

LEADING TEAMS

SETTING THE STAGE FOR
GREAT
PERFORMANCES



J. RICHARD HACKMAN

HARVARD BUSINESS SCHOOL PRESS

Copyright 2002 Harvard Business School Publishing Corporation

All rights reserved

Printed in the United States of America

06 05 04

5 4 3 2

Requests for permission to use or reproduce material from this book should be directed to permissions@hbsp.harvard.edu, or mailed to Permissions, Harvard Business School Publishing, 60 Harvard Way, Boston, Massachusetts 02163.

Library of Congress Cataloging-in-Publication Data

Hackman, J. Richard.

Leading Teams : setting the stage for great performances / J. Richard Hackman.

p. cm.

Includes bibliographical references and index.

ISBN 1-57851-333-2

1. Leadership. 2. Teams in the workplace. 3. Performance. I. Title.

HD57.7 .H336 2002

658.4'092—dc21

2002000757

The paper used in this publication meets the requirements of the American National Standard for Permanence of Paper for Publications and Documents in Libraries and Archives Z39.48-1992.

The second account invites a group-level interpretation: The *crew* got itself into trouble by not paying attention to changes in the situation, by not planning and organizing the team's work, and by poor coordination between members. Focusing on the crew also highlights the captain's team leadership—not so much his behavioral style, but the kinds of expectations about behavior that he had established and enforced. Such matters are unlikely to be revealed by the first account, which implicitly casts Phil in the role of savior. The second version of the story also raises questions about the airline's overall strategy for structuring and supporting its crews. Might there have been something about the way crews were set up and staffed at that airline, or about the amount of authority they had for managing their own affairs, or about company-specified performance routines that could have contributed to what happened to Charlie and Phil?

To answer such questions requires that we move beyond explanations of team performance that rely mainly on the attributes or behaviors of individuals, and instead focus on how the teams themselves are designed and supported. It may sound silly to say, but if you are going to lead a team well, you must first make sure that you actually have a team to lead—and that you then deal with it as a *team* rather than as a set of individuals.

ESSENTIAL FEATURES OF REAL TEAMS

Real work teams in organizations have four features: a team *task*, clear *boundaries*, clearly specified *authority* to manage their own work processes, and membership *stability* over some reasonable period of time.⁶ The first and perhaps most important task of those who create or lead work teams is to make sure that these four essential features are in place.

Team Task

Everybody would agree that a string quartet is a team. Its work simply cannot be accomplished without all four members present and playing. Everybody also would agree that a collection of people on a street corner waiting for a traffic light to change is not. Beyond such extremes, however, the terms *group* and *team* are akin to projective tests: People read

into them what they wish, and conversations about teams can be frustrating because people have different things in mind when they talk about them.⁷

The first group I encountered in my research career was in the central office of a telephone company. My colleague Ed Lawler and I were being shown around the company in preparation for some research we were planning. One group we visited was introduced to us as "Supervisor Szczarba's team." Arrayed before us were a dozen or so telephone operators, each at her own console (there were no men in the group), each talking to her own customers, each taking her break at a time negotiated with Ms. Szczarba. Members of the group did perform their work alongside one another. But the only thing they actually shared was Ms. Szczarba.

Such teams are called *co-acting groups*. It is easy to tell who is in a co-acting group because members usually work in proximity to one another and have the same supervisor. But each member has an individual job to do, and that job's completion does not depend on what the others do. Co-acting groups are barely groups at all, and they are not what we are concerned with here. Our focus is on groups whose task requires them to work *together* to produce something—a product, service, or decision for which members are collectively accountable and whose acceptability is potentially assessable. The kind of outcome produced is not critical; it can be a physical product, a service, a decision, a performance, or a written report. Nor is it necessary that the outcome actually *be* assessed; all that is required is that the team produce an outcome that can be identified as its product and that it be theoretically possible to measure and evaluate that product. If Ms. Szczarba's team had been assigned collective responsibility for handling all service requests for, say, a specific part of the company's service area, and if members had been held collectively accountable for how quickly and well those requests were processed, then it would have been a real work team.

A great deal of organizational work is performed these days by sets of people who are called "teams" but who really are co-acting groups. Managers in organizations where this is done may harbor the hope that they can harvest the widely touted benefits of teamwork while continuing to directly manage the behavior of individual members. That hope is misplaced: If you want the benefits of teamwork, you have to give the *team* the work. So there is a choice here: Either design the work for a team, or design it for individuals. If done well, either strategy can yield fine results.

