

## Chapter 4

## Evidence-based Management for Entrepreneurial Environments: Faster and Better Decisions with Less Risk

JEFFREY PFEFFER

**E**NTREPRENEURSHIP IS RISKY. Most new technologies and new businesses fail. Shane (2008) reported that 25 percent of new businesses failed in the first year and that by the fifth year, fewer than half had survived. In the United Kingdom, Stark (2001) presented data showing a 75 percent failure rate for small and medium-sized enterprises in the first three years. The risk and high failure rate is because most new ideas and technologies are not good and are, therefore, rejected by the marketplace.

High failure rates have become accepted as an inevitable cost of entrepreneurial activity, offset by the jobs, wealth, and ideas created by those new ventures that are successful. So the venture capital industry's business model is premised on getting a few exceptional returns ("home runs") among the multitude of failures in each portfolio. For instance, a German venture capital fund begun in the late 1990s showed a cumulative internal rate of return of negative 3.8 percent as of 2009, almost break-even over the period. But of the 28 investments the fund had made, 11 had no value at all and four were worth less than 15 percent of the value of the initial investment. The almost break-even return was the result of one investment worth four times and another six times the amount invested as well as some smaller positive returns. A study of 128 exited investments in the United Kingdom also reported a highly skewed distribution of returns, with 34 percent being a total loss, 13 percent of the exits at break-even or a partial loss, and 23 percent of the investments having an internal rate of return of above 50 percent (Mason and Harrison 2002).

As a consequence of this high rate of failure for new ventures, both human and financial resources go to waste. Many talented people, including engineers, scientists,

## 58 JEFFREY PFEFFER

and others with advanced degrees spend enormous time and energy on entrepreneurial activities with little to show for it other than what they learned from the experience. The wasted effort derives in part from the fact that it is often difficult to know when a new venture is beyond hope or when the investment of a little more time and money can make it successful. There are numerous examples, the Apple Newton being just one, of a product idea that failed because it was too early for the market, where subsequent variations of the same basic idea turn out to be huge commercial successes. Consequently, the temptation to persist is strong. Such persistence reflects the psychology of escalating commitment (e.g. Staw 1976), which argues that people do not want to admit they have made a mistake with the negative implications for their self-concept and therefore become psychologically identified with their decisions. This persistence also reflects the uncertainty of not knowing when a small incremental investment will actually make the earlier efforts pay off (Heath 1995). And there is a natural tendency to not quit and consequently risk having others capitalize on the unrealized potential of one's efforts. If it were possible to more quickly and accurately forecast the likelihood of success and make decisions that would increase success rates, at least some of that human capital would not go to waste.

A similar waste of resources characterizes the financial capital that is plowed into entrepreneurial ventures. The evidence shows that many investors do not earn returns commensurate with the risks they take. Kaplan and Schoar examined returns to private equity—venture capital and leveraged buyout funds—over the period 1980–2001. They found that the median internal rate of return for VC funds was 11 percent and that the median venture capital fund's performance was only about two-thirds that of the public market equivalent, measured as the return to the Standard & Poor's 500 (Kaplan and Schoar 2005). Cochrane (2005), looking at individual transactions rather than funds, a methodology that admittedly leaves out management and performance fees accruing to the general partner, concluded that VC returns were similar, in their means, standard deviations, and volatility, to the returns shown by smaller NASDAQ-traded stocks. However, lacking a public market, the venture capital investments were inherently riskier and less liquid.

Industry-wide estimates of financial returns to entrepreneurial investments are highly skewed by a few prominent, early and successful entrants to the venture capital and for that matter the hedge fund industry that have earned exceptional returns. Kaplan and Schoar (2005) noted the substantial heterogeneity in returns to venture capital funds. Indeed, one manager of a fund of funds investing in other venture capital partnerships commented that more than 100 percent of the industry's returns were earned by the top 20 (in terms of performance) firms, which means that there are literally hundreds of venture capital firms that have returned nothing to their investors.

