



Commercial Productivity

August 2006
Confidential White Paper
© Copyright 2006 McKinsey & Company

Lisa Ellis
Tom Herbig
Eric Kutcher
Stephen Moss
Nir Polonsky
Michael Yoo

Contents

| | |
|---|---|
| Introduction | 3 |
| Reclaiming growth from within | 4 |
| Commercial productivity in action | 6 |
| Case 1: Revving up account growth | 6 |
| Case 2: The route to rep excellence | 8 |
| Case 3: Attach rates, up close and personal | 9 |

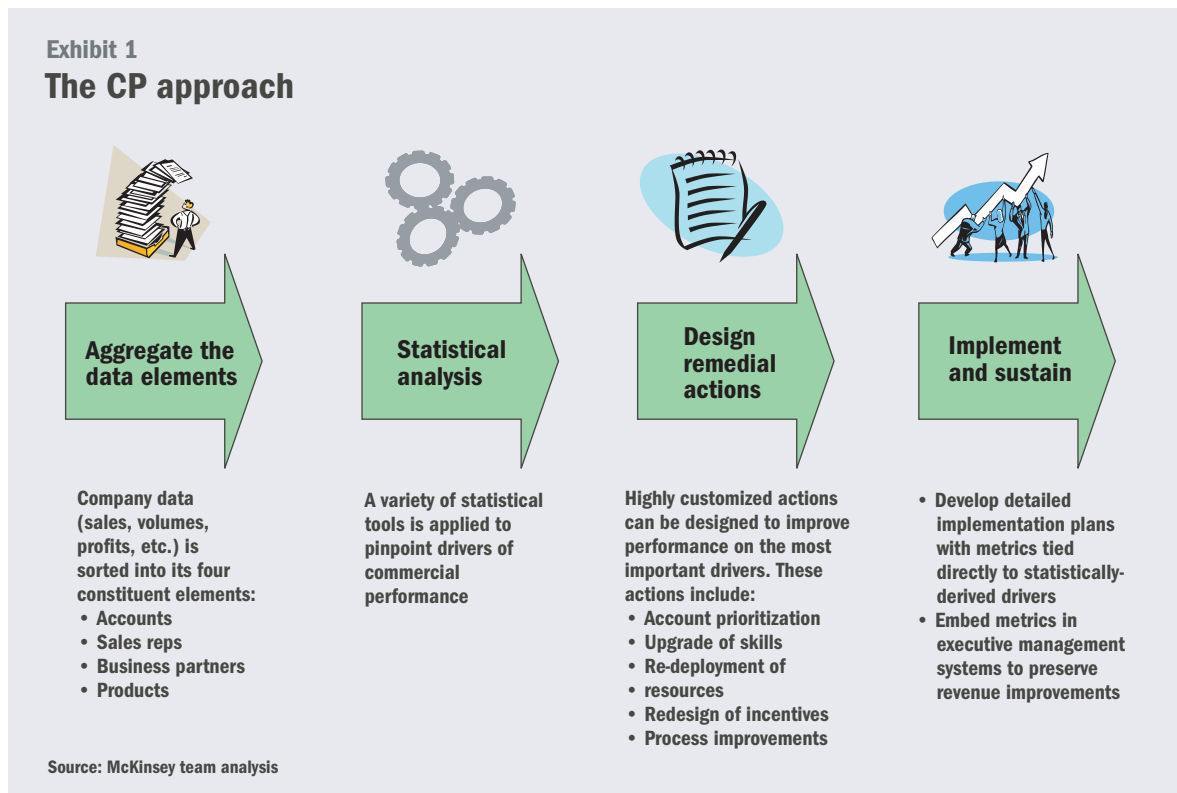
Introduction

Companies often look to new products and channels, as well as to acquisitions, to achieve the revenue growth they need. In the hunt for new revenue, one important point is often overlooked: for companies with complex sales and distribution, an excellent place to search for growth is actually inside the organization, especially by optimizing the performance of the direct and indirect sales forces.