What is not fine is to send mixed signals: to use the rhetoric of teams when the work really is performed by individuals, or to directly supervise individual members when the work really is a team's responsibility.⁸

Sometimes managers' choices about work design are constrained because some tasks can only be done by teams (such as the string quartet mentioned earlier, or carrying out a multiparty negotiation) whereas others are inimical to teamwork. More often, however, the work could be carried out either by an individual or by a work team. In such cases, wise managers do not rush to assign the work to a team. Instead, they carefully weigh the advantages and disadvantages of using a team to perform the work, and take care not to assign to a team work that actually would be better performed by an individual.

Creative writing is one type of task that is often assigned to a team but should not be. So far as I know, not a single great novel, epic poem, or symphonic score has ever been written by a team. This kind of work involves bringing to the surface, organizing, and expressing thoughts and ideas that are but partially formed in one's mind (or, in some cases, that lie deep in one's unconscious), and these tasks are inherently better suited for individual than for collective performance. Even committee reports—mundane products compared with novels, poems, and musical scores—invariably turn out better when written by one talented individual on behalf of a group than by the group as a whole working in lockstep.

The same is true for many aspects of executive leadership. For all the attention being given to top-management teams these days, my reading of the management literature is that successful organizations almost always are led by a single talented and courageous human being rather than by a team, no matter how many well-qualified members the team contains. As numerous senior managers have discovered, bucks rarely stop at teams. Moreover, the most engaging and powerful statements of corporate vision invariably are the product of a single intelligence, set forth by a leader willing to take the risk of establishing collective purposes that lie just beyond what others believe to be the limits of the organization's capability.

Beyond creative writing and executive leadership, there are many other kinds of tasks that are better done by individuals than by teams. It is a mistake—a common one and often a fatal one—to use a team for work that requires the exercise of powers that reside within and are best expressed by individual human beings. A manager's first responsibility in

creating a work team, then, is to make sure that the work to be done is appropriate for team performance and that it requires members to work together interdependently to achieve an identifiable collective outcome. If that cannot be done (and many times it cannot), then the wise choice is to design and manage the work for individual performers rather than for an interacting work team.

Bounded Team

To work well together, team members need to know who they are. Members are sure to run into difficulties if there is so much ambiguity about who is actually on the team that they cannot reliably distinguish between the people who share responsibility and accountability for the collective outcome and others who may help out in various ways but are not team members. Having a clearly bounded team does not mean that members must do all their work in the same place at the same time, nor does it mean that membership cannot change as circumstances change. It merely means that members know who is actually on the team—a seemingly simple matter but one that trips up a surprising number of teams.

A task force at a financial institution where I once did some research got in trouble for exactly this reason.⁹ The team's task was to develop a plan for standardizing certain policies and practices across multiple locations that had, over time, developed their own ways of doing things. The team was formed by a senior vice president who also was its primary client. To ensure that all major points of view were represented, the vice president appointed to the team representatives of several key functional areas in the organization, including a few individuals he knew and trusted from other countries where the company did business. He further specified that the team could recruit to its ranks members from additional functions or locations if they would bring useful perspectives to the team's work, and he encouraged the team to work closely with professionals from a consulting firm who were on site to help the firm focus and streamline its offerings.

The first meeting of the task force was held at company headquarters and was attended by most of the functional area representatives but no one from a nonheadquarters location. Subsequent meetings, also at headquarters, were attended by different subsets of members, sometimes including

overseas representatives who happened to be in town. These meetings were frustrating to the task force leader and to other regular attendees because members who were attending for the first time invariably suggested new ideas or raised for consideration possibilities that had been considered and dismissed in a previous meeting. On a couple of occasions, a senior officer not on the team came to a team meeting to advocate a particular idea. One or two of the outside consultants also occasionally attended, although they behaved more as observers than as active participants. About a month into the work of the task force, the team leader took a trip to Europe to talk with overseas members whose contributions up to that point had been made mainly by telephone or e-mail. He returned with yet another set of new ideas and perspectives for the headquarters-based members to consider.

The task force never jelled as a performing unit. Members never developed the specialized roles and shared norms of conduct that characterize real work teams, nor did the team develop the collective momentum that would have enabled it to generate a high-quality proposal satisfactory to the team's client. Eventually, the team leader got so frustrated that he went off by himself and, with input from a couple of other members, personally wrote the proposal and submitted it to the senior vice president. There never actually was a team product because there never actually was a *team*.