These poor results from much, although obviously not all, entrepreneurial activity occur in spite of the hard work and diligence of many talented individuals. Moreover, little seems to have changed over time, indicating that there has been little learning or improvement in decision-making quality. These facts suggest that there may be potential to improve the decision-making process associated with developing and building new enterprises.

## WHY HIGH FAILURE RATES FOR ENTREPRENEURSHIP PERSIST

There are several causes for the persistently high failure rates for new businesses. One problem is that it has become conventional wisdom, accepted by all the parties ranging from entrepreneurs to those who provide them financing, that a high rate of failure is an inevitable consequence of doing new things, inventing new technologies, and opening up new markets—activities which are inherently risky and uncertain because they involve doing things that have not been successfully done before. Because this conventional wisdom suggests that a high failure rate is inevitable, there is often little effort expended trying to improve decision-making in new venture activity.

A tremendous amount of the culture of high technology entrepreneurship is carried in and influenced by the venture capital community. Many of these firms do what they do without much introspection or reflection, partly as a result of the egos and self-confidence of the VC partners. One of the more consistent findings in social psychology is the so-called “above-average effect,” in which more than half of most people believe they are above average on virtually all positive qualities, even including height and income (see, for instance, Kruger 1999; Chambers and Windschitl 2004). People who have survived and prospered in the venture industry have obviously done well, and those VCs who don’t do well generally don’t last. Therefore, it is axiomatic that most fund managers believe they are much above average in their abilities and in their decision-making. Consequently, many believe they don’t need to learn much or have much to learn. This attitude exists even though VC success may be as much a function of the particular firm where one works, one’s timing in both entering the industry and when investments were made, and random good luck as a consequence of any particular individual skill. There is much research that suggests that when good performance outcomes occur, positive qualities get attributed to the people, groups, or companies that enjoy those good outcomes (e.g. Staw 1975; Rosenzweig 2007). This association of positive attributes with good performance occurs whether or not such attributes were causally related to the good results or even whether or not the high-performing entities actually possess the positive qualities. This means that high-performing VCs will be perceived as having individual skill as a consequence of their performance, whether or not such skill actually exists.

Precedent and the way things are done in the entrepreneurial financing industry have been substitutes for thinking for quite a while (Pfeffer and Sutton 1980). Yet another issue that constrains improvement in decision-making is that pressures to do what others in the industry are doing, because of the assumption that the crowd is invariably wise, are strong. Entrepreneurs, too, mostly have strong egos, which is what is required to take on something new where the risks of failure are high. But this overconfidence among entrepreneurs and those that back them makes it difficult for people involved in creating new businesses to question things and to learn from setbacks and other experience.

Moreover, most venture capitalists and entrepreneurs believe that outstanding individual people make the difference, leading them to focus on finding and recruiting

## 60 JEFFREY PFEFFER

stars and to eschew much attention to process, including decision-making processes. In addition, most investors engage in a set of ritualized due diligence practices with little effort to close the loop and learn from the results of their decisions—after-action or after-event reviews are reasonably rare. With little effort expended to improve entrepreneurial decision quality, not surprisingly, decision quality doesn't improve. Therefore, failure rates don't change. This persistence of failure rates over time seemingly reinforces the validity of the conventional wisdom that high failure is an inevitable consequence of entrepreneurial activity. And the cycle continues.

Yet another possible reason for such small changes in rates of failure is that few of the participants in entrepreneurial activity suffer significant consequences from unsuccessful decisions, and therefore many players have less incentive than one might expect to improve their decision-making. As has been documented, much of the return to the principals or general partners in both hedge funds and venture capital funds come from the guaranteed annual percentage they earn, typically 2 percent of the amount of the fund's principal (Mackintosh 2009). Entrepreneurs often, although not always, are working with other people's money, so their financial downside, except in terms of the opportunity costs of their time, are also limited. And, because failure is most often seen as an unavoidable risk of being an entrepreneur, there are few if any career risks for starting something that doesn't work out. Many entrepreneurs go on to work at least temporarily in VC firms and few have much difficulty finding subsequent jobs or, for that matter, investment capital. John Lilly, for instance, currently the CEO of the Internet browser company Mozilla, was first the CEO of Reactivity, a company that was ultimately unsuccessful.

