2011 11th IEEE International Conference on Data Mining Workshops

Estimating Post-Event Seller Productivity Profiles in Dynamic Sales Organizations

Kush R. Varshney, Moninder Singh, Mayank Sharma, and Aleksandra Mojsilović
Business Analytics and Mathematical Sciences Department
IBM Thomas J. Watson Research Center
Yorktown Heights, NY, USA
Email: {krvarshn,moninder,mxsharma,aleksand}@us.ibm.com

***Abstract*—In modern sales organizations, the salesforce is constantly in flux due to sellers retiring or leaving to take positions in other organizations, and due to new sellers being hired from universities, from other organizations, and transferring in from different divisions within the enterprise. The productivity of sellers in the time period after these human resource events is not the same as that of entrenched sellers. It takes time for new sellers to ramp up their productivity and the ramping up profile of productivity for new sellers of different categories varies. Also, sellers that attrit have a lasting effect on sales in the pipeline after they leave. Thus revenue projections and other business planning decisions cannot simply be made from the total headcount of sellers; post-event seller productivity profiles must also be taken into account. In this paper, we propose a regularized estimation technique using linear programming to minimize maximum residual error for determining seller productivity profiles from aggregate sales revenue and headcount data that can then be used in future planning. We demonstrate the estimation technique on real sales data from a global enterprise and compare results to productivity profiles elicited from salesforce leaders in the enterprise.**

***Keywords*-linear programming; minimax estimation; productivity profiles; salesforce analytics;**

I. Introduction

Workforce productivity is a critical driver of the success of enterprises; improvements in worker productivity can lead to substantial gains in both revenue and profit. For example, the speed at which calls can be correctly handled by a call center associate has a direct impact on the operational success of the organization. Similarly, in an enterprise with a large services or product portfolio, the effectiveness with which its salesforce can sell services or products to its clients directly affects its bottom and top line [1], [2]. This is also true for a non-client-facing workforce, such as assembly line workers or software developers [3], where the productivity of the workers is a major determiner of the throughput, and thus of the financial performance of the organization.

Thus, strategies for increasing worker productivity have been widely studied over the past decade, with much of the focus being on the development of various information technology based platforms and techniques for this purpose. These include various customer relationship management (CRM) systems [4], [5] for the integration and management of data relevant to the complete marketing and sales processes, as well as salesforce automation (SFA) systems [6] for allowing sales managers to better balance available seller resources against identified opportunities by automating the various stages of the sales process.

While such tools provide improvements in productivity by improving the overall efficiency of (and automating) the sales process, additional major improvements in productivity require insights that can be obtained through predictive and prescriptive analytics using enterprise data. For example, sales demand forecasting algorithms can be used to estimate the quantity of a product or service that consumers might purchase, a critical input into sales resource planning [7]. Similarly, sales territory design and optimization techniques can be applied to enable more effective allocation of sellers to customer accounts or sales districts [8] while sales recommendation systems [9], customer propensity modeling and market opportunity estimation algorithms [10] can be used to aid sales resources in more effectively identifying and serving client opportunities.

These methodologies mostly focus on a single dimension of the sales process: the client dimension. However, productivity of sellers, especially in a large company with a large product portfolio and a diverse client base, is affected by many additional factors along multiple dimensions, such as seller characteristics (e.g., skills, experience) and marketplace opportunity (e.g., competitive landscape, size of potential sales) as well as external factors (e.g., macroeconomic conditions). This often leads to large differences in sales performance across various business units, product groups, or geographies within an organization.

Businesses therefore look for novel ways to obtain an accurate assessment of sales performance, benchmark different sales groups and units, and provide guidance to identify, design and manage performance improvement initiatives. Several studies have investigated the effects of these factors on sales performance. However, the work has primarily focused on analysis using a single factor at a time, cf. [11], [12], [13], [14], [15].

Substantial work has also been done for the identification of the most relevant factors behind sales performance. This work has mostly been qualitative (e.g., interviews, case studies and literature review) and therefore has been unable to quantify the underlying relationships. For example, a

978-0-7695-4409-0/11 $26.00 © 2011 IEEE
DOI 10.1109/ICDMW.2011.77

meta-analysis is presented in [16] where relevant findings from published and unpublished work are aggregated into a set of major determinants of seller performance. Issues relevant to measuring the performance of industrial sellers are reviewed in [1] through a set of interviews and self-assessments conducted with sellers and managers from five major industrial firms. More recently, [2] describes a unifying framework for modeling the relationship between various productivity drivers/inhibitors and business success along multiple dimensions for identifying actionable drivers, and determining the best set of actions to improve both seller and organizational performance given various constraints.