The difficulties encountered by this task force are common when a team's boundaries are so unclear that its membership is uncertain, or when they are so permeable that there is a never-ending flow of people in and out of the group. Organizational psychologist Clayton Alderfer refers to such systems as "underbounded." The primary threat to underbounded systems, Alderfer shows, is that they "become totally caught up in their environmental turbulence and lose a consistent sense of their own identity and coherence."¹⁰ It is virtually impossible for an underbounded team to develop and implement a coherent strategy for carrying out a piece of work.

The reverse state of affairs also can occur. A team with tight, impermeable boundaries is what Alderfer calls an "overbounded" system. An overbounded team is something of an island unto itself. Although membership and team identity are completely clear, the team is closed off to its environment. Members risk overlooking significant environmental changes, they have limited capacity to respond appropriately to such changes even

if they are noticed, and they do not engage in the kinds of cross-boundary exchanges that can be critical to team performance.¹¹

The risk of overboundedness is real, especially for geographically isolated teams such as a new plant startup far from headquarters or a consulting team that operates for an extended time at a remote client site. Underboundedness, however, is the more common problem for organizational work teams. The dysfunctions that stem from unclear or excessively permeable team boundaries are especially prevalent for professional and managerial teams such as the financial institution task force described earlier. Curiously, managers generally give less attention to clarifying the boundaries of teams made up of their peers or immediate associates than they do to creating clear, moderately permeable boundaries for front-line teams composed of rank-and-file organization members.

A few years ago, for example, I was invited to meet with the top-management team of a large insurance company. The team consisted of about two dozen senior executives, each with his or her own part of the business to run or staff support function to lead. The chief executive officer (CEO) told me he was getting increasingly frustrated as he tried to get his team to take on collective responsibility for the well-being of the company as a whole. Team members, for their part, were just as unhappy with the team's once-a-month meetings as was the CEO. Everyone wanted to know what it was that was keeping this collection of highly talented and committed individuals from pulling together as a team.

But were these senior managers actually a team at all? From one perspective, they were merely co-actors: Each member's main responsibility and accountability was for the performance of his or her own organizational unit. Although it surely was a good idea for the firm's senior managers to get together regularly, it was not clear to any of them just why they were meeting or what they were supposed to *do* as a team. One thing was certain: The CEO did not want them to actually run the company. Although that assuredly would have turned the group into a real task-performing team, running the organization was *his* job.¹²

Nor was it clear exactly who was on the team—there were so many members that it was hard for them to keep track of one another, especially since some members would send deputies to meetings they could not personally attend. This top-management team had neither a clear

identity nor a shared and coherent view of its purposes and processes. It was an underbounded set of people who had meetings rather than a real team that performed work.¹³

Teams of managers and professionals are at special risk of underboundedness because their main work invariably involves extensive and often intensive engagement with a variety of other individuals and groups. That necessity can blur the boundaries of the core group, as Martine Haas discovered in her research on project teams that performed knowledge-intensive work for a large international development organization.¹⁴ When Haas asked managers for membership lists of the teams she would be studying, she often would get a response along the lines of "Well, that's not entirely clear—it depends on how you want to define the team." Even individuals who held team leader roles frequently could not give her a definitive list of members. That there were certain individuals who were "core" members of the team was clear to everyone. But beyond the core, membership was ambiguous because some individuals were partial members, others provided information or resources to the team but were (perhaps) not actually on it, and still others were clearly not members but just as clearly were critical to its work. It was as if there were an undifferentiated gradient extending out from the core of the team to all manner of other organization members whose participation was, in varying degrees, important to the team's performance. Haas eventually solved her research problem by distinguishing between the core members and everyone else, but the organizational problem she inadvertently surfaced remained a potentially significant impediment for some of the teams she studied.

Teams of many different sizes and durations dot the managerial and professional landscape of almost every purposive organization. Managerial and professional work does not lend itself to the formation of single teams whose members work only on those teams for extended periods. Instead, one person is likely to serve simultaneously on a variety of different teams that form, reform, and disappear as frequently as sand dunes on a windy beach. In one financial institution, for example, almost all professionals served on multiple teams, some of whose expected life spans extended indefinitely into the future, some of which were created to accomplish a particular task and would disband when that task was finished, and some of which were created on the fly to solve an unanticipated but

pressing problem immediately. That is as it must be when teamwork is performed in uncertain or fast-changing environments. But even temporary teams whose members serve simultaneously on others are not immune from the requirement that they be thoughtfully composed and appropriately bounded. Sand-dune teams, as much or more than permanent teams, need the advantages of having clear and moderately permeable boundaries.

