010 Census. The LinkUp search engine did not include any census job listings due to issues with Monster.com and USAJobs.gov, so we needed to remove the census jobs from NFP to obtain a more valid data set for comparison.

Step 5 - Examine correlation between the LinkUp Index and Nonfarm Payrolls with no time lag (t+0)

Before testing the correlation between NFP and the LinkUp Index using various time-lags, we first tested the correlation with no time lag. Not surprisingly, there is no correlation between the LinkUp index in a given month and the NFP for that month.

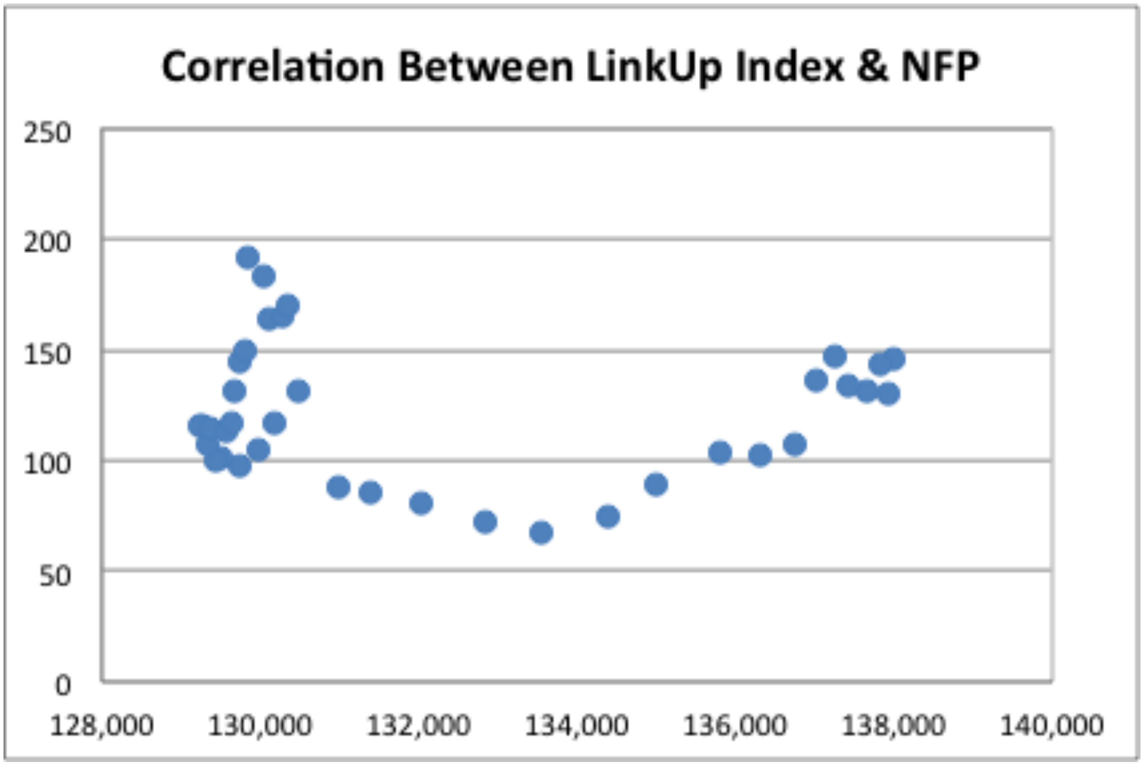

While it could be the case that there might be a correlation between job listings and job growth in a rapid job-growth environment in which jobs openings were posted on company websites and filled very quickly, in less than 2 weeks, for example, this would be a highly unusual circumstance. On average, companies typically take at least 30 days to fill job openings in any type of labor market environment, and in the aggregate across the entire economy, it is hard to imagine that speed-to-hire would ever accelerate enough to trigger correlations between job openings and job growth in the same month.

In any event, LinkUp has the ability in its search engine to calculate how long job listings stay on an employer's website and they would be able to accurately track such an increase and integrate that accelerating speed-to-hire into their forecasts. In the current data set used for this study, the average speed-to-hire for the jobs indexed by LinkUp was 27 days.

Step 6 - Examine the correlation of the LinkUp Index and NFP with 1-month time lag (t+1)

We then tested the correlation between the LinkUp Index and NFP for the following month. This makes intuitive sense because a job listing on a company website indicates an employer's desire or intent to make a hire. But there is a delay between the posting of a job listing and the actual point in time when that position is filled.

Below is the scatter plot of LinkUp Index via NFP's monthly difference (ΔNFP) from January, 2008 to December 2010. Figure 4 shows that ΔNFP is positively correlated with the LinkUp index.