Many companies contend that sales effectiveness

is largely limited by external factors beyond their control – customer industry, size, and geography, for example. On the contrary, our research shows that sales effectiveness is mainly driven by factors *within* a company’s control, and that most large sales organizations are leaving money on the table.

Commercial Productivity (CP), a new approach to improving revenue performance, relies on statistical analysis of detailed, transaction-level data to isolate the drivers of performance in four commercial areas: accounts, sales representatives and management, channel partners, and products. A company can then address the deficiencies identified in the analysis with specific actions to secure revenue improvements (*Exhibit 1*). Using this approach, many companies have found that they can break through the myths that undermine their revenue performance and achieve substantial near-term and long-term improvement.



Many people are now familiar with the statistical “magic” of *Freakonomics*, the economics best-seller by Steven D. Levitt and Stephen J. Dubner. What these authors do – explaining the drivers of a range of cultural phenomena, from a plunging crime rate to the performance of Sumo wrestlers in late rounds of tournaments – is essentially what **Commercial Productivity** does for revenue performance: leveraging statistical root analysis to explain observations, many of which run counter to conventional wisdom.

A distinct advantage of CP is that it leverages transaction-level data assembled from several sources within a company’s own data systems and thereby amasses a wealth of information rarely accessed holistically by sales organizations. CP then supplements this transactional data with sales-force survey data, tying actual revenue performance to specific policies, processes, skills, or behaviors.

Another advantage is that the work is self-funding, as some revenue improvements begin in 30 days and most can be captured in three to nine months. Companies that have taken up this work have enjoyed a 5%-to-20% improvement in return on sales.

Reclaiming growth from within

Many sales managers believe that they are already extracting near-maximum performance from their charges. And most subscribe to the belief that there are natural constraints that prevent sales reps from doing more. “That’s a bad territory,” managers say, or “That rep is stuck with a lot of bad customers.” Although bad territories and bad customers exist, these constraints are not nearly so numerous or extensive as sales managers imagine.

To better understand the nature of these constraints, companies can turn to a new form of statistical analysis, borrowed from mainstream academics and customized for use in real-world commercial enterprises. **Commercial Productivity** looks within performance and identifies those factors that can be influenced or controlled. The analysis then suggests structural changes and incentives that can help the company find growth from an unlikely source – its own sales organization.

To begin, CP identifies the key revenue elements upon which commercial performance rests: *accounts, sales reps, business partners, and products*. These are the elements companies typically measure to assess sales.

Next, for each element, CP uses statistical analysis to identify three forms of performance drivers (*Exhibit 2, overleaf*):

- (1) *structural factors outside a company’s control*, such as an account’s industry sector, size, and geographic region
- (2) *structural factors within a company’s control*, such as the number of sales reps per account or region, the amount of sales support, the tenure of a sales rep within an account, and the frequency and nature of promotions or partner incentives
- (3) *unnecessary variance*, often due to variations in behavior and skills, which can be influenced through changes in structure and incentives

The first set of factors is ignored. The second and third sets present opportunities for the company. Exhibit 2 shows how this separation is made in a company facing a typical problem: a wide variance in rates of account growth. The charts shows the

distribution of accounts by their rate of growth; most are clustered in the middle, while some at left are growing slowly or not at all, and those at right are growing quickly.

On this, we superimpose a representation of the relevant external structural factors on that distribution. What falls inside the shaded area is what we call necessary variance: performance differences caused by external factors that the company cannot easily alter. We then overlay a representation of the additional variance due to structural factors within a company's control. These differences are caused by the company's basic decisions that go into its sales model. What's left over is the residual or unnecessary variance that cannot be explained by any legitimate business factor. This variance is a result of differences in skills and behaviors within the sales force.