Managers and professionals sometimes talk and act as if they do not require the level of clarity about team boundaries that they would insist on for rank-and-file teams in their own organizations. "We're experienced at this kind of thing; it's what we do every day," one senior manager told me, "so we don't need to have all the i's dotted and t's crossed." It may be true (although I doubt it) that people who do managerial and professional work are better than others at handling ambiguity and uncertainty. But the underbounded character of many such groups makes it harder than it need be for them to demonstrate how good they really are at teamwork.

Ironically, many senior teams could learn a great deal about establishing and maintaining clear team boundaries from teams that carry out front-line work. Recall, for example, the teams of flight attendants at the domestic airline described in chapter 1. The four-person teams formed during training had clear boundaries and stable membership. When passenger loads were light, however, not all four members were required for in-flight work. On those days, one member could stay at company offices to perform other work while his or her teammates handled the team's in-flight duties. The team's boundaries remained clear despite the fact that members were flexibly deployed as circumstances changed.

A production team in a semiconductor manufacturing plant, by contrast, had to contend with constant changes in membership that were entirely out of the team's control. The number of people needed for the work was a direct function of the number of orders the company received for the team's product, and that varied widely (some would have said wildly) from month to month. When the backlog of orders got large, management would hire additional staff; when production capacity was greater than needed, less senior team members would be laid off or furloughed. Life on the team was a yo-yo of expansion and contraction separated by occasional periods of relative stability of membership.

An opportunity to establish clearer team boundaries developed when the production manager decided to give teams in his area increased

responsibility for managing their own work. One of the team's new responsibilities was the management of its boundaries. The manager determined the number and mix of members who were needed to accomplish the team's work when production requirements were at an absolute minimum, and he designated those individuals as "core" team members. Core members were given employment security in that they would not be laid off unless the entire production area was shut down. In exchange, they were required to manage the expansion and contraction of their team as business needs changed.

After some initial resistance (adding and subtracting staff is emotionally demanding and always had been a management responsibility), members stepped up to the challenge and handled it well. The team gradually developed its own pool of people who would be invited to augment the team when production demands increased. Core team members invested time and energy to develop the skills of these individuals so they would be able to join in the work quickly and smoothly. And, over time, the temporary staff came to identify with the team and develop a measure of loyalty to it.

There was nothing this organization could have done to eliminate the fluctuations in demand that frequently required that the team expand or shrink in size. But by giving the team control of its own membership, the production manager clarified and affirmed the boundaries of the team and thereby reduced significantly its vulnerability to the dysfunctions of underboundedness.

Making music is about as different from making memory chips as it is possible to get. Yet the Orpheus Chamber Orchestra, a self-governing orchestra based in New York City, had much the same problem as the semiconductor production team and solved it in much the same way. The orchestra consists of twenty-six members. Few pieces in the chamber orchestra repertoire call for exactly that number of instrumentalists, however. Pieces that require fewer players pose no problem since the orchestra can use the same strategy as the flight attendants and let members who are not needed do other things. But pieces for larger ensembles require that the orchestra use additional players, who are recruited from the ranks of New York-area freelance musicians. Many of these additional players also serve as substitutes when an Orpheus member will be unavailable for a performance or tour, and over time members have come to know the best of them quite well. Although the substitutes have neither

any guarantee of future employment nor the right to participate in orchestra decision making, a number of them have become de facto associate members of the ensemble. Moreover, on those rare occasions when there is an opening in the core orchestra, the substitutes invariably are among the prime candidates to fill it. Like the semiconductor team, the orchestra maintains a clear membership boundary, but does so in a way that provides both the flexibility that the work requires and a pool of reserve members who are well known by the core team, familiar with the team and its work, and usually ready to pitch in when needed.

The three teams just described—the flight attendants, the semiconductor production teams, and the chamber orchestra—all performed front-line work in their organizations. Although the nature of that work could hardly have been more different, each of these teams faced and overcame a significant challenge in establishing and maintaining clear team boundaries. Many management and professional teams, by contrast, have similar challenges but are less inclined to step up to them. Those who lead or serve on high-level teams in organizations would be well advised to notice and learn from those front-line teams that manage to achieve, simultaneously, both clear team boundaries and flexibility in obtaining and deploying human resources as task demands change.