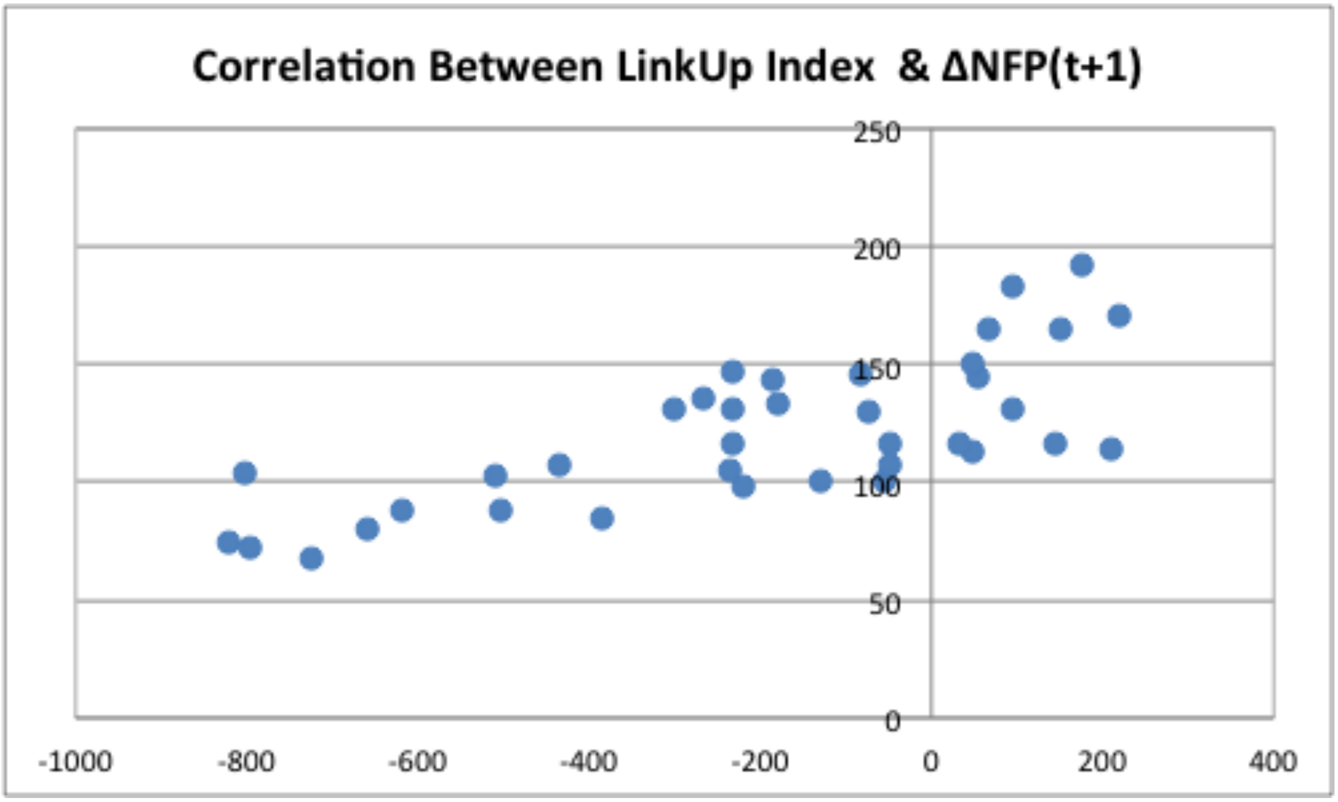

To get a more detailed view of the correlation between LinkUp Index and ΔNFP, we then calculated the Pearson product-moment correlation coefficient (PPMCC) using the following formula:

$$\rho_{X,Y} = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}}$$

Step 7: Determine the correlation between LinkUp Index and ΔNFP for various time lags (t:t+5)

We then determined the correlation between the LinkUp Index and ΔNFP across various time lags up to and including t+5. The time period for these initial correlations is January 2008 to January 2011. Somewhat surprisingly, the correlation for t for the period from January 2008 to January 2011, was higher than the correlation for t+1. As expected, however, the correlations decline as the time lag increases to t+5.

LinkUp Index & ΔNFP						
Period \ Lag	ΔNFP	ΔNFP(t+1)	ΔNFP(t+2)	ΔNFP(t+3)	ΔNFP(t+4)	ΔNFP(t+5)
Jan 08 - Jan 11	0.747961	0.736308	0.670490	0.576904	0.448423	0.357637

Step 8: Calculating the correlations for various time periods and time lags (t:t+5)

Because LinkUp has been continuously improving and adding companies to its search engine, we wanted to test whether or not correlations would be stronger for more recent time periods. In particular, LinkUp made a number of significant improvements to its search engine in 2008 as the site was essentially released to the public. We also wanted to determine if there was a point in LinkUp's history after which correlations for t+1 were stronger than t. For both of these inquiries, we determined a series of correlations between LinkUp Index and ΔNFP using different start dates in 2008. For each period, the end date was January 2011.

In the table below, it is apparent that correlations do improve as the start date approaches December 2008, rising from .736 for the period between January 08 to January 11 to .782 for the period between December 08 to January 11. The table also indicates that beginning with a start date of July 08, correlations between the LinkUp Index and ΔNFP t+1 are consistently higher than between LinkUp Index and ΔNFP t. For the entire series of 12 start dates, the mean correlation of ΔNFP t+1 is .754.

LinkUp Index & ΔNFP						
Period \ Lag	ΔNFP	ΔNFP(t+1)	ΔNFP(t+2)	ΔNFP(t+3)	ΔNFP(t+4)	ΔNFP(t+5)
Jan 08 - Jan 11	0.747961	0.736308	0.670490	0.576904	0.448423	0.357637
Feb 08 - Jan 11	0.743657	0.735981	0.669322	0.582527	0.456923	0.361520
Mar 08 - Jan 11	0.742782	0.735074	0.669977	0.584926	0.457587	0.363865
Apr 08 - Jan 11	0.740859	0.739683	0.678283	0.589438	0.465226	0.373924
May 08 - Jan 11	0.741297	0.742296	0.679175	0.592572	0.469949	0.387795
Jun 08 - Jan 11	0.744383	0.743581	0.683178	0.598892	0.487511	0.411505
Jul 08 - Jan 11	0.750172	0.755176	0.698713	0.633964	0.532068	0.505680
Aug 08- Jan 11	0.754926	0.763361	0.722817	0.667741	0.611258	0.561625
Sep 08 - Jan 11	0.754927	0.763605	0.725512	0.692832	0.621523	0.595295
Oct 08 - Jan 11	0.753111	0.763618	0.746606	0.700256	0.650581	0.621515
Nov 08 - Jan 11	0.752749	0.787497	0.756280	0.737488	0.682481	0.687694
Dec 08 - Jan 11	0.750917	0.781744	0.766498	0.744847	0.710547	0.719029
Mean	0.748145	0.753994	0.705571	0.641866	0.549506	0.495590

Because it takes time for an employer to make a hire after the job has been listed on the company's corporate career portal, it is logical that there would be a time delay of some kind between the rise and fall of new and total job listings on LinkUp and job growth (or decline) in the U.S. as reflected in the BLS Establishment Survey. Because the average job listing indexed by LinkUp appears in the search engine for 27 days, the correlation for the t+1 period would be expected to be the strongest of all the time periods.

It is worth noting, however, that the time lag between posting a job and making a hire can, and most likely does, change over time. In a low unemployment environment, it would be expected that companies need to move more quickly to make a decision about a qualified candidate in order to avoid losing that candidate to another employer. Conversely, in a high unemployment environment, it would be expected that employers would be far more selective in their hiring, taking their time with the process and individual candidates to make sure that they selected the optimal candidate given the abundant supply of labor. In such an environment, it might not be unusual to see the average length of time a job appeared on LinkUp extend to 60 or possibly even 90 days.

As a result of this phenomena, we recommend that LinkUp continually monitor the average length of time that a listing remains in the search engine and adjust its forecasting model accordingly.

In the graph below, we have included the data for ΔNFP (t+1) and the LinkUp index for the period between January 2008 and January 2011 which has a correlation as indicated in Table X of .736.