However, all these approaches focus on modeling and understanding *steady state* workforce productivity such as identifying attributes that affect such productivity (e.g. training and skill of workforce) and actions that can be taken to increase it (e.g. hire more skilled workers). One aspect, though, that has traditionally not been focused on is understanding the temporal nature of the productivity of a typical worker after certain HR-related events, such as the hiring of the worker, a transfer to another business unit within the same organization, or the retirement of the worker from the organization. In such cases, the productivity of a worker changes over time post-event until it reaches steady state and the worker performs at the level of a typical worker.

Thus, in order to accurately model and estimate workforce productivity, it is imperative to estimate how worker productivity changes over time after such events [17]. For example, a new hire fresh from a university (with no experience) cannot reasonably be expected to become productive from day one. One would expect an initial period of little or no productivity during which the seller learns the ropes, gets familiar with the organization, products, etc., followed by a period of gradually increasing productivity until the seller reaches an average, steady state productivity level. Similarly, an experienced seller (who may have had very high productivity at his or her previous position) would still be expected to have a period of little to low productivity as the worker gets familiarized with the products and processes of the new organization. However, on average, this period would be lower than that for an inexperienced hire and the experienced hire would be expected to have a much shorter time period to reach the productivity of a typical seller. On the other hand, the productivity of a seller who leaves an organization will not typically fall to zero immediately but would be expected to do so gradually over some time as deals in advanced stages that the seller had been working on close despite his or her leaving.

In any modern sales organization, the workforce is in a state of constant flux with often rapidly changing underlying headcount dynamics, with sellers being hired, moving across jobs within the organization, joining via acquisitions, or leaving due to various reasons (e.g. retirement, another job). This is often true even though the overall headcount may be static or show little change, especially in large organizations with complex and rich product/service portfolios and large client-facing workforces. Thus, modeling the post-event productivity profiles of workers is necessary, both at an individual level as well as at an organizational level.

At the individual level, post-event productivity profiles help in setting the right expectations (especially for new hires) as well as help the organization identify the kind of help and training need (and the best time to provide it). However, a much more significant impact of such post-event productivity profiles is at the organizational level, and is the main motivation behind the work described in this paper.

Traditionally, organizations have made strategic as well as tactical decisions and conducted planning, such as making revenue projections, making budget allocations and hiring decisions, and setting financial targets on the basis of the sales headcount. However, the headcount is not a true representative of the effective sales capacity due to the underlying sales dynamics as explained above. Each salesforce event, such as a new hire, has an immediate impact on the total headcount; however, its effect on the sales capacity of the workforce is not felt immediately. Modeling and using productivity profiles for each type of headcount dynamic can help determine the true sales capacity (the *effective* number of sellers that are selling). As an example, consider a scenario where 100 current sellers leave and are replaced by 100 inexperienced, university hires. Assuming that inexperienced hires have zero productivity upon starting, the headcount will remain unchanged but the sales capacity will drop by 100.

Salesforce productivity, as measured by revenue per unit of sales capacity, thus offers a much stronger metric to base decisions on, as compared to the traditionally used measure of productivity as revenue per unit of sales headcount. This has a strong impact on decisions made at multiple levels and functions within an organization. For example, executives often compare various regions or brands to identify areas of poor performance and to make decisions regarding future investments and deployment/hiring of sellers. Comparing revenue per unit of sales capacity offers a much more accurate measure to do so than the traditional revenue per headcount measure. Similarly, financial planners often use past revenue per headcount to set sales targets for the future. Once again, better and more accurate estimates can be obtained by basing such targets on productive capacity.

Moreover, productivity curves can also help organizations in carrying out 'what-if' analyses to see the effect of various hiring plans and simulate the effect of implementing such plans on future revenue attainment. For example, an organization may decide to hire a number of sellers to meet near term revenue estimates (or fill a gap between targeted revenue and expected revenue). While it may seem to be an appropriate decision assuming each seller will bring in the same amount of revenue that each current seller does, it will likely be a gross overestimate since most new sellers will likely have varying levels of zero or little productive capacity in the near future. Thus, if the organization hires a number of inexperienced sellers, their addition will likely add nothing to the immediate term revenue attainment (assuming zero productivity for such sellers for the first few months) but the models based on headcounts will make inaccurate assumptions of the available salesforce capacity and overestimate the revenue potential of the salesforce.