I am unconvinced that high rates of failure are inevitable and that improvement in decision-making is impossible. Consequently, in this chapter, I outline a case for applying evidence-based management to entrepreneurial activity. After first defining the elements of an evidence-based management approach, I consider a few commonly-voiced but largely inaccurate objections to its use and then provide some examples of how evidence-based decision-making has been and could be used to improve the quality of entrepreneurial decision-making. My argument is premised on the idea that people can improve the quality of their decision-making in all environments, and that an evidence-based approach is one reasonable way to accomplish this.

## THE FOUNDATIONS OF EVIDENCE-BASED MANAGEMENT

Evidence-based management (hereafter EBM) is modeled on the evidence-based movement in medical practice. Although evidence-based decision-making in medicine is growing in its acceptance, in medicine as well as in other contexts such as criminology and education, evidence-based practice has historically faced resistance to its implementation. That resistance continues to the present (Domurad 2005) and makes implementing evidence-based management more challenging than it should be.

EBM seeks to apply the best currently available data and theory to managerial decision-making (Pfeffer and Sutton 2006). The underlying assumption is that although it is the case that at any given point in time information is incomplete and

## EVIDENCE-BASED MANAGEMENT 61

over time the evidence on what to do and how to do it changes as new data come in, in general, decision quality will be higher if people make fact-based decisions. Moreover, it is incumbent on both individual organizations and larger communities of practice to systematically gather and learn from actual experience so that, over time, decision quality progressively improves. EBM emphasizes gathering and paying attention to the data, understanding the best current theory about the subject of a particular decision, and continually updating both theory and evidence as new information becomes available. Although such an approach seems logical and, indeed, almost like common sense, it actually requires a different mindset than is common in most organizational management.

As we know, decisions are not always based on data and theory (Pfeffer and Sutton 2006). Instead, organizations frequently rely on casual benchmarking—following what others are doing regardless of whether the experience of others in possibly quite different circumstances is relevant to their own case. Leaders also make decisions based on their own experience, even though such experience is often unreliable as a guide for subsequent action for several reasons. Experience is a problematic guide to action because there is a tendency to see what we expect to see, the basis of all magic acts, which means learning from experience is difficult and requires effort. Few organizations outside of the military, with its after-event or after-action reviews and medicine, with its mortality and morbidity conferences, engage in the systematic, structured reflection that would be required to learn from experience. Experience is also inherently idiosyncratic, reflecting a particular case and set of circumstances, and therefore suffers as a guide to action from the problem of trying to derive general principles from very small samples. Finally, experience at its best is a guide for decision-making in situations that mirror the past from which the experience comes, but past experience may be unreliable in providing guidance in very different or novel contexts.

In addition to experience, decisions often reflect what leaders believe to be true—their ideology (Tetlock 2000)—and what they have done in the past and seems to have worked. Ideology colors what people see and how they apprehend the world around them, as well as how they incorporate their observations into decisions. As such, ideology, and by this mean I mean political ideology, colors what people do even in business decision contexts (Tetlock 2000). In addition, people naturally tend to advocate doing things that favor their own competencies and interests and that are consistent with enhancing their self-image. None of these bases for making decisions leads to particularly sound, fact-based choices.

With its emphasis on taking action on the basis of the best knowledge available at the moment while recognizing that all knowledge is imperfect and therefore we need to learn from experience, evidence-based management is consistent in its underlying philosophy with an attitude of wisdom. As psychologists John Meacham and Robert Sternberg have argued, wisdom means knowing what you know and what you don't know, and acting on the basis of what is known at the moment while being open to changing your mind (e.g. Meacham 1983; Sternberg 1985).