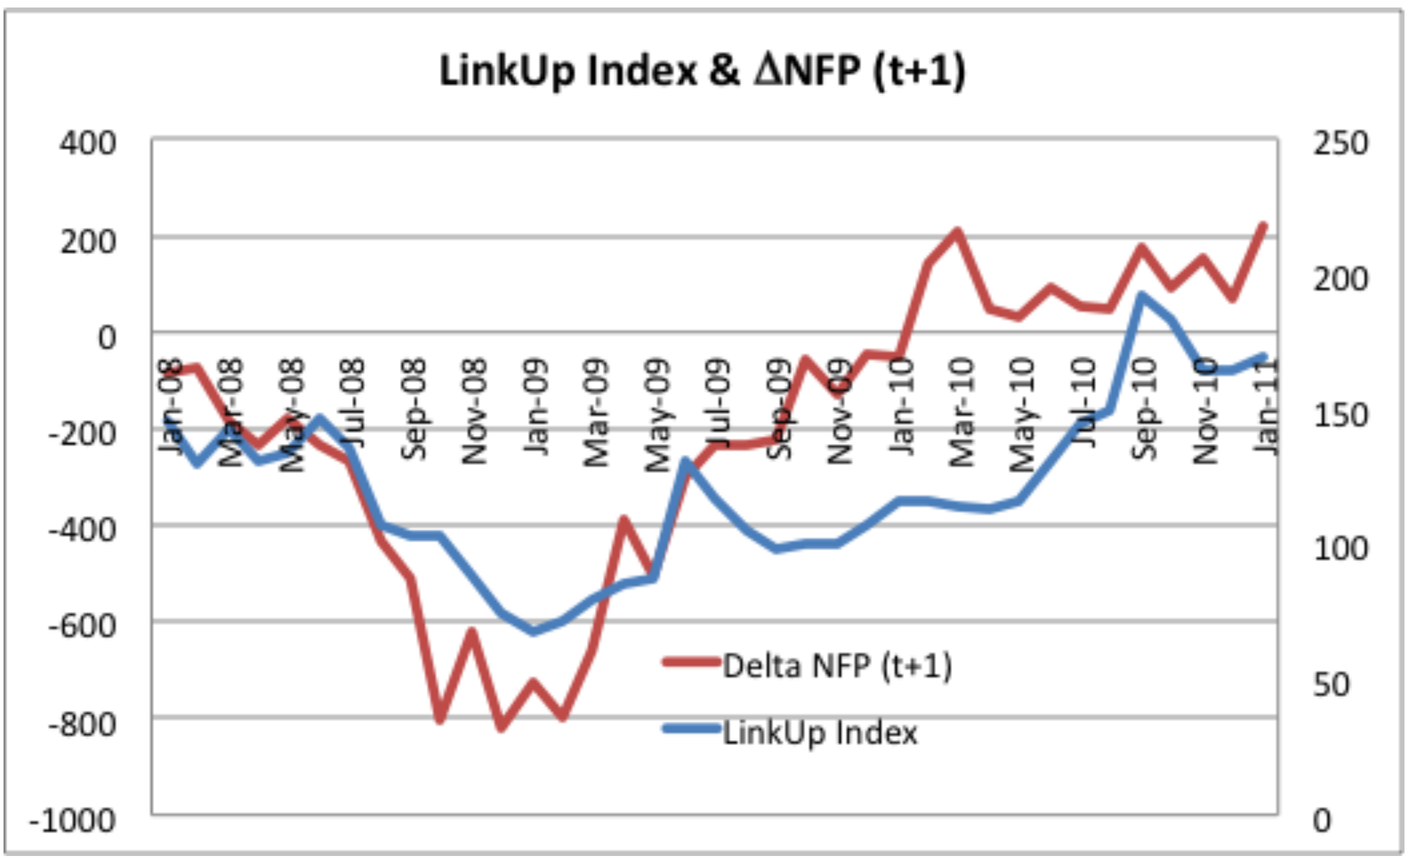

Step 9 - Analyze correlations between ΔNFP and various other jobs data sets

There are at least a dozen popular and widely cited jobs data sets published each month, including, among others, the Conference Board, ADP, Intuit, Monster, and Indeed. Each jobs data set possesses unique attributes and characteristics, but we felt that the Conference Board's Help Wanted Online Data Series (HWOL) was one of the most appropriate to examine for a variety of reasons. The HWOL claims to have the largest data set of job listings available online, including job listings sourced from job boards, company websites, daily and weekly newspaper websites, job aggregator sites, and a variety of other sources. We felt that the HWOL data should encompass the entire U.S. economy, from both a geographic perspective as well as a vertical industry perspective, and should therefore demonstrate a reasonably high correlation to ΔNFP.

As the table below indicates, however, we discovered that the Conference Board's HWOL index contains little to no correlation to the BLS' nonfarm payroll data. As well, there were no time periods in which the correlation between HWOL and ΔNFP was stronger for t+1 than for t.

Help Wanted Online & ΔNFP (unadjusted)						
Lag / Period	ΔNFP	ΔNFP(t+1)	ΔNFP(t+2)	ΔNFP(t+3)	ΔNFP(t+4)	ΔNFP(t+5)
Jan 08 - Jan 11	0.473836	0.333742	0.185425	0.044835	-0.113014	-0.243475
Feb 08 - Jan 11	0.462190	0.328706	0.176166	0.045909	-0.109821	-0.249124
Mar 08 - Jan 11	0.460551	0.322259	0.180631	0.053919	-0.113677	-0.250076
Apr 08 - Jan 11	0.457148	0.330542	0.192198	0.054990	-0.110588	-0.246271
May 08 - Jan 11	0.464981	0.342545	0.195197	0.062057	-0.102975	-0.223503
Jun 08 - Jan 11	0.478193	0.347697	0.204276	0.073446	-0.076252	-0.191814
Jul 08 - Jan 11	0.483650	0.357704	0.216637	0.103010	-0.041760	-0.132998
Aug 08- Jan 11	0.494664	0.371835	0.249031	0.143090	0.028465	-0.083164
Sep 08 - Jan 11	0.514226	0.413997	0.300738	0.240625	0.096579	0.011790
Oct 08 - Jan 11	0.558655	0.470218	0.410275	0.320494	0.210936	0.114545
Nov 08 - Jan 11	0.615469	0.588298	0.495379	0.454406	0.330342	0.261690
Dec 08 - Jan 11	0.708288	0.652748	0.606979	0.559236	0.464134	0.381327
Mean	0.514321	0.405024	0.284411	0.179668	0.038531	-0.070923

We believe there are a number of reasons that explain this lack of a correlation between HWOL and ΔNFP. First, the HWOL includes job listings from hundreds and hundreds of sources, and it is reasonable to assume that there are duplicate job listings included in their data. Despite even the best efforts to eliminate duplicate listings, it is almost impossible to remove them entirely, and the data set is negatively impacted with their inclusion. Secondly, the HWOL includes listings from pay-to-post job boards, newspaper classified sites, and job aggregators such as Indeed and Simplyhired. Each of these sources include job listing pollution such as fake jobs, work-at-home scams, money-mule jobs, phishing jobs, staffing positions, freelance project work, and old job listings. These 'garbage' listings introduce an excessive amount of 'noise' into the data set which limits its ability to accurately forecast the BLS data. And finally, we believe that the HWOL is overly reliant on fading business models such as daily newspapers and pay-to-post job sites which companies are increasingly moving away from in favor of paid search models, social media, and their own corporate career portals on their company website.