The estimation of typical productivity profiles from very noisy samples of individual sales productivities is similar to estimation of time-response signals to events that has been studied in various domains, and can thus be potentially approached using similar methods as described in [18], [19]. However, there are two issues that prevent such methods from being applied here. First, individual sales data is often very difficult to get within an organization due to privacy rules as it is closely related to individual performance and compensation. Second, even if such data was acquired (say, by anonymization), individual revenue data is often not reflective of the individual's true productivity. This is because sellers often work in teams and take credit jointly and individually for any transaction, a practice called 'stacking,' which often ends up highly inflating the revenue data that is associated with an employee. Moreover, the magnitude of such stacking may vary widely with geographies, brands, job roles, etc. As such, estimating productivity profiles from such data will likely generate incorrect results. On the contrary, aggregate revenue data from the ledger gives a true picture of the actual revenue attained by the entire salesforce. Hence, in this paper, we describe an approach for estimating productivity profiles using aggregate revenue and headcount data.

The remainder of the paper is organized as follows. In Section II, we describe salesforce data that would typically be available from a CRM to use in performing analytics. Additionally, we propose an estimation scheme to learn temporal seller productivity profiles from this data for different classes of new sellers as well as sellers that attrit. Section III validates the proposed approach through application to a real-world data set from an enterprise's sales organization. In Section IV, we summarize the paper and provide perspectives on future directions of research.

II. Productivity Profile Estimation from Aggregate Data

In this section, we describe aggregate salesforce data of revenue and headcounts that is available in CRMs and present an estimation approach for determining sales productivity and attrition profiles from such data.

A. Aggregate Data

Historical revenue and headcount data for the entire salesforce is typically available within any sales organization. Moreover, in addition to the total headcount at any point of time, organizations also track the underlying salesforce dynamics such as the number of sellers being hired, acquired via acquisitions or joining via internal transfers, as well as number of sellers leaving due to retirement or moves to external organizations. We present a formulation to infer productivity profiles of new sellers as a function of time after joining the organization and attrition profiles of residual productivity after sellers depart the organization.

We denote the total revenue for all sellers in the salesforce at time k by R_k. Then for a series of ξ times, the available revenue data is $R_1, \ldots, R_\xi$. We index productivity profiles by times $j = 1, \ldots, \nu \leq \xi$, and we index classes of new sellers by $i = 1, \ldots, \mu$. The available headcount data is as follows. The value $n_{i,j,k}$ is the headcount of sellers at time k that joined j times earlier as a member of class i. The value $n_{\text{old},k}$ is the headcount of sellers at time k that joined more than ξ times ago. The value $n_{A,j,k}$ is the attrition headcount of sellers that left the salesforce j times before the time k. With these definitions, $n_{i,j,k} = n_{i,j+1,k+1}$ and $n_{A,j,k} = n_{A,j+1,k+1}$.

B. Regularized Linear Programming Formulation

Given the total salesforce revenue and headcount data, we would like to determine the ν-time long productivity profiles of new sellers for each of the μ classes. (We will come back to attrition profiles in Section II-C.) We approach this problem as one of data fitting and focus on minimizing the maximum residual error using linear programming [20, Sec. 1.3]. The general approach is as follows.

Consider the collection of data points $\{(\mathbf{n}_1, R_1), \ldots, (\mathbf{n}_\xi, R_\xi)\}$, $\mathbf{n}_k \in \mathbb{R}^\nu$, $R_k \in \mathbb{R}$, that we wish to model linearly through a parameter vector $\boldsymbol{\alpha} \in \mathbb{R}^\nu$ as $R = \mathbf{n}^T\boldsymbol{\alpha}$. The optimization problem to minimize the maximum residual error is:

$$\min_{\boldsymbol{\alpha}} \max_{k} |R_k - \mathbf{n}_k^T\boldsymbol{\alpha}|,$$

and can be equivalently written [20, Sec. 1.3]:

$$\begin{array}{lll} \text{minimize} & z & \\ \text{subject to} & R_k - \mathbf{n}_k^T\boldsymbol{\alpha} \leq z, & k = 1, \ldots, \xi \\ & -R_k + \mathbf{n}_k^T\boldsymbol{\alpha} \leq z, & k = 1, \ldots, \xi. \end{array}$$

