Robert E. Tucker, "Is Coaching Worth the Money? Assessing the ROI of Executive Coaching," in Howard Morgan, Phil Harkins & Marshall Goldsmith (Eds.), The Art and Practice of Leadership Coaching: 50 Top Executive Coaches Reveal Their Secrets, New York: John Wiley & Sons, 2005. pp. 245-254.

CHAPTER 11

Is Coaching Worth the Money? Assessing the ROI of Executive Coaching

Any person looking at the value of executive coaching must bear in mind that its practices are changing alongside broader business practices. Since the recession of 2001, much conventional business wisdom has become open for reexamination. The cult of the strong, forceful leader, for instance, has been replaced by a more humble model. At such a turning point, we need to examine not only how coaching has worked in the past but where it is headed in the future. This does not give us a reason to avoid exploring coaching best practices and ROI, but it does provide a caution that we are unlikely to reach any final answers in the near term. As Fredrick Reicheld recently commented, the careful study of business is a relatively new thing. We write this book at the beginning of the study of executive effectiveness—not at the end. The coaches represented in this book are pioneers. It is only in the decades ahead that we will come to a full understanding of the ground that they cover. Some would see this as daunting. We see it as exciting.

What Is the Business Impact of Executive Coaching?

We know that executive coaching is increasingly popular. On its own, this fact would tend to argue for its value—companies don't have a habit of throwing away money and those that do rarely live to throw it away for long. Still, there are nagging concerns about just how much coaching adds back to the bottom line. To borrow an argument from the Nobel Prize-winning economist Robert Solow, there is a fundamental difference in the return on investment in, say, a new factory and a new recreation room.[1] Both, technically, may contribute to increased productivity. One does so rather more directly.

Which is coaching? Is it more like investing in the new factory—investment that makes a direct and measurable impact on financial performance? Or is it more like buying the new rec room—helping the bottom line because happy employees tend to be productive employees? The answer has everything to do with how coaching is done. The available evidence suggests that many companies are investing heavily in coaching hoping to get the first kind of return (the factory) and actually getting the second (the new rec room). In this section, we will examine some of the research done over the last few years that helps to clarify how companies can get the most economic impact from their investment in coaching.

First, we need to look at what research has been done and separate the wheat from the chaff. Ten minutes on the Internet will prove that there is no shortage of what calls itself research on coaching. To be charitable, most of this material is not useful. Anecdotal stories of individuals who found greater effectiveness through their relationship with a coach, or elaborate case studies of coaching effectiveness told from the standpoint of the coach don't advance our understanding of the root effectiveness of coaching as a practice. What we need are careful quantitative studies of actual coaching engagements. We want to ensure that we are measuring the effectiveness of coaching and not of some other learning experience that might be going on while coaching is underway. We need a better sense of the periods in an executive's life when coaching is most beneficial. And we want to determine the relationship between what might be called customer satisfaction, and actual performance improvement.

How do we determine that it is coaching that is leading to performance improvement, and not some other activity, experience, or pressure? Particularly in situations where coaching is encouraged on an executive either by his or her superiors or by some broader part of the organization, we need to make sure that it is the coaching that is helping the executive improve, and not just his or her awareness of the expectation of change. Put on the spot, most executives are intelligent enough to temporarily change their performance in desired ways. What is less certain is that this forced change leads to permanent behavioral change. One of the claims made by effective coaches is that the coachee is learning new *habits*—that is, that the points of discussion and the lessons learned will be lasting. Good research allows us to zero in on the impact of coaching—and coaching alone—an important question for determining the return on investment.

Second, good research will reveal quite a bit about the timing of effective coaching. In coaching, as in life, timing is everything. Three questions are important in this regard. What is the best time to expose an executive to coaching? Are there windows in an executive's career when the organization would be earning better ROI than at others? Finally, are there unique periods when coaching is unlikely to be productive? These questions can only be answered by carefully studying how coaching works in large numbers of cases across a variety of companies.

Related to this is the question of performance improvement lag. Briefly stated, lag has to do with the amount of time that passes between the coaching event and the performance improvement. Adult education tends to have quite a bit of lag built into it. After learning something, adults tend to digest it—to make it their own and put it into practice in their own environments. The actual performance impact of a useful developmental experience, then, will likely start off low, gradually climb and then diminish as the executive wrings the lasting benefit from that experience. Well-conducted research on coaching will give a better sense for what these lags might be and how we can design our coaching interventions to impact our organizations as much as possible.