As we did for the LinkUp Index, we also graphed the correlation between the HWOL and ΔNFP. While it is apparent that there is some correlation between HWOL and ΔNFP, it is not nearly as strong as the correlation between the LinkUp Index and ΔNFP.

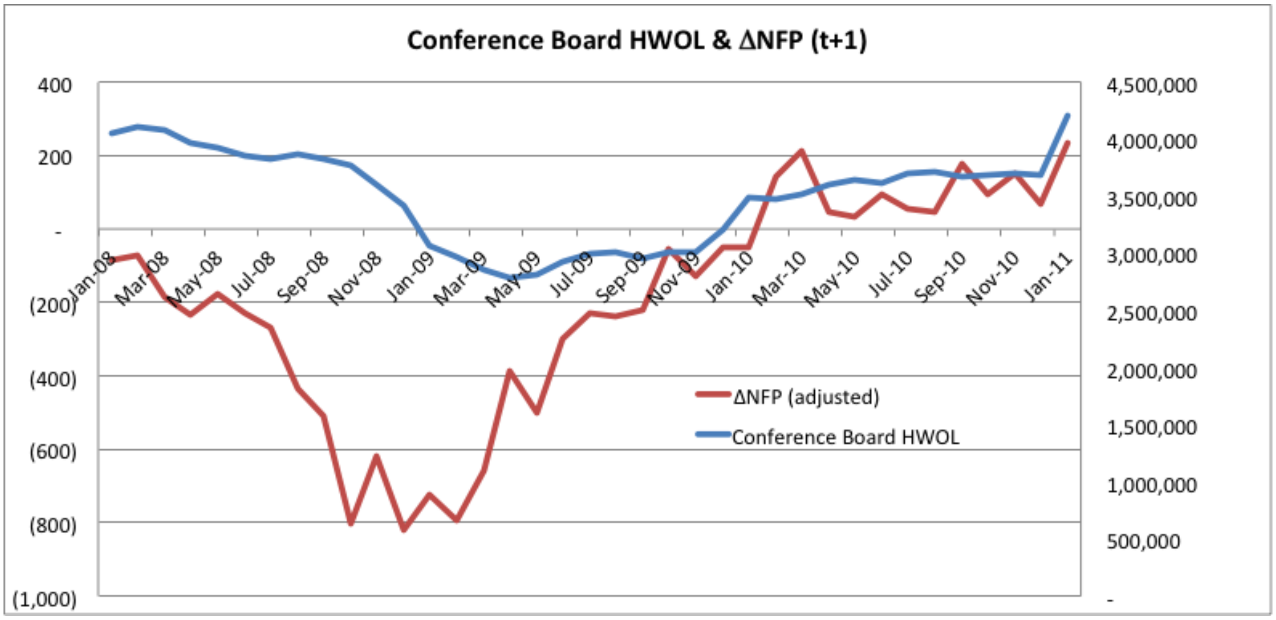

Finally, we calculated the correlations between ΔNFP and various jobs data sets including the LinkUp Index, HWOL, and ADP. The ADP National Employment Report is based on payroll data from over half of ADP's U.S. business clients. The ADP data represents about 24 million employees from all 19 of the major North American Industrial Classification (NAICS) private industrial sectors.

Correlation Between ΔNFP & Various Jobs Data Sets (Jul 08 - Jan 11)

Time:t	ΔNFP	LinkUp	HWOL	ADP
ΔNFP	1.000			
LinkUp Index	0.750	1.000		
Conference Board	0.512	0.607	1.000	
ADP	-0.626	-0.375	0.213	1.000

Correlation Between ΔNFP & Various Jobs Data Sets (Jul 08 - Jan 11)

Time:t+1	ΔNFP	LinkUp	HWOL	ADP
ΔNFP	1.000			
LinkUp Index	0.755	1.000		
Conference Board	0.385	0.607	1.000	
ADP	-0.741	-0.375	0.213	1.000

Step 10 - Model the relationship between ΔNFP & LinkUp Index using linear regression model

In order to further examine the predictive attributes of the LinkUp Index, we prepared a linear regression model to analyze the LinkUp Index and ΔNFP using the formula:

$$\Delta NFP_{t+1} = \alpha LUI_t + \varepsilon_t$$

In particular, we were interested in assessing the extent to which the LinkUp Index could be effectively utilized to predict future nonfarm payrolls not only from a directional standpoint, but also from a degree standpoint. While there might be a variety of techniques to test the correlations of relative sensitivity between the LinkUp Index and ΔNFP, we felt that a linear regression would provide some insight along these lines.