Delimited Authority

Once it is clear what a team's work is and who is on the team, the two features just discussed, the next item of business is to determine the extent of the team's authority. It almost always is a good idea to do this explicitly. Otherwise, team members will do it implicitly as they proceed with their work—and, in the process, run a significant risk of either excessive timidity in making decisions or overstepping the actual bounds of their authority. The only way to avoid such miscalibrations is for managers to specify when a team is formed just how much authority the team initially will have and to make sure that members understand clearly what decisions are and are not theirs to make. There are many aspects of teamwork for which ambiguity is a good thing (such as leaving room for members to bring their own meaning to the team's purposes, as we will see in chapter 3), but the extent of a team's authority is not one of those aspects.

Both leaders and team members typically experience a measure of

anxiety, angst, and ambivalence when authority is transferred from management to a work team. To say "This is now *your* team's responsibility, and you are fully accountable for how you handle it" is virtually to guarantee that interesting social and emotional juices will begin to flow, both between the team and its manager and among team members. For this reason, there is no guarantee that the team will exercise no less, but also no more, than the full amount of authority it has been assigned. I have observed numerous instances of both underuse of the authority a team actually has and inappropriate exercise of authority the team does *not* have. The former is more common early in the life of a team, when managers may be overoptimistic about members' readiness to decide about matters that previously were someone else's responsibility and therefore give the team too much authority too soon. The latter is more common after a team has gained some experience and logged a series of successes, at which point members may decide (especially if they are located far from headquarters) that they can just go ahead and run the whole show. For all of these reasons, the process of helping a team come to terms with both the extent and limits of its authority requires leaders to exercise significant leadership skill and to exhibit no small measure of emotional maturity (see chapter 7).¹⁵

In deciding the extent of a team's authority, it is helpful to consider deliberately and thoughtfully who is in the best position to handle each of four functions that must be fulfilled by any organizational unit that has work to accomplish. The first and most basic function, of course, is to *execute* the work—applying physical or mental energy to accomplish tasks. The second function is to *monitor and manage* the work process—collecting and interpreting data about how the work is proceeding and initiating corrective action as needed. The third is to *design* the performing unit and arrange for needed organizational supports for the work—structuring tasks, deciding who will be involved in performing them, establishing norms of conduct for work behavior, and making sure team members have the resources and assistance they need to carry out their work. The fourth is to *set direction* for the team—specifying the collective objectives and aspirations that spawn the myriad of smaller tasks that command the attention and energy of members of any purposive group or organization.

There is great variation across teams and organizations in how many of these four functions are assigned to teams themselves and how many

are handled by managers. (In this book, *manager* refers to any individual whose main responsibilities have to do with directing and structuring the work of others.) As is shown in figure 2-1, four levels of increasing team self-management can be identified, depending on who has the authority for each of the four functions.¹⁶

Members of a *manager-led* team have authority only for actually executing the task. As seen in the left-hand column of the figure, it is managers who monitor and manage team performance processes, structure the unit and its context, and set overall direction. In such teams, managers manage, workers work, and the two functions are kept as separate as possible.

Next come *self-managing* teams, whose members have responsibility not only for executing the task but also for monitoring and managing their own performance. This type of unit is often seen in new, high-commitment plants and is commonplace in professional work—for example, a team of service providers who share responsibility for promptly and efficiently

FIGURE 2 - 1

The Authority Matrix

FOUR LEVELS OF TEAM SELF-MANAGEMENT

Setting overall direction				
Designing the team and its organizational context	Management Responsibility			
Monitoring and managing work process and progress		Team's Own Responsibility		
Executing the team task				
	Manager-led teams	Self-managing teams	Self-designing teams	Self-governing teams

Source: Copyright (c) 1986 by the American Psychological Association. Adapted with permission.

handling requests from clients, or a construction team whose members track their progress and make any adjustments in their work strategy that are needed to get the job completed on time and in accord with the project plans.¹⁷

Members of *self-designing* teams have the authority to modify the design of their team or aspects of the organizational context in which they operate, or both. Managers set the direction for such teams but give members full authority for all other aspects of the work. Management task forces often are self-designing units, such as the four-person team of senior bank managers studied by Connie Gersick who were charged with creating a new kind of money market account. Individual members of this team began talking with one another informally as soon as legislation was passed that allowed banks to offer this type of account. After consulting with the bank CEO and managers of those departments whose work would be affected by the new product, the team settled on its own composition and defined its own task—which was to do whatever was needed to get the product ready for market by the first day the bank would be permitted to offer it. The team's overall direction was prespecified by the terms of the legislation authorizing the product, but members had control of virtually everything else. Moreover, Gersick reports that the CEO personally affirmed the task force's self-designing character by informing the team that "It [the product] had better be good!" but refraining from issuing any instructions whatever about how the team should organize itself or carry out the work.¹⁸

Finally, *self-governing* teams have responsibility for all four of the major functions listed in figure 2-1. Team members decide what is to be done, structure the unit and its context, manage their own performance, and actually carry out the work. Examples of self-governing units include some legislative bodies, corporate boards of directors, advisory councils of community service agencies, worker cooperatives, and professional partnerships.