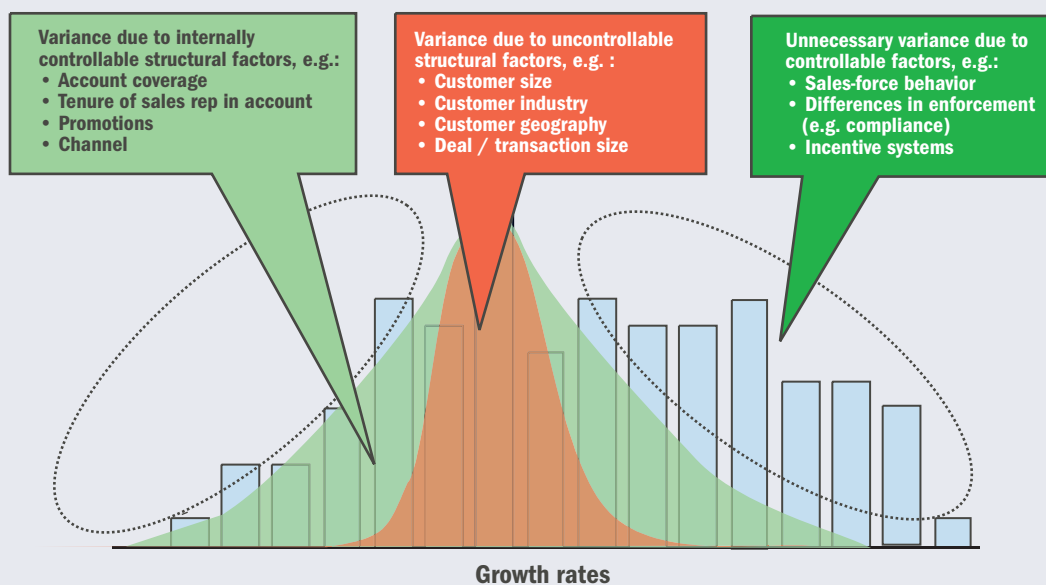
It is these last two areas – structural factors the company can influence and unnecessary variance – that are the foundation of CP's performance improvement potential. Our research shows that these factors often drive more than 70% of the variance in a company's sales effectiveness.

Each of these factors is then assessed for its contribution to revenue, revenue growth, and profitability. Having identified the sources of variability, the company can then design tactical actions to address these performance drivers. For structural factors within a company's control, we alter the structure (e.g., coverage or deployment), for example, to drive desired behavior; for unnecessary variance, we can among others things change incentives, institute training, and introduce a strong structure to compensate for differences in behavior and skills.

Exhibit 2

Identifying the controllable drivers of performance

Distribution of accounts by rate of growth Accounts



Source: McKinsey team analysis

CP can address inefficiencies and performance drivers that span the entire sales process, from customer coverage through opportunity generation and management and finally to pricing execution (*Exhibit 3*).

Commercial productivity in action

The essence of CP is the statistical testing of controllable factors, using detailed, transaction-level data, that yields an understanding of the commercial drivers. These statistical tests often yield a simple correlation – but the key to success is to continue with the analysis until the company finds a correlation so strong that it amounts to causality.

This is not a statistical wild-goose chase: Causality is often found only one level down from the original correlation produced. For example, the root cause of unnecessary variation in account growth might not be the number of reps per se (despite a promising correlation between number of reps and revenue growth variation) but rather managers' time with sales reps. CP allows companies to drill down and uncover more specific factors embedded in the data. Armed with the true drivers of performance, the company can take specific actions to enhance performance – and the additional rewards of doing so, as we shall see, can be substantial.