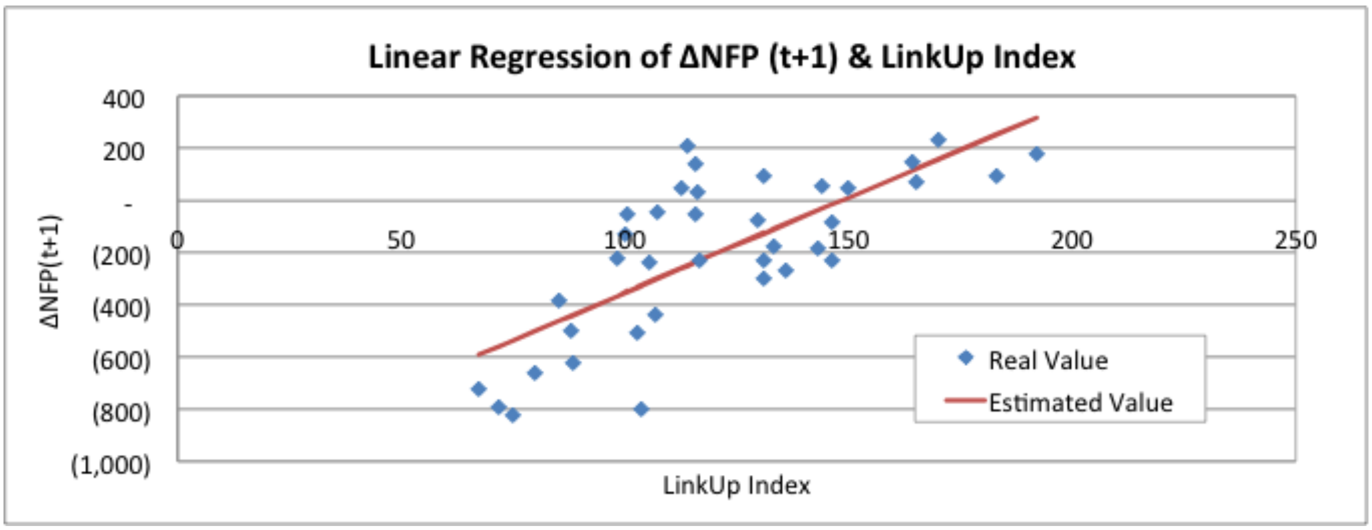

As the regression model above and the regression statistics at right indicate, it can reasonably be concluded that the LinkUp Index is highly correlated to ΔNFP (t+1) from both a directional and degree perspective. While the data set is somewhat limited and does not include data from across a wide range of economic conditions, we would expect that strong correlations would hold up across a full range of expected economic environments.

Regression Statistics	
Multiple R	0.736
R Square	0.542
Adjusted R Square	0.529
Standard Error	210.639
Observations	37.000

Step 11 - Remove data anomalies & determine new correlations

As mentioned earlier, we identified 2 data anomalies in our analysis that warranted further investigation. In graphing monthly job count, we noticed unusual spikes in June 2009 and again in September 2010. In June 2009, LinkUp's job count jumped 45.05% from 726,876 jobs to 1,054,364 jobs. In September 2010, LinkUp's job count jumped from 1,042,676 job listings to 1,320,180 jobs, an increase of 26.61%.

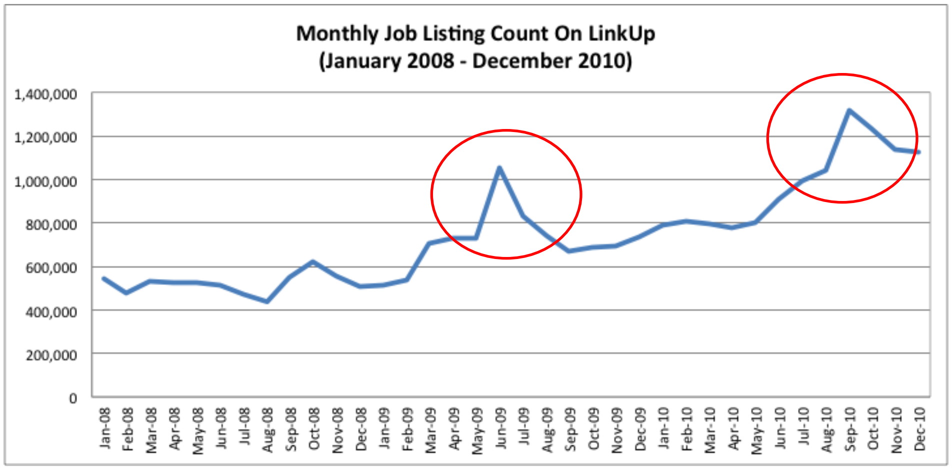

Given the severe spike in job count, we wanted to determine if the unusual data for the periods in question was 'legitimate,' or an incident that was the result of an internal factor that could be corrected or an external factor outside of LinkUp's control. We notified LinkUp of our findings and asked them if they could identify any possible explanations for the high job count. After looking into the issues further, Eric Caron, LinkUp's Web Master provided us with the following explanations:

Quirk on June 2009 (5/28 – 6/10, 6/22-6/23)
"This primarily is a problem from growing pains, where for the first time our existing logic was causing bottlenecks because we just could not process as many jobs for all the companies in the system as we were trying to do. The first notion that this was becoming a problem was on 5/8, but it really started on 5/28 and was fixed on 6/10. There was also an error later in the month that caused a greater-than-normal number of jobs to be deleted, but since it occurred completely in the month of June it may not have had an impact on the dataset." - Eric Caron, LinkUp's Web Master

Quirk on September 2010 (8/16 – 9/10)
"Problem first recognized on August 25th but the problem was introduced on August 16th by introduction of a new location system. The new location system, when implemented, had significant performance problems, and wreaked havoc by slowing down the speed at which jobs were being parsed. This caused larger companies to run less frequently because they took longer to complete, meaning the data was skewed both in that not enough new jobs were being added and not enough old jobs were being deleted – especially troublesome because the company's spider frequency became greatly out of whack from the norm. The problem was reasonably fixed by the 10th and completely fixed by the 14th." - Eric Caron, LinkUp's Web Master.

In addition to the explanations, Eric provided us with the following charts from their servers that further highlighted some of the underpinnings of the data anomaly for September 2010. (Similar charts exist for the June 2009 issue).

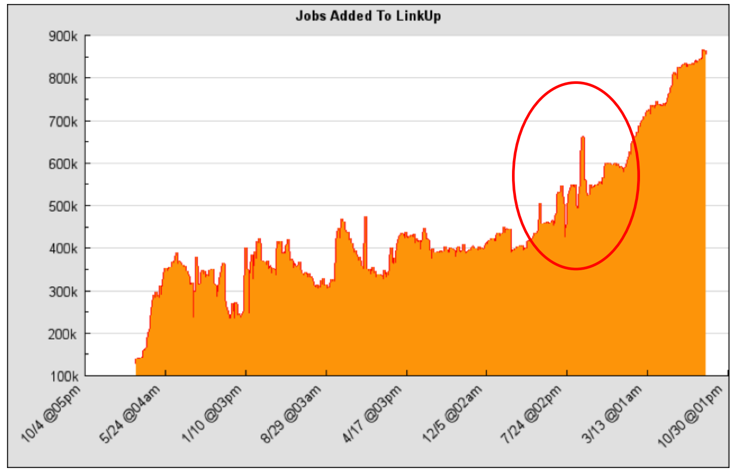

The "Jobs Added To LinkUp" chart represents an hourly snapshot of active jobs stored in the LinkUp database. Because of the graph's high granularity, any jagged edges likely indicate abnormalities in either the health of the system storing the data or the technologies acquiring the data. Fortunately due to the repetitive nature of LinkUp's system, such anomalies can be subtracted from the dataset and corrected by extrapolating the surrounding points.