For our problem at hand, the parameter vector is the productivity profile measured in revenue per seller, and we have μ different parameter vectors: one for each class. Thus, we denote the productivity profile for sellers in class i by $\alpha_{i,1}, \ldots, \alpha_{i,\nu}$. Additionally, we assume that after ν months in the salesforce, sellers become 'old-timers' with productivity L. Here L is a saturation level of productivity (which we also optimize for data fit). Putting these pieces together, our linear model for revenue is:

$$R_k = \sum_{i=1}^{\mu} \sum_{j=1}^{\nu} \alpha_{i,j} n_{i,j,k} + L n_{\text{old},k}, \quad k = 1, \ldots, \xi.$$

$$
\begin{aligned}
&\text{minimize} && z && (1)\\
&\text{subject to} && R_k - \sum_{i=1}^{\mu}\sum_{j=1}^{\nu}\alpha_{i,j}n_{i,j,k} - Ln_{\text{old},k} \le z, && k=1,\dots,\xi\\
& && -R_k + \sum_{i=1}^{\mu}\sum_{j=1}^{\nu}\alpha_{i,j}n_{i,j,k} + Ln_{\text{old},k} \le z, && k=1,\dots,\xi\\
& && 0 \le \alpha_{i,1}, && i=1,\dots,\mu\\
& && \alpha_{i,j} \le \alpha_{i,j+1}, && i=1,\dots,\mu,\ j=1,\dots,\nu-1\\
& && \alpha_{i,j+1} - \alpha_{i,j} \le L/m, && i=1,\dots,\mu,\ j=1,\dots,\nu-1\\
& && \alpha_{i,\nu} \le L, && i=1,\dots,\mu.
\end{aligned}
$$

For purposes of regularization, we restrict productivity profiles to belong to a certain function class through the following constraints. We constrain them to be monotonically non-decreasing in time and in the range $[0, L]$, and also include a maximum slope constraint parameterized by a fixed value m. The constrained linear program that we solve to determine productivity profiles is given in (1).

C. Formulation with Attrition

We may also include the effect of attrition on the total revenue because the sales from the pipelines of departed sellers keep coming in after they have left the salesforce. We assume this after-effect lasts for ν_A times. The attrition profile is denoted $\beta_1, \dots, \beta_{\nu_A}$. The model for total revenue with attrition is:

$$
R_k = \sum_{i=1}^{\mu}\sum_{j=1}^{\nu}\alpha_{i,j}n_{i,j,k} + \sum_{j=1}^{\nu_A}\beta_j n_{A,j,k} + Ln_{\text{old},k}, \qquad k = 1,\dots,\xi.
$$

We include similar regularization for the attrition profile as for the productivity profiles. We constrain the attrition profile to be monotonically non-increasing in the range $[0, L]$ and have a minimum slope constraint parameterized by the value m_A. The linear program with attrition is given in (2).

III. Empirical Results

We apply the estimation methodology discussed above to real selling data from a large global enterprise. The data represents headcount and revenue data for five and a half years, and covers over 10,000 sellers for one business unit spanning all geographies. The revenue we consider is transactional revenue; it is adjusted for seasonal effects, such as more sales in the fourth quarter of the year. The headcount data is for all seller job categories, including technical and client-facing sellers. The $\mu = 4$ different classes of sellers that join the salesforce are: transfers in, new hires with experience, new hires from universities, and sellers from acquisitions. We differentiate between these kinds of seller additions based on feedback from sales leaders that indicate significant differences in productivity between these classes of sellers after joining a sales organization.

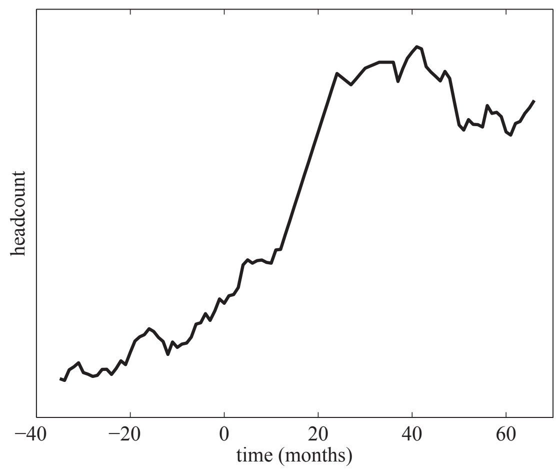

Figure 1. Total headcount data.

Fig. 1 shows the total headcount of sellers in the organization during the period we examine.[1] Fig. 2 shows headcount dynamics for the different classes of new sellers and sellers that attrit. The values shown indicate the number of sellers of each category that joined the organization during the month. The attrition headcount dynamic is shown as a positive quantity but represents a decrease in the total headcount. In this analysis, we take ν to be 36 months, so the number of 'old-timers' at month 1 can be calculated by subtracting the sum of new sellers in months -35 to 0 in all classes from the total number of sellers at month 1.