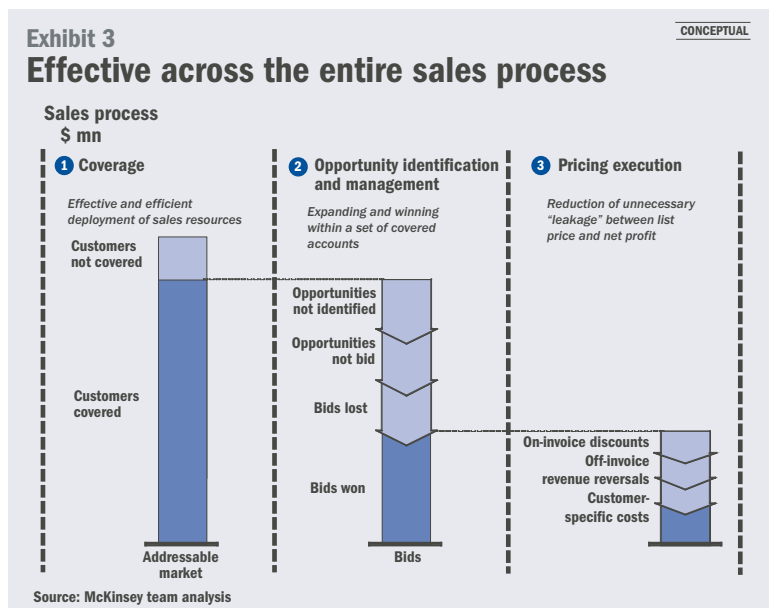
CP works on sales data that most companies collect and examines these across the entire sales fun-

nel, from coverage to pricing, as noted above. Three case examples, which touch on several of these areas, will help explain the sophisticated analytical machinery of **Commercial Productivity** and how it can be realistically applied in the field.

CASE 1: REWINDING UP ACCOUNT GROWTH

Consider a company with substantial variation in the revenue growth of its accounts. Account revenue was growing at some accounts by 70% annually, while at others it was falling by 20% a year – but why? Could these fast-growing accounts teach commercial managers anything that might help improve the performance of the others?

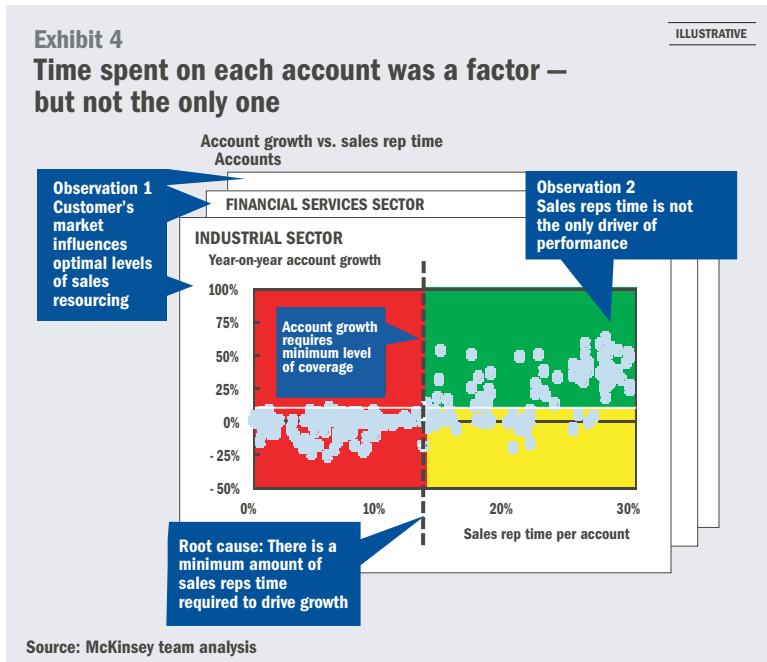
Many sales managers and most sales reps will contend that such variation relates mainly to structural factors beyond the company's control – the nature of the territory, say, or a particularly balky set of customers. To get past this kind of guesswork, the company hypothesized factors that



could lead to higher or lower growth. These included: number of sales reps per account, amount of sales support, and sales rep behavior, as well as end-use industry, customer size, customer geography, and the timing of a recent large purchase.