The flight attendants at the international airline described in chapter 1 are located near the lower boundary of the self-managing category. Although they did have responsibility for monitoring and managing many parts of their work, that was only because the work was performed out of management's sight. The management strategy of their company, clearly, was for cabin service teams simply to execute work that headquarters managers had conceived, designed, and engineered in detail. Flight attendants at the

domestic airline, by contrast, were self-designing. These teams not only executed and managed their customer service work but also had the authority to make decisions about their own structure. Members selected their own teammates when groups were formed, they decided whether the day's flying required the presence of all four members, and they chose which team member would serve as team leader each day. So long as their decisions were consistent with the company's overall direction (and stayed within the bounds of the handful of "must do" and "must never do" constraints that management specified), teams were free to structure their work as they saw fit.

The levels of authority held by flight attendant teams at these two airlines roughly bound the scope of this book. We do not spend much time on manager-led teams, the leftmost column in figure 2-1, because such teams invariably waste human resources. For managers to specify and control all aspects of a team's work in real time is but a group-level version of scientific management—the dysfunctions of which, for both people and organizations, have been amply demonstrated over the last several decades.¹⁹ Nor do we analyze in depth the special decision-making dynamics that occur in self-governing groups, the rightmost column in the figure.²⁰

Self-managing and self-designing teams, the middle two columns of the matrix in figure 2-1, are found in great numbers throughout all sectors of society—in businesses, in public agencies, and in nonprofit organizations. When well designed and well led, they can be a self-correcting and renewable collective resource, an organizational feature that builds rather than expends human capital. But when they go sour, they can do real damage both to organizations and to people. Managers who clearly specify a team's authority, and who make certain that team members understand both the extent and the limits of their responsibilities, go a long way toward making the former outcome more likely than the latter.

Stability over Time

We already have seen that teams whose boundaries are so unclear that it is impossible to know for sure who is actually in the group are almost certain to encounter difficulties in carrying out their work. But how about groups whose membership is reasonably clear at any given moment but likely to change at the next moment? The answer to this question is clear:

Teams with stable membership perform better than those that constantly have to deal with the arrival of new members and the departure of old ones.

The evidence documenting the validity of that assertion is, as I read it, incontrovertible. That is why, as I board a commercial airline flight, I occasionally have the impulse to stick my head into the cockpit and casually ask, "First trip flying together?" The impulse is as silly as acting on it would be embarrassing, and I never actually would do such a thing. But the reason the thought comes to mind derives from some analyses performed by the National Transportation Safety Board (NTSB) that have profound implications for both airlines and their passengers. NTSB staff combed the agency's database to identify the circumstances under which aircraft cockpit crews are most likely to get into trouble. They found that 73 percent of the incidents in the NTSB database occurred on a crew's first day of flying together, and 44 percent of those took place on a crew's very first flight.²¹ Crews are most vulnerable to mishaps when they are just starting out, before they have had the chance to learn through experience how they can best operate as a team.

This NTSB analysis reinforced earlier findings obtained at NASA's Ames Research Center in a project intended to assess the effects of fatigue on pilot behavior but that unexpectedly wound up showing the benefits of keeping members of task-performing teams together. The NASA researchers recruited a number of crews as they returned from several days on the line, and compared their performance on a moderately demanding scenario in an aircraft simulator to that of crews whose members had just completed several days off duty. As one would expect, the pilots returning from multiple-day trips were indeed tired—and individual members of those crews did make more mistakes than did the rested pilots. The surprising finding was that the fatigued crews, *as crews*, made significantly fewer errors than did crews composed of rested pilots who had not yet flown together. Having experience working together as a team more than overcame the debilitating effects of individual fatigue.²²