We run the proposed linear program without attrition and with slope constraint parameter $m = 1/6$ to estimate productivity profiles for the different classes of new sellers. While smaller values of m (such as $1/10$) lead to smoother productivity curves, no major differences in the shape of the solution profiles was observed when changing m. As such,

[1]Due to the confidential nature of the data used, we do not show the actual values of revenue, total headcount or headcount dynamics in the figures. Also, the productivity profiles are shown after normalizing the average revenue per salesperson.

$$
\begin{aligned}
\text{minimize} \quad & z \qquad\qquad (2)\\
\text{subject to} \quad & R_k - \sum_{i=1}^{\mu}\sum_{j=1}^{\nu}\alpha_{i,j}n_{i,j,k} - \sum_{j=1}^{\nu_A}\beta_j n_{A,j,k} - Ln_{\text{old},k} \le z, && k = 1,\ldots,\xi\\
& -R_k + \sum_{i=1}^{\mu}\sum_{j=1}^{\nu}\alpha_{i,j}n_{i,j,k} + \sum_{j=1}^{\nu_A}\beta_j n_{A,j,k} + Ln_{\text{old},k} \le z, && k = 1,\ldots,\xi\\
& 0 \le \alpha_{i,1}, && i = 1,\ldots,\mu\\
& \alpha_{i,j} \le \alpha_{i,j+1}, && i = 1,\ldots,\mu,\ j = 1,\ldots,\nu-1\\
& \alpha_{i,j+1} - \alpha_{i,j} \le L/m, && i = 1,\ldots,\mu,\ j = 1,\ldots,\nu-1\\
& \alpha_{i,\nu} \le L, && i = 1,\ldots,\mu\\
& 0 \le \beta_{\nu_A},\\
& \beta_{j+1} \le \beta_{i,j}, && j = 1,\ldots,\nu_A-1\\
& \beta_j - \beta_{j+1} \le L/m_A, && j = 1,\ldots,\nu_A-1\\
& \beta_1 \le L.
\end{aligned}
$$

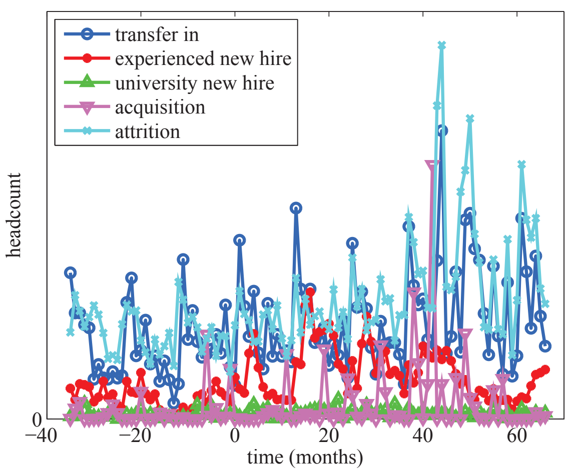

Figure 2. Headcount dynamics.

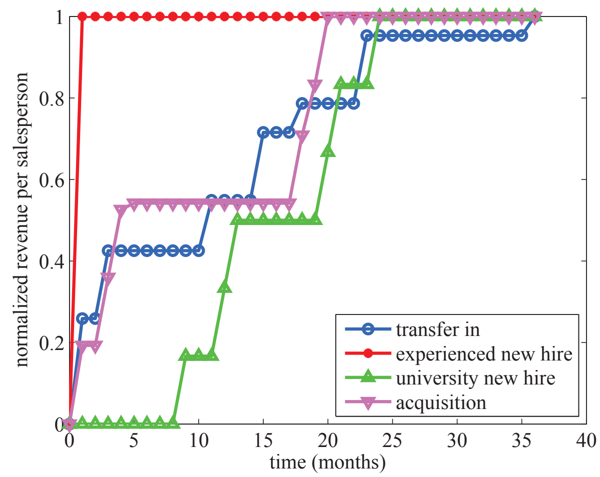

Figure 3. Estimated productivity profiles without attrition.

we chose $m = 1/6$ based on feedback from sales leaders that sellers typically took at least 6 months to ramp up to their final productivity. The estimated profiles are shown in Fig. 3. With these estimated productivity profiles, we can examine the reconstructed revenues and note the data fit and residual error in Fig. 4.

Similarly, including attrition profiles of length $\nu_A = 12$ months and with attrition slope constraint parameter $m_A = 1/6$, we obtain the productivity profiles shown in Fig. 5. The revenue data fit with attrition is shown in Fig. 6.