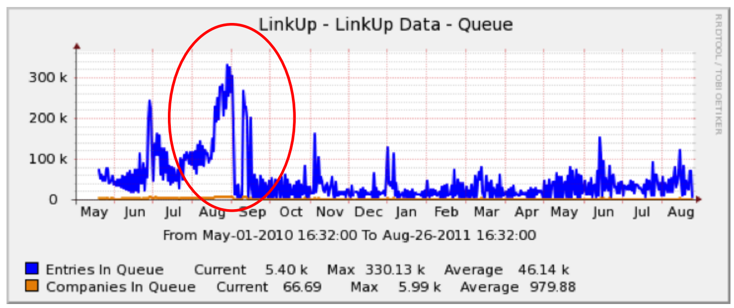

The "LinkUp – LinkUp Data – Queue" graph shows the number of jobs being analyzed by the system at any given moment. Given that the system should have a recurring, continual incoming stream of jobs to analyze and a continual outgoing stream of jobs post-analysis, the graph should look similar to a heart rate monitor showing similar peaks and valleys on a regular schedule, changing in amplitude according to the labor market. Any unusual, extended outliers on the graph indicate time when the system analyzing jobs was not performing as it should.

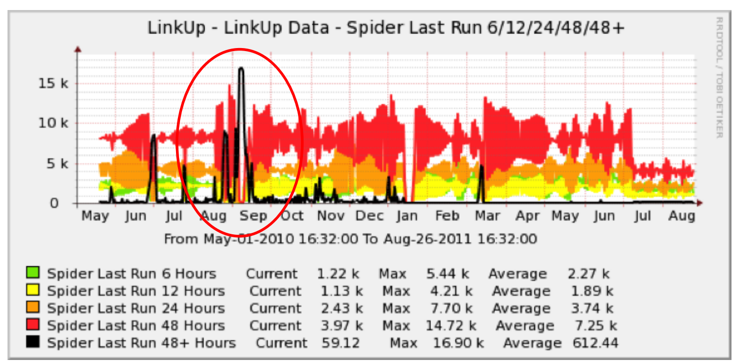

The "LinkUp – LinkUp Data – Spider Last Run 6/12/24/48+" graph represents a snapshot of the number of companies that fall into each segment at any moment. Although some companies update the jobs multiple times per day while others update less frequently, LinkUp has determined that the majority of companies update their jobs section on a near-weekly basis. Because of this, most companies get re-indexed every 24-48 hours. When all systems are behaving normally, the black-line should be near zero and the other lines should be relatively consistent. When the green line becomes too prominent, it means the system is being overly agressive in re-indexing sites; likewise when the black line is too large, it means that a disproportionate number of companies have not been completely indexed in a timely manner.

Upon examination, the data anomalies were both caused by bugs that prohibited the search engine from removing jobs from its index that no longer appeared on company websites. These bugs resulted in an artificially high total job count as new jobs were appropriately added each day, but expired jobs were not removed as they should have been. Given our interest in obtaining the most accurate correlation calculation, we decided that it was worth determining if we could reasonably adjust the data to remove the anomalies and arrive at a more accurate job count for the periods in question.

The problem then became how to adjust the job count data in a statistically defensible manner. After some thought and discussion with LinkUp, we determined that the best means to do this was to calculate a range of statistics related to the expected duration of a job listing on LinkUp. Once that information was obtained, we could create an 'artificial' overlay that would essentially reduce the job count by estimating the expected number of expired job listings each day for a given period and artificially removing them from the job count.

Because it is likely that expected job duration could change over time depending upon economic conditions, we decided to calculate job duration using the 6 month period prior to each anomaly. The distribution of job durations for the period between December 2008 and June 2009 is as follows:

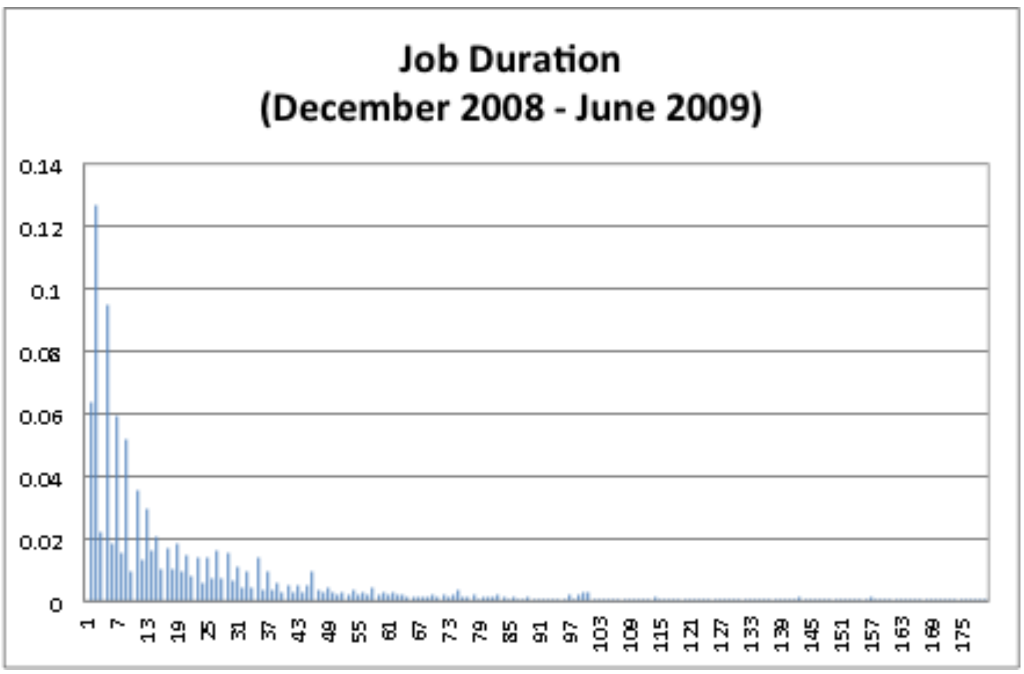

We then applied the data from the Job Duration chart above to the jobs data for the month of June 2009. We constructed a daily 'overlay' that essentially calculated the estimated number of jobs that should have been removed for each day, assuming that the statistical pattern observed above for the prior 6 months also occurred in June.