There are many reasons why reasonably stable teams perform better. Members develop familiarity with one another, their collective work, and the work setting, so they are able to settle in and focus on working together rather than waste time and energy getting oriented to new coworkers or circumstances. They develop a shared mental model of the performance situation, one that, with time and experience, is more integrative than the

individual models with which they began. They develop a shared pool of knowledge, accessible to all, and build what social psychologist Dan Wegner calls "transactive memory" (that is, members themselves serve as memory aids to one another, providing the possibility of collective recollection that exceeds the capacity of any single individual). They learn who is especially knowledgeable or skilled about which aspects of the work and thereby build the team's capability to actually use what members know and know how to do. They learn how to deal appropriately with those of their members who may be less skilled, either in teamwork or on the task itself, without excessively disrupting the team's progress toward achieving its objectives. And, gradually, members are quite likely to build shared commitment to the team and a measure of caring for one another.²³

These are all positive features of well-functioning teams, but none of them by itself is the main or only mechanism responsible for the superior performance of stable teams. Instead, the features come in a package, with each process reinforcing the constructive impact of the others. There is no specific action a leader can take to ensure that these processes will develop; like all other human and social systems, work teams develop in their own, idiosyncratic ways. Research does suggest, however, that training team members *together*, rather than as separate individuals, can jump-start the development of performance-enhancing team processes.²⁴ A team whose members learn together how to work together, and who then stay together to further build their collective competencies, is almost certain to develop into a more effective performing unit than otherwise would be the case.

The research findings summarized earlier have clear implications for management policy and practice. In commercial aviation, for example, crews should be trained together and then remain intact for a considerable period of time, giving members the opportunity to develop themselves into the best-performing unit that they are able to become. Moreover, on any given trip they should fly the same aircraft and work with the same cabin crew. And the leader of the crew, the captain, should conduct a team-oriented briefing before each trip to reduce as much as possible the crew's exposure to the liabilities of newness.

Yet in most airlines, crew members are trained as individuals and crew composition is constantly changing because of the long-standing practice, enforced by labor contracts, of assigning pilots to trips, positions, and aircraft

as individuals—usually on the basis of a seniority bidding system. In one airline my colleagues and I studied, for example, a normal day's flying could involve two or even three changes of aircraft and as many different cabin crews. In another carrier, it was not uncommon for cockpit crews themselves to have one or two changes in composition during their one- or two-day life spans. Perhaps the most vivid example of crew instability that I have observed occurred when a crew lost its first officer for personal reasons while members were setting up the cockpit prior to departure. During the wait for a reserve pilot to appear, I asked the captain if he was concerned about the last-minute change in crew composition. "No problem," he responded. "Every pilot in this company knows his job, and the new first officer will pick up right where Bob left off." Because we were already late departing, the captain called for pushback and the prestart checklist immediately after the reserve pilot arrived. Engine start, taxi, and takeoff proceeded normally, and only when we were well into the climb were introductions made all around.

Why have airline managements, pilot unions, and federal regulators, all of whom are deeply committed to improving the safety of flight, not jumped to implement policies and practices based on the research findings just summarized? My conversations with representatives of these groups suggest two reasons. First, they do not want to believe the findings. Everyone knows that if a team stays together too long members will become too comfortable with one another, lax in enforcing standard procedures such as checklists, and too ready to forgive teammates' mistakes and lapses. Yes, teams may become better at working together as they move through the early phases of their lives. But that learning happens quickly, then plateaus, and then, at some point, overfamiliarity sets in and dominates members' subsequent interaction. It is better, therefore, to have a constant flow-through of new members to keep teams on their collective toes.

Everyone knows such things—but they are not true. Members of competently designed teams do learn fairly rapidly how to work together, as claimed. But, except for one special type of team, I have not been able to find a shred of evidence to support the view that there comes a point at which the learning stops and the positive trend reverses, when compositionally stable teams function decreasingly well the longer members stay together. (The exception is research and development teams. Organizational researcher Ralph Katz found that the productivity of such teams peaked

when members had worked together for about three years, and then began to decline. It appears that research is a type of teamwork for which a moderate flow-through of new members really does help, probably because the new arrivals bring with them fresh ideas and perspectives to which the team might not otherwise be exposed.²⁵) The very best teams get better and better indefinitely, like a great marriage that is stronger on the couple's fiftieth anniversary than it was on their first, or like the Guarneri String Quartet, whose members have continuously improved their music making over more than three decades of playing together.²⁶

A second reason managers and policy makers do not take the benefits of team stability as seriously as they might is that doing so often would disrupt traditional organizational practices, constrain members' choices about their schedules and assignments, and cost a great deal of money. The bidding systems that are in place at most U.S. carriers, for example, provide individual pilots, especially senior pilots, with far greater latitude in choosing their work schedules than would be possible under any system that promoted crew integrity and stability. Moreover, it is much more efficient in making work assignments to treat each crew member as an individual than it would be to deploy crews as intact units—especially if one adds the constraint that each crew must remain with a particular aircraft for some period of time. The financial benefits of this scheduling strategy can be computed, and they are substantial. (An analyst at one airline did a back-of-the-envelope calculation for me and reported that keeping crews together would cost his company several million dollars a year; that, he said, made the idea a nonstarter in the industry's highly competitive environment.)