In both sets of results, with and without attrition, we see that new hires with experience behave as 'old-timers' in that they are at steady-state productivity L from the outset. Also in both sets of results, new hires from universities are not productive at all in the beginning. In the first eight months, they have zero productivity before eventually having some productivity in the ninth month after joining the salesforce.

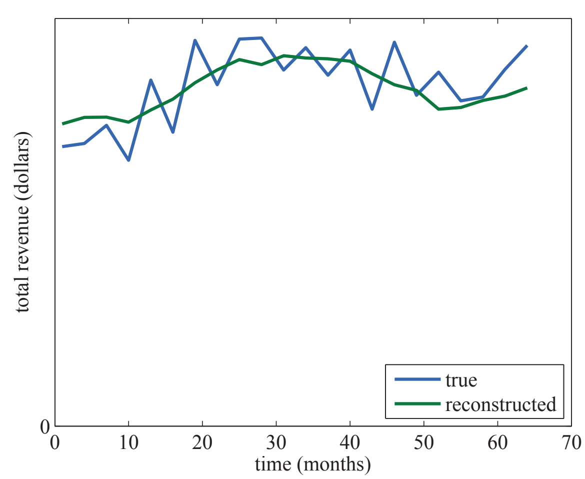

Figure 4. Revenue fit using estimated productivity profiles without attrition.

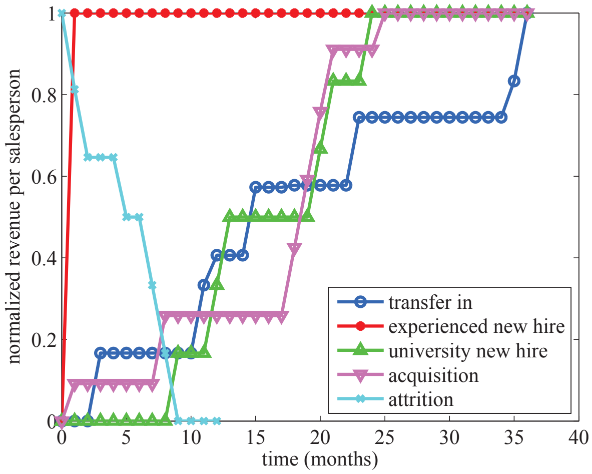

Figure 5. Estimated productivity profiles with attrition.

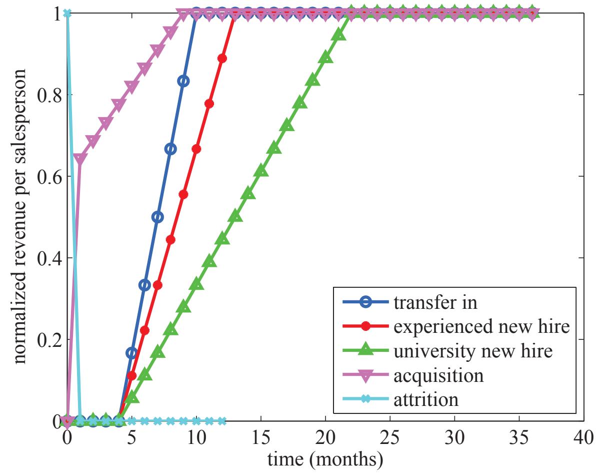

Figure 7. Productivity profiles elicited from leaders in sales organization.

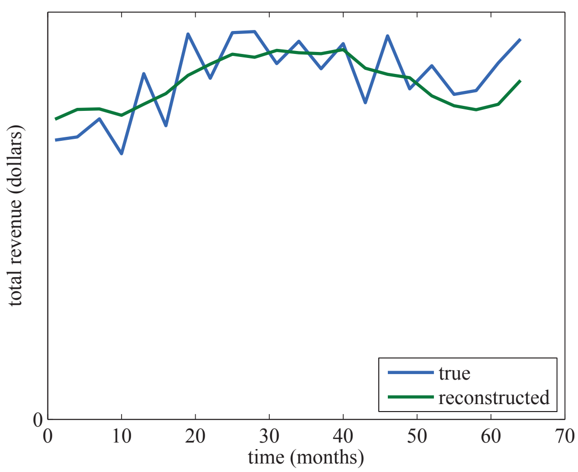

Figure 6. Revenue fit using estimated productivity profiles with attrition.