Finally, good research will allow us to differentiate between improved effectiveness as an outcome of coaching and mere customer satisfaction. Think, for a moment, about our above example of the recreation room. If asked, employees may say that their satisfaction is enhanced by a new recreation room. They may even report that they are more productive as a result of increased relaxation, happiness, or other related factors. In the end, however, it is not their own conclusions that are definitive, but actual hard data showing an increase in productivity. Stated bluntly, people can be wrong about how much more effective they are becoming. If people were able to accurately assess and remedy their own performance, there would be less need for executive coaching. In other words, determining whether or not executives are happy with their coach is not the same thing as carefully measuring the improvement each executive achieves as a result of coaching. Executive satisfaction will tend to measure, among other things, how well the coach and executive get along. At best, questioning the coachee will obtain his or her self-evaluation of increased effectiveness. We shouldn't lose sight of the fact that this runs the risk of significant bias. Good quality research will keep this fact in mind.

What's Out There Now: Early Research on the Value of Coaching

What would quality research look like, and does it exist now? To answer the first question, we should make a list of the qualities that the best research would have. Once we have such a list, it will be much easier to evaluate what's out there.

First, we are looking for efforts to measure the exact impact of coaching in a relatively large sample. The larger the sample, the more confident we can be that we are measuring actual effectiveness, as opposed to something that might be happening for local reasons. We also become more confident that best practices work from company to company.

Second, we'd love to find research where participants are selected randomly. This is important because it will help us avoid something called selection bias. Selection bias occurs when individuals are selected for possessing some initial quality or attribute; inevitably, subsequent measurement efforts pick up this quality or attribute. An example might be selecting a company's high potentials, for coaching and then comparing them after coaching with employees not on the high potential list. Understandably, you'll never be sure whether you are measuring the effect of the coaching or the fact that they were already high potentials. Selection biases can be very subtle but can have dramatic effects. The most well-known statistical example concerns a headline from the 1948 presidential race. On Election Day, the *Chicago Tribune* famously printed the mistaken headline "Dewey Defeats Truman," despite the fact that Truman had actually won the race. What happened? The *Tribune* relied on a poll in which participants were drawn from Department of Motor Vehicle records. In 1948, however, people who owned cars were a bit wealthier than the average American. The poll sampled these slightly wealthier individuals, who voted for Dewey slightly more than the actual rates in the population. Because the sample had selection bias, the results were useless—and quite embarrassing for the *Chicago Tribune*.

There are many contexts where true random sampling is difficult or impossible to achieve. In a coaching setting, randomness is unlikely to make sense. You give coaching to people who need it and avoid it for people who do not. Although it is true that the expense of coaching as a practice makes random sampling less feasible, we should look to cases where researchers have done what they can to avoid selection bias. At a minimum, what we would hope for is some comparison between a group of managers subject to executive coaching and those who weren't. We'd want these two groups (those selected for coaching and those not) to be as similar as possible in terms of rank and context. We'd then want to define what we mean by coaching as clearly as possible—making it easier to figure out what exactly was working if we find any effect. Is there research out there that meets these demanding standards? Some.

We chose to look at one study that is representative of good early research on the subject. It appeared in the journal *Personnel Psychology* in the spring of 2003. In this study, the authors looked at the effectiveness of 360-degree feedback when participants met with a coach versus the effectiveness of the 360-degree feedback alone.[2] The authors found that individuals who met with a coach to discuss their feedback and their action plan were more likely to set specific goals and achieve them in the year ahead. What is good about this research is its specificity. At the end of reading it, we know what has and what has not been demonstrated, and thus, what does and does not work. Other research that would fit this bill includes such work as:

- Gerald Olivero, Denise K. Bane and Richard E. Kopelman, "Executive Coaching as a Transfer of Training Tool: Effects on Productivity in a Public Agency" that appeared in *Public Personnel Management*, vol. 26, no. 4 (1997), pp. 461–469.

  Summary: Describes the advantages of one-on-one executive coaching in positively influencing transfer of training. Examines the effects of executive coaching in a local government agency. Thirty-one managers took a management development program, followed by eight weeks of one-on-one executive coaching. The study finds that training increased managerial productivity by 22.4 percent, while coaching increased productivity by 88 percent.
- Andrea D. Ellinger, "Antecedents and Consequences of Coaching Behavior," *Performance Improvement Quarterly*, vol. 12, no. 4 (1999), pp. 45–70.