Keeping teams together does not inevitably increase costs, however. For digging tunnels, for example, the use of intact teams has been shown to generate significant financial benefits. Journalist Fred Hapgood interviewed Howard Handewith, an analyst for the American Underground-Space Association, about the matter. According to Handewith, intact teams were able to construct a mile of superconducting supercollider tunnel for \$5 million, about a third of the typical cost of big-tunnel construction. Hapgood goes on to report that "the tunneling teams arrived fresh from working on the Chicago water system and the Dallas subway. They came already tuned to one another and to their tools. Handewith believes that

these results suggest the kind of returns that can be expected as the opportunity to work steadily, as part of a team, becomes more common."²⁷

Even when keeping work teams together would be costly, those costs must be weighed against the also real costs of substandard team performance that are incurred when scheduling and rostering practices make it impossible for members to become established as a stable work team. Our research suggests that this also is a substantial cost—albeit one that is harder to compute in dollar terms. Indeed, when the team's work is important enough and when money is no object, intact teams *are* used in aviation operations. The Strategic Air Command (SAC), whose aircraft would have delivered nuclear bombs had that become necessary during the Cold War years, did train and schedule crews as intact units. Aircraft maintenance crews also were intact teams, and both the flight crews and the maintenance crews were assigned to particular aircraft. Fortunately, SAC crews were never called on to perform their ultimate mission. But they did fly frequent training missions that were scored on objective criteria such as bombs on target. A colleague who helped train those crews reports that their performance was, with rare exceptions, superb. Inertial forces are powerful, however, and traditional practices persist in both civilian and military aviation despite the accumulated evidence that they almost certainly underutilize the enormous pool of talent and motivation present in the pilot workforce.²⁸

Is all of the above relevant only to teams that fly airplanes? No; unclear team boundaries and instability of membership is a pervasive and pernicious problem for many different types of organizational work teams: project task forces, senior management teams, sales teams, and, increasingly, virtual teams whose membership may be as fluid and uncertain as it is dispersed. As we will see throughout this book, the most important contribution that leaders can make to the teams they create is to establish and maintain the handful of organizational conditions that foster and support competent teamwork. Doing that, however, requires that there be a team in place to create those conditions *for*. A reasonably stable team gives members the opportunity to accumulate shared experiences and to learn from them, to be sure. That in itself is sufficient reason for keeping teams together. But even more important, it is team stability that makes excellent team leadership possible.

LAYING THE FOUNDATION FOR EFFECTIVENESS

Creating real work teams is akin to laying a solid foundation for a building. If the foundation is well conceived and solid, the builder can proceed to erect the rest of the structure with confidence. If it is not, the building will never be as sturdy as it could have been. Neither an eye-catching above-ground design nor elegant interior finishings can ever compensate for a flawed foundation.

The same is true for work teams in organizations. Real teams can be small or large, can have wide-ranging or restricted authority, can be temporary or long-lived, can have members who are geographically co-located or dispersed, and can perform many different kinds of work. But if a group is so large, or its life is so short, or its members so dispersed and out of touch with one another that it cannot operate as an intact social system in carrying out its work, then prospects for effectiveness are dim.

Wise leaders therefore lay the foundation for team effectiveness by making sure they have created a real work team. As we have seen in this chapter, this means making sure that the task actually is appropriate for teamwork and that it requires members to work together interdependently. It means establishing clear but moderately permeable membership boundaries. It means providing the team with substantial but clearly delimited authority for managing its work. And, finally, it means ensuring that the team will be reasonably stable over time as members carry out that work.

Creating a real work team establishes the first of the five performance-enhancing conditions discussed in this book. Moreover, a real team is the prerequisite for the other conditions. Managers who establish real teams have more choices, and better choices, about setting a good direction for the team (chapter 3), creating an enabling team structure (chapter 4), arranging organizational supports for the team's work (chapter 5), and providing hands-on coaching that helps members work together productively (chapter 6). To give insufficient attention to getting the essential features of real work teams in place not only puts the team and its clients at some peril but also makes the team leader's own job harder than it need be.