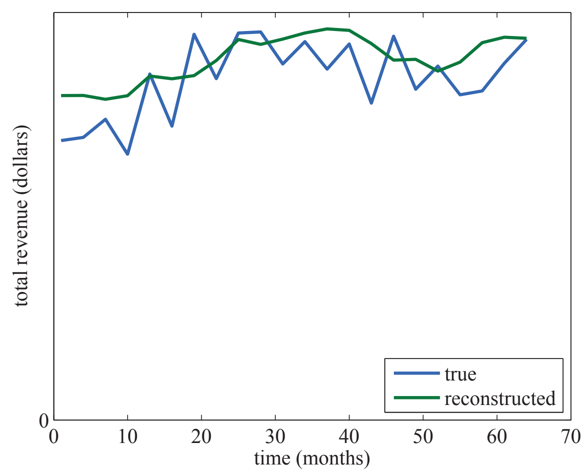

Figure 8. Revenue fit using productivity profiles elicited from leaders in sales organization.

Transfers in and those that join due to acquisitions begin with productivity greater than zero. Those from acquisitions reach the steady-state productivity L after approximately the same time as new hires from universities whereas those that transfer in reach L after new hires from universities. The attrition profile reaches zero after nine months. The residual error for both sets of regularized productivity profiles is not too severe and matches fairly well, and in fact smooths out some of the noise in the total revenue.

In addition to estimating productivity profiles from CRM data, we also elicited productivity profiles from leaders in the sales organization. These elicited profiles are shown in Fig. 7. The human experts indicated no residual effect of attrition sellers, and indicated four months of no productivity for new hires (experienced and university) and transfers in. They also indicated an intermediate initial capacity for acquired sellers. The ramping up period indicated by the sales leaders was slowest for university hires, followed by experienced new hires, and fastest for transfers in. With these elicited productivity profiles, we obtain the revenue reconstruction shown in Fig. 8.

The productivity profiles we obtain through data mining share many characteristics with the elicited profiles, such as the general shape of the profile, but also have certain key differences. One main difference is that for sellers joining the organization, the estimated profiles mostly have a longer period until the seller reaches full capacity. For example, university new hires have eight months of no productivity rather than four months, and transfers in barely reach full capacity after three years. On the other hand, experienced new hires reach full capacity very rapidly (within the first month itself), something that is quite counterintuitive as even experienced sellers would be expected to have some non-productive time. Similarly, the effect of a seller who attrits is felt much longer than the very short period estimated by sales leaders.

The minimax reconstruction error of the elicited profiles is 1.6280 larger than the data-optimized solution without attrition and 1.6673 times larger than the solution when attrition is included in the model. Thus through optimization, we recover profiles that are more reflective of the behaviors present in the data. The gut instinct of sales leaders is accurate to a point, but through business analytics, we are able to produce more refined and exact productivity profiles.

IV. Conclusion

In this paper, we have discussed the issue of seller productivity in sales organizations with constant flux of human resources. In particular, we have focused on temporal productivity profiles of sellers after HR events such as hiring and attrition, which have not received much attention in the literature from a quantitative or data mining perspective. Effectively modeling productivity as a function of time is essential for strategic as well as operational enterprise planning. With such models, it is possible for organizations to get an accurate estimate of the performance of a salesforce within a particular unit (such as geography or brand) and compare performance within units, set better sales targets for its salesforce, as well as try scenarios and ask 'what-if' questions such as the near-term effect of its hiring plans on revenue (e.g. what will the impact of hiring 50 university sellers be on revenue in the following quarter, or what kind of new-hire composition can help it meet sales targets in two quarters).

Towards developing post-event productivity profile models, we obtained enterprise CRM data of aggregate sales revenue and headcounts, including monthly hiring and attrition numbers. For this type of data, we developed a productivity model estimation procedure that minimizes the maximum residual error of aggregate sales revenue using linear optimization. The solutions we find through data mining yield smaller minimax error than productivity profiles elicited from leaders within the sales organization.

In future work, we plan to use the estimated productivity profile models for forecasting and prediction to inform hiring decisions. In doing so, we will examine the generalizability of the estimated profiles to future times. Another avenue of future work is to examine productivity profile estimation with optimization criteria other than minimax, for example least squares. We may also examine various other advanced formulations of regression.

From a business viewpoint, there are several avenues for future work as well. Firstly, we assumed that the same productivity profiles were applicable across all geographies, brands and job roles. However, this is likely not true. For example, a new teleseller (i.e. one who conducts sales via telephone) will likely ramp up differently than a client-facing seller as well as a technical seller. As such, we are planning on applying this methodology to learn different sets of productivity profiles for different geographies, brands and job roles. It will also be interesting to see if similar profiles are generated as we do this exercise, or if some of the differences observed between the learned and elicited profiles (notably experienced hires and transfers in) disappear.