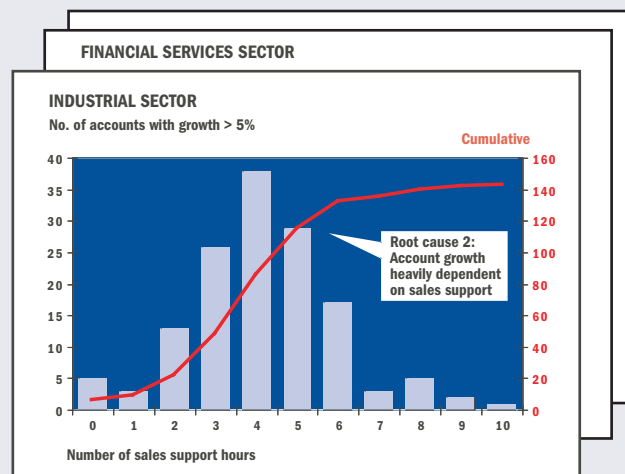
It then performed a multivariable regression to assess which of these factors had an impact on overall account-level performance. This regression showed that little of the observed variance was due to structural factors beyond the company's control, namely customer size and geography. Instead, most of the variance resulted from two factors that the sales force could affect: the number of sales reps per account (or, put another way, the amount of sales-rep time per account) and the amount of sales support (represented by FTEs per account).

In Exhibit 4, we see the result of the most revealing regression, which matched account growth vs. sales rep time per account. The analysis shows at a glance a root cause of variance: there is a minimum amount of sales rep time (about 13%) required for an account to grow. It also yielded two observations: (1) the customer's market matters when defining the optimal deployment to drive account growth; and perhaps more interestingly, (2) the amount of sales-rep time was a necessary but not sufficient condition (denoted by the green vs. yellow areas) – there was another factor critical to driving growth. This statistical signature, combined with interviews, confirmed that many of the company's accounts receive insufficient sales coverage to be able to grow.



Another regression unearthed a second critical factor or root cause (*Exhibit 5*). When the distribution of growing accounts was matched with the amount of sales support time, it became clear that a minimum level of support was also critical to account revenue growth. At the same time, there is a ceiling on the amount of sales support time re-

Exhibit 5
Sales support was a second critical factor



Source: McKinsey team analysis

quired to drive account growth. Anything beyond this level of support shows diminishing returns. In this case, this optimum level was about five hours per account.

Based on this CP analysis, the company identified and prioritized the actions it needed to take. These included: defining the number of sales resources needed by industry, limiting the number of accounts assigned to sales reps to allow them to spend time with the most promising ones, and supplementing reps' service to these promising accounts with more sales support.

CASE 2: THE ROUTE TO REP EXCELLENCE

In the previous example, we looked at accounts. Now, we turn our attention to sales reps – and by extension, sales management. In this case, the sales organization found that there was enormous variation in the performance of its reps, and it believed that the set of accounts a rep had been assigned (i.e., the rep's territory) was the principal factor determining performance – and the variation in revenue attainment.

On the contrary, a CP analysis soon found that at this company, all territories were created equal – i.e., each had the same revenue potential. A second round of CP analysis pointed to sales managers' span of control as the main reason for the huge variation; it found an inverse correlation between the two, as shown at left in Exhibit 6.

Still, the correlation was not yet established as causal, and so managers decided to investigate further: What was it about span of control that was driving variation? Testing – supplementing the data with a survey of sales and sales-management activity – yielded a much better idea of the time each frontline manager was spending on different activities in a given week or month. Sales reps reported that they highly valued the coaching they received from sales managers, and so managers decided to investigate that.

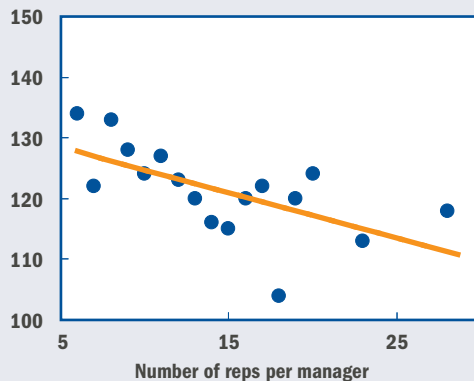
Exhibit 6 (*right-hand side*) shows the result of that inquiry. Revenue per sales rep depends very strongly on the amount of time the sales manager spends per quarter coaching on sales calls. In fact, there is a step-function improvement in sales-rep performance when a sales manager moves from

Exhibit 6

Variance in sales rep performance

Variance in performance is driven by span of control...