  Summary: Discusses the use of coaching to facilitate the development of learning organizations. Presents the results of a study to determine the outcomes of coaching interventions. Finds that managers' commitment to coaching can impact employee, manager, and organizational performance.
- Carol Patton, "Rating the Returns," *Human Resource Executive*, vol. 15, no. 5 (April 2001), pp. 40–43.

  Summary: Outlines a nine-step ROI process that determines the value of executive coaching. Claims that this process must be applied consistently through the organization. Includes a list of measurement tools and important ROI measurements.

What we see when we look at this research is a fairly compelling case for the cost effectiveness of executive coaching. We would add that this effectiveness can be enhanced if a number of careful steps are followed. First, the coaching should be as specific as possible in terms of its goals. This doesn't rule out what we would call opportunism—cases where we find something worth working on in the course of our engagement—but it does help to have a specific set of goals agreed to early on in the engagement. Second, there should be some sort of agreement about what will count as meeting those goals. If the executive has a problem motivating team members, an agreement should be reached prior to the engagement about how a change in this skill will be measured. This system of measurement need not always be quantitative. Some of the best coaches rely far more on interview-based feedback. What is important is that this agreement be reached before the engagement begins. Otherwise, there will be the temptation to define success based on where you end up rather than on where you wanted to go. Third, getting to know your coach as well as you can before an engagement is a good idea. Ask questions. Talk about results. Evaluate the decision to hire a coach as carefully as the decision to buy a piece of major equipment. Do you know the brand? Do you need the features, and do they fit in with other aspects of your business?

The sheer diversity of coaching research tells us where we are in the developmental phase of the industry. Like the automobile industry in the early twentieth century, we are in a situation where the need for the product has become clear, but where the exact form that product will take is still very much under development. What is important in this development phase is reliability of the product and reputability of the person or company providing the service. In the next 20 years, we will see a consolidation of the coaching field around accepted best practices and a further refinement of the research studying coaching. For now, the best advice for those who would seek coaching is to know your practitioners; evaluate his or her reputations; and have detailed conversations about expectations before the engagement begins.

Conclusions and Recommendations

What becomes clear when we look carefully at the existing research is that asking the question "Is coaching worth the money?" is like asking the question "Is training worth the money?" The question itself is too broad. Unless you know quite a bit about the proposed intervention—what are its specific goals? who will do the training? how will the success of the training be measured?—you will not know enough to answer the ROI question. Coaching *can be* worth the money, but the consumer must purchase wisely. The two best reasons that coaching *can be* worth the money have to do with (1) the disproportionate influence exerted by executives and (2) the increased effectiveness of one-on-one training.

The Disproportionate Influence of the Executive

Many studies of coaching ROI assume what we are calling the disproportionate influence exerted by executives. The argument goes something like this:

John Smith, executive, controls the fate of a company that generates $5 billion annually. A 1 percent improvement of his effectiveness, it is suggested, will translate into a $50 million dollar savings for the company. Even if the rate of return is less (or even far less) than the estimated $50 million, the savings would still be substantial—and worth it for the organization. Small increases in executive effectiveness are able to have large effects on organizational performance. This makes executive coaching appealing because a little effort can have a large impact. Although it's probably a good idea to be cautious about the wilder claims made on the basis of this assumption (it's not uncommon to see a 100 to 1 return on investment claimed for coaching), there is certainly something to this rationale.

Two suggestions are important here for getting the most bang for the buck. First, coaching should have a discernable effect on business-critical skills. If the executive is regarded as too abrasive, for example, the net result of the coaching should be a reduction in the perception of abrasiveness on the part of the people with whom the executive works. If the issue is an unwillingness to engage in conflict, colleagues, superiors, and subordinates should regard the executive as better able to engage in effective conflict at the end of the coaching engagement. If all of this sounds very obvious, it is. Nonetheless, most coaching engagements are not structured so clearly. Accordingly, even though the executive has an enormous amount of influence, the coaching engagement often becomes too watered down to have concrete business impact, and thus any appreciable ROI. As many of the coaches in this volume will tell you, one key way to enhance return on investment is specific and prearranged goals and a way of measuring whether or not those goals have been accomplished. This can take the form of a written action plan with specific follow-up, or it can take the form of an agreement between the coach and the executive to find some way of measuring progress.

Second, coaches and executives alike should be aware of what could be called nonlinear skills or attributes. A nonlinear skill or attribute is one where one of two things is true: either a small change will have a dramatic effect on the effectiveness of the executive, or a large change will have a negligible effect on the effectiveness of the executive. The first case is called an amplifier, where a small amount of improvement facilitated by the coach has a large effect on performance. The second is called a damper where a coach can work and work and not make much of a mark on the effectiveness of the executive.