Secondly, the productivity profiles elicited from sales leaders are currently being used within predictive models that are used within the organization for help with predicting expected revenues, planning and setting targets, and evaluating/comparing salesforce performance. We would also like to examine the use of the mined productivity profiles within these models and compare the predictions, both with those predicted by current models as well as actual results.

References

[1] D. N. Behrman and W. D. Perreault, Jr., "Measuring the performance of industrial salespersons," *J. Bus. Res.*, vol. 10, no. 3, pp. 355–370, Sep. 1982.

[2] M. Singh, D. Bhattacharjya, L. Deleris, D. Katz-Rogozhnikov, M. Squillante, B. Ray, A. Mojsilović, C. Barrera, D. Kakrania, J. Richard, A. Saha, and J. Fu, "The growth and performance diagnostics initiative: A multidimensional framework for sales performance analysis and management," *Serv. Sci.*, vol. 3, no. 1, pp. 82–98, Spring 2011.

[3] S. E. Sim and R. C. Holt, "The ramp-up problem in software projects: A case study of how software immigrants naturalize," in *Proc. Int. Conf. Softw. Eng.*, Kyoto, Japan, Apr. 1998, pp. 361–370.

[4] W. G. Zikmund, R. McLeod, Jr., and F. W. Gilbert, *Customer Relationship Management: Integrating Marketing Strategy and Information Technology*. Hoboken, NJ: Wiley, 2002.

[5] M. J. A. Berry and G. S. Linoff, *Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management*. Indianapolis, IN: Wiley, 2004.

[6] A. J. Morgan and S. A. Inks, "Technology and the sales force: Increasing acceptance of sales force automation," *Ind. Market. Manage.*, vol. 30, no. 5, pp. 463–472, Jul. 2001.

[7] S. Makridakis, S. C. Wheelwright, and R. J. Hyndman, *Forecasting: Methods and Applications*. New York, NY: Wiley, 1998.

[8] A. A. Zoltners and P. Sinha, "Sales territory design: Thirty years of modeling and implementation," *Market. Sci.*, vol. 24, no. 3, pp. 313–331, Summer 2005.

[9] P.-N. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*. Boston, MA: Pearson Addison Wesley, 2006.

[10] R. Lawrence, C. Perlich, S. Rosset, I. Khabibrakhmanov, S. Mahatma, S. Weiss, M. Callahan, M. Collins, A. Ershov, and S. Kumar, "Operations research improves sales force productivity at IBM," *Interfaces*, vol. 40, no. 1, pp. 33–46, Jan.–Feb. 2010.

[11] J. Barling and R. Beattie, "Self-efficacy beliefs and sales performance," *J. Organ. Behav. Manage.*, vol. 5, no. 1, pp. 41–51, Jun. 1983.

[12] M. R. Barrick, G. L. Stewart, and M. Piotrowski, "Personality and job performance: Test of the mediating effects of motivation among sales representatives," *J. Appl. Psychol.*, vol. 87, no. 1, pp. 43–51, Feb. 2002.

[13] H. M. Greenberg and J. Greenberg, "Job matching for better sales performance," *Harvard Business Rev.*, pp. 128–133, Sep.–Oct. 1980.

[14] G. L. Hunter, "Information overload: Guidance for identifying when information becomes detrimental to sales force performance," *J. Personal Selling Sales Manage.*, vol. 24, no. 2, pp. 91–100, Spring 2004.

[15] Z. Ton and R. S. Huckman, "Managing the impact of employee turnover on performance: The role of process conformance," *Organ. Sci.*, vol. 19, no. 1, pp. 56–68, Jan.–Feb. 2008.

[16] G. A. Churchill, Jr., N. M. Ford, S. W. Hartley, and O. C. Walker, Jr., "The determinants of salesperson performance: A meta-analysis," *J. Marketing Res.*, vol. 22, no. 2, pp. 103–118, May 1985.

[17] M. Sharma, A. Mojsilović, M. Singh, and K. R. Varshney, "A framework for sales force productivity profile estimation," in *INFORMS Annual Meeting*, Austin, TX, Nov. 2010.

[18] D. Telesca and L. Y. T. Inoue, "Bayesian hierarchical curve registration," *J. Am. Stat. Assoc.*, vol. 103, no. 481, pp. 328–339, Mar. 2008.

[19] D. A. Katz-Rogozhnikov, K. R. Varshney, A. Mojsilović, and M. Singh, "MCMC inference of the shape and variability of time-response signals," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing*, Prague, Czech Republic, May 2011, pp. 3956–3959.

[20] D. Bertsimas and J. N. Tsitsiklis, *Introduction to Linear Optimization*. Belmont, MA: Athena Scientific, 1997.