After applying this overlay, we determined that the June 2009 job count was inflated by 167,696 job listings due to the bug identified in the code that prevented those jobs from being removed as they should have been. The inflated job count represents 15.90% of the total job count for the month.

Corrected LinkUp Job Count (Jan 09 - Aug 11)		
Month	Company Count	Job Count
Jan-09	7,646	515,911
Feb-09	7,496	539,358
Mar-09	8,794	704,014
Apr-09	8,596	731,006
May-09	8,284	726,876
Jun-09	8,042	886,668
Jul-09	7,139	833,760
Aug-09	7,025	739,565
Sep-09	6,843	671,648
Oct-09	6,847	689,204
Nov-09	6,914	692,970
Dec-09	6,863	736,957
Jan-10	6,793	787,335
Feb-10	6,974	807,012
Mar-10	6,996	797,879
Apr-10	6,892	776,669
May-10	6,877	800,911
Jun-10	6,924	906,796
Jul-10	6,903	996,147
Aug-10	6,955	1,042,488
Sep-10	6,861	1,212,455
Oct-10	6,736	1,234,351
Nov-10	6,920	1,137,054
Dec-10	6,812	1,123,133
Jan-11	7,041	1,190,875
Feb-11	7,092	1,116,882
Mar-11	7,257	1,347,618
Apr-11	7,066	1,275,451
May-11	7,966	1,336,337
Jun-11	7,648	1,323,608
Jul-11	7,299	1,230,662
Aug-11	7,048	1,066,104

We applied the same methodology for the 2nd data anomaly in September 2010. After calculating the job duration distribution pattern for the prior 6 months, we constructed an overlay to reduce the inflated job count based on the expected job duration pattern for the month. In this case, we removed 107,725 job listings that we estimated should have been removed from the LinkUp search engine. This represents 8.16% of the reported job listings for that month.

After applying the statistical overlay to the 2 months in question, we arrived at the job and company count for LinkUp as indicated in the table at left.

While it appears that the job count for June 09 and September 10, as well as the months immediately following (July 09 and October 10), might still be slightly elevated, we have no means by which to further refine the counts in an effort to obtain a more accurate data set for those months. We are confident, however, that the approach we took to improve the accuracy of the data represents a valid and statistically defensible methodology, while further steps might compromise the integrity of our results.

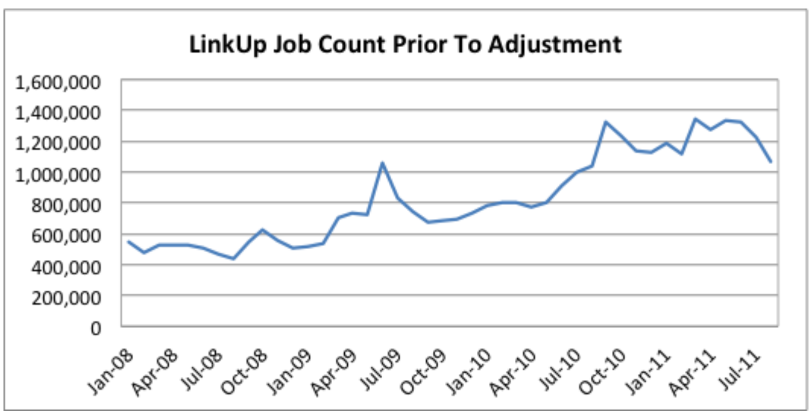

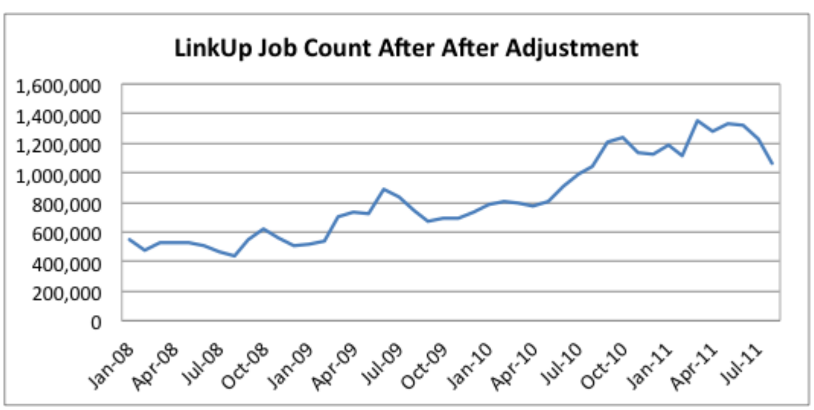

The data table (top left) results in a graph of the job count on LinkUp as seen in the graps at bottom left. We included the same graph of the job count data prior to the adjustment of the 2 data anomalies for reference.

After preparing the adjusted job count data, we calculated the revised LinkUp Index that accounts for the data anomalies in June 09 and September 10. Again, we included the LinkUp Index prior to the adjustment for reference purposes.

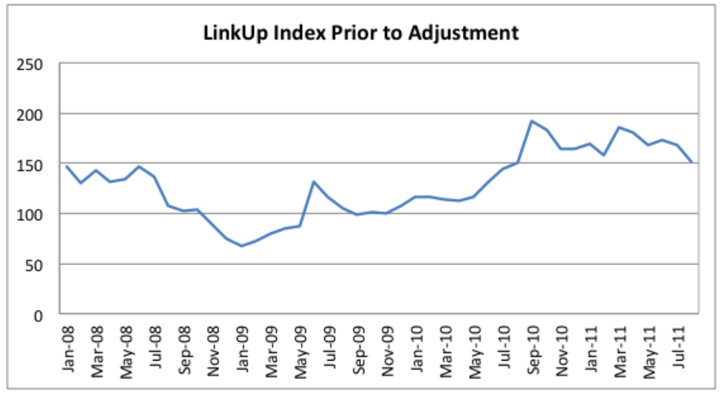

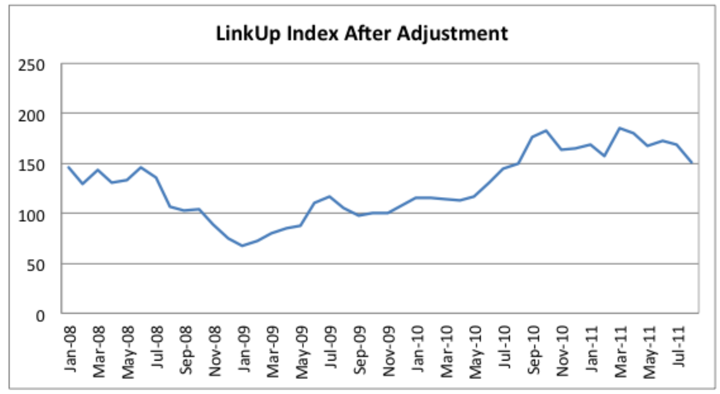

With the revised LinkUp Index, we then recalculated the correlation between the LinkUp Index and ΔNFP (t+1). The revised correlation increased from .754 prior to the adjustment to .769 after the adjustment.