Coaches and executives should be on the lookout for these amplifiers. A good example of this is some engagements we have seen that have centered on effective conflict management. Because conflict is, in most cases, somewhat rare, and because most people have a negative emotional response to conflict, small improvements regarding conflict management skills tend to result in large changes in perceptions of effectiveness. Indeed, too much change and the perception of the executive's effectiveness can go down—we start running into concerns about assertiveness. Amplifiers help to achieve good ROI because they are cases where a little bit of work by the coach and the executive result in quite a bit of perceived improvement.

Dampers, on the other hand, are pitfalls to be avoided. These are skills or attributes in which large amounts of work can be invested, and actual change is achieved, but there is little perception of increased effectiveness. A good example of this might be work we have seen in increasing the organization skills of an executive. Even relatively large shifts in an executive's ability to be organized are often not noticed by his or her peers or superiors. This is true for a number of reasons. First, it is assumed that executives are going to be organized to at least a certain level. Below this level, even improved performance is unlikely to earn much credit. Second, the skill is often seen as cost-effectively delegated. If a coaching engagement will cost more than other solutions that are more likely to lead to effectiveness, such as delegating more authority to an office manager, then coaching won't make a great deal of sense. Self-esteem probably works like a damper as well—it takes quite a bit of increased self-esteem to make a noticeable impact on others. This has led many of our authors to be wary of deeper psychological issues. They are more likely to require large amounts of work for what may well be only small improvements. This is not an argument against an executive working through these issues. It is an argument that says that a coach might not be the person who can help most effectively.

The Effectiveness of One-on-One Training

Second, much of the research seems to assume that executive coaching works as a form of one-on-one training. The claim goes something like this: since we know that training in a group setting works, and since executives tend to be too busy to participate in group training, one-on-one training (which we'll call coaching), will work as well. This makes sense. Just as a one-on-one session with a teacher is likely to get certain concepts and practices across, so too the one-on-one relationship established as part of the coaching process is an effective way to impart new practices and ideas.

A couple of guidelines are important here. First, we have to be careful to include enough people in the coaching process to get the advantages of both the classroom and the one-on-one session. Many who have experience in group educational settings can tell you that some of the most valuable lessons come not from the teacher, but from the other students in the class. This unexpected benefit often comes from the quiet members of the class—the one's you'd least expect to provide a frame-changing comment. Coaching can be similar. We need to make sure to include enough input to allow for surprises, both for ourselves as coaches and for the executive. This is a good argument for the inclusion of interviews in at least some component of the coaching process.

- ❑ As an organization, are you committed to coaching as a process rather than just an event?
- ❑ Is the coachee's immediate supervisor committed to the coaching process?
- ❑ What are the types of changes that you hope will result?
- ❑ Have you established internal measurements to identify when you have achieved success?
- ❑ Do you have benchmarks on those measures to identify the baseline?
- ❑ Do you have a control group identified?
- ❑ Are you using the right period of time to properly achieve the results you are looking for (at least 18 to 24 months)?
- ❑ Have you considered indirect measures such as employee satisfaction or turnover?
- ❑ Are you measuring the coach on the results that the coach achieves or the time that the coach spends?
- ❑ Have you ensured that one of the measurements is perceived improvement, as viewed by those who work with the coachee on a frequent basis?
- ❑ Based on everything that you know about the coachee, is there a reasonable probability that the coachee will change?

FIGURE 11.1 Coaching Checklist: Enhancing ROI

Second, we need to be a bit careful about the fundamental differences between training and coaching. Training assumes a lesson in place and ready to go. It assumes that there will be a more or less effective reception of this lesson and an application after the fact. Coaching often works differently. The coach works to facilitate progress in a number of areas, but it is the experience of the executive that serves as the foundation for any lessons learned. The coach needs to walk a thin line between too little and too much structure—between adding value and turning the coaching discussion into a monologue. Good coaches have great skill at walking this delicate line.

Conclusion

A brief look at the quality research available on the ROI of coaching gives us grounds to be optimistic. Measuring impact remains an underexplored issue; and yet there are solid reasons that coaching can make a significant impact on the bottom line of our organizations. Despite the significant cost associated with executive coaching, such costs can be carefully and wisely managed to the benefit of all concerned. Coaching is a good way to improve individual executive effectiveness. When planned and executed well, it will likely lead to more effective organizations.

Figure 11.1 on page 253 provides a useful coaching tool for enhancing ROI.