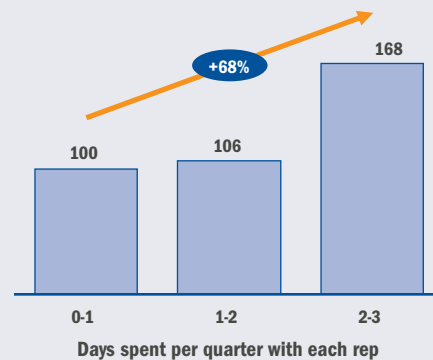
Average annual sales per rep vs. span of control
Sales reps, \$ thousands



Source: McKinsey team analysis

...and within span of control, managers' time with reps counts most

Sales per rep vs. time spent with managers
\$ thousands



1-2 to 2-3 days per quarter coaching an individual rep on sales calls. The root cause of rep performance was not span of control per se, but rather the time spent training reps in the field.

This, of course, represents a structural opportunity. By changing a sales manager's responsibilities, outsourcing some job elements to lower-paid administrative staff, and decreasing span of control, the company was able to increase the amount of time available for coaching sales calls. The question was how far to go. Too many reps per manager would render reps less productive; too big a jump in managerial head count would be inordinately expensive. Economic analysis – juxtaposing operating margin and span of control – helped determine the optimal number of reps per manager, and nine managers were hired to guarantee reps the training time they needed. This new, optimal span of control allowed the company to reallocate and increase its sales management resources profitably.

CASE 3: ATTACH RATES, UP CLOSE AND PERSONAL

The final example looks at the familiar problem of how to improve “attach,” or cross-selling, rates of key products. Measuring these rates is not new. What is new is that **Commercial Productivity** uses a statistical method, called “actual minus expected,” to address one of the most common concerns in sales.

One company was troubled by what it viewed as a low rate of cross-selling across its sales force. It used CP to determine that individual sales-rep behavior and skills are a key driver of attach rates and their improvement. The company created an “expected” performance – that is, performance normalized for legitimate business factors, such as products sold, account size, and size of deals. The difference between this expected performance and the actual performance of the reps represented the opportunity for improvement – significant for the

company, since 11 of 16 reps performed worse than expected.

CP then identified what was driving this variance in sales-rep attach rate performance – misaligned incentives, differences in rep skills (i.e., knowledge of the product to be attached), and differences of opinion as to the value of increasing attach rates.

Once again, inserting a robust structure, including revised incentives, to address this unnecessary variance was the key to improving overall attach rates. Now focused on the root causes of the problem, the company implemented a plan to change behavior by offering greater incentives for improving attach rates and increasing the profit contribution component of quotas.

While we focused in the above examples on the direct sales force, similar issues and opportunities exist in the indirect channel. For example, account growth and attach rate show similar, if not greater, variance in accounts that are primarily managed by VARs, distributors, and other business partners. The same issues – time and resources dedicated to the account, skills, and incentive alignment – resurface in the indirect case. Incentive misalignment can be particularly important to address in cases where business partners also sell competitors' products. However, addressing the root causes in the indirect channel – which is an external entity – tends to be more complex than in the direct case: fewer control levers are available and the issues must be addressed both with the internal channel organization and with external business partners, which have a broader set of incentives.

* * *

How can a company know whether CP is relevant to its situation? One way to find out is by answering several key questions. How complicated is your business – do you have more than 50 ac-

counts? Do you prioritize accounts or deploy resources for historical reasons, rather than analyzing customer potential? How well do you manage your network of business partners – especially those that sell a wide range of products? Is your sales management centrally controlled or highly

distributed (and possibly highly variable)? Do sales-rep incentives focus more on unit volume than on revenue? Is the range of product margins wide (e.g., more than 20 points)? Companies that answer “yes” to these questions may very well be leaving money on the table.

Lisa Ellis is an associate principal, **Tom Herbig** is a principal, **Eric Kutcher** is a principal, **Stephen Moss** is an associate principal, **Nir Polonsky** is an associate principal, and **Michael Yoo** is a consultant, all in the Stamford office.