We also produced a graph of the correlation between the LinkUp Index and ΔNFP (t+1) as seen below (bottom right). We have included the same graph prior to the adjustment for reference purposes.

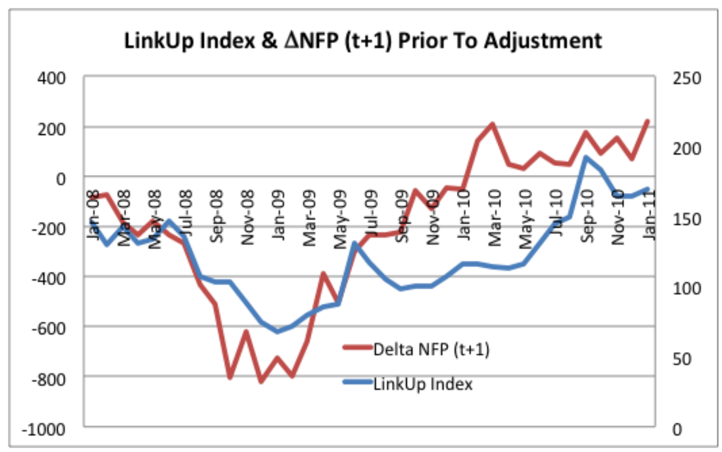

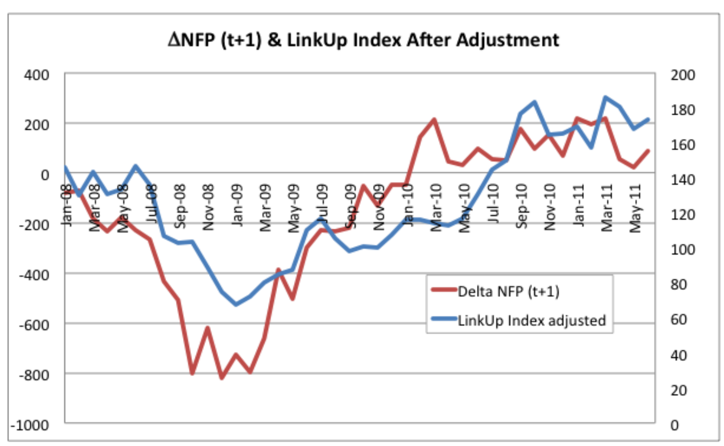

As the correlation calculations and the graphs at right clearly indicate, the LinkUp Index contains strong predictive attributes as it relates to ΔNFP (t+1). The correlations improved meaningfully after accounting for the 2 identified data anomalies.

Based upon the examination and correction of the 2 data anomalies, we recommend that LinkUp continue to closely monitor its job and company count data to assure the integrity of the data and the accuracy of its forecasts. We believe that with close monitoring, careful implementation of code changes, and a continued focus on the jobs data, LinkUp can avoid the types of incidents that created the data anomalies in the first place.

More importantly, we have a high degree of confidence that as LinkUp continues to invest additional time, resources, and expertise into the data portion of its business, the correlations identified in our study will continue to increase further.

Through our study, we have demonstrated that the LinkUp job search engine has very meaningful predictive attributes in relation to the Department of Labor's nonfarm payroll (NFP) report issued by the Bureau of Labor Statistics. LinkUp's highly unique job search engine, which indexes more than 850,000 jobs found on 22,000 corporate career websites across the U.S., delivers a data set that is substantially different from and far cleaner than any available data set related to the U.S. labor market.

LinkUp's job listings contain none of the noise that afflicts the other job openings data sets such as old or expired jobs, duplicate listings, and job pollution created by phishing jobs, work-at-home scams, and fraudulent posts. LinkUp's index contains every single job opening contained on a company's corporate career website, not simply those jobs that the company is actively advertising for in the daily newspaper or on pay-to-post job boards. As the recruitment advertising industry evolves towards newer models such as paid search, social media, and mobile media, the traditional job board model and the data sets that rely on those job listings will become as obsolete as the daily newspaper already is.

In analyzing the predictive attributes of LinkUp's data, we established a model by which we could compare LinkUp's monthly jobs data to the NFP report. We constructed the LinkUp Index to reflect the average number of job listings per employer organization in LinkUp's index. This methodology allowed us to account for the fact that LinkUp is constantly adding new companies to the index which would otherwise influence the data in a traditional month-to-month comparison.

After calculating the LinkUp Index for each month between January 2009 and January 2011, we calculated the Pearson product-moment correlation coefficient between the LinkUp Index and the change between NFP (ΔNFP) for various time periods (t:t+5). We determined that the strongest correlation occurred for the period ΔNFP (t+1), a finding consistent with our analysis that the average job opening is listed in LinkUp's search engine for 27 days. The very high correlation of .754 makes intuitive sense given that a job posting on a company's website indicates a strong intent to make a hire and that companies take an average of 27 days to fill that position. Through regression analysis, we also determined that the LinkUp index is also sensitive to degree. That is to say, the relative changes in the index provide an indication of the expected degree of change in the nonfarm payroll report.

We also determined through our analysis that the LinkUp Index is far more predictive of ΔNFP than the Conference Board's Help Wanted Online Index (HWOL) which demonstrated almost no correlation to ΔNFP. Finally, we corrected the LinkUp data set to account for 2 identified data anomalies caused by bugs in the LinkUp code and recalculated the correlations. After constructing an overlay for June 2009 and September 2010 that removed expired jobs in those 2 months that the search engine had not removed in actuality, the correlation to ΔNFP (t+1) rose to .769.

In summary, we believe that the LinkUp Index can be effectively utilized to predict the nonfarm payroll report issued by the Department of Labor each month. This very strong correlation between the LinkUp Index and ΔNFP exists because of the very unique attributes of LinkUp's job search engine which also makes LinkUp the strongest indicator of what is and will be happening and in the U.S. labor market of any data provider in the market today. We also believe that as LinkUp continues to improve its technology and add new companies, employer organizations, and government job listings into its index, the correlations will grow even stronger.