## 62 JEFFREY PFEFFER

With the emphasis on data and feedback processes, EBM is also consistent with the principles of design thinking as practiced in places like product design company IDEO and other firms such as Procter & Gamble (e.g. Brown 2008; Kelley 2001; Martin 2009). A design-oriented approach emphasizes prototyping and systematically learning from experience and also getting into the field to see how real people actually use products and services so that new versions can be based on the issues people face as they interact with the company's products. EBM is quite consistent with this notion of running experiments—building prototypes and seeing how people react—and also with embedding design in learning from real situations.

Finally, the evidence-based management idea is also consistent with many of the ideas of the total quality movement. Just as in the case of quality efforts, EBM stresses diagnosing the root causes of problems and addressing those fundamental sources rather than just treating symptoms or acting without doing any diagnosis at all. Also, like the quality movement, EMB emphasizes the gathering of systematic data to the extent possible so that actions can be formulated using the best information available.

The quality movement and its approach has fallen into some disuse—even Toyota has recently experienced substantial product problems—mostly because an emphasis on quality requires systematic, persistent discipline that is difficult to maintain when confronted by the temptation to try new ideas and the boredom and fatigue that results from close attention to detail and process. And although design thinking has been featured in numerous books and articles, it, too, is not as widely practiced as its apparent publicity success would suggest.

### OBJECTIONS TO AN EVIDENCE-BASED MANAGEMENT APPROACH

If evidence-based management and other, complementary approaches are not widely used, it is important to understand why. We frequently hear the same issues raised as objections to using evidence-based management, and some of these concerns would seem particularly relevant for entrepreneurial decision-making. One concern is that the current business environment changes more rapidly than in the past, and the high-velocity of competitive dynamics make any process that takes a long time virtually irrelevant. Because it relies on facts, theory, and analysis, evidence-based decision-making takes more time than just acting on gut instinct or recalled experience. A second, related issue is that there are many decision circumstances for which good evidence and theory simply do not exist, rendering EBM largely moot. This would seem to be particularly the case for entrepreneurial decisions. What data or evidence can possibly be brought to bear on decisions about launching new technologies and new products into a competitive environment fraught with uncertainty? If entrepreneurs have succeeded in the past against all odds on the basis of their persistence, sometimes in the face of evidence that would argue against what they had done, this success only convinces them to ignore data, particularly data contrary to their intuitions, in the future. This often-misguided reliance in their own intuition occurs

## EVIDENCE-BASED MANAGEMENT 63

because history is often ambiguous and organizational learning is a process fraught with difficulties (e.g. Levitt and March 1988; Levinthal and March 1993).

Third, very much as the case with evidence-based medicine, executives are extremely reluctant to substitute theory or data for their own personal clinical experience and judgment. This latter point helps explain why there is so little transfer of knowledge between the research and practitioner domains of management.

None of these issues seems particularly compelling and there are things to be done in any event to mitigate their relevance. Even though many leaders complain about the length of time ostensibly required to implement evidence-based decision-making, many of these same leaders, even those in relatively small, high-technology enterprises, seem perfectly content to hire management consulting firms to provide advice and executive search firms to go outside to find additional talent. Such engagements typically not only cost a great deal of money, they often take months to complete and in the case of executive search, often result in either no hire or a poor one.

Practicing evidence-based management need not consume a lot of time in any event. With online databases and libraries that cover virtually every conceivable question, searching for the best theory and data takes comparatively little effort. Google Scholar is one such site that brings relevant research from peer-reviewed academic journals to a person's fingertips. Although not all of the content found on that site is free, the cost for accessing most articles is modest and pales in comparison to the consequences of the decisions companies make. And there is much information about products, services, company financial results, and markets available for a modest amount of effort or cost.

Often applying evidence-based management thinking is simply a matter of uncovering the assumptions that underlie some potential choice and then accessing the collective wisdom of one's colleagues to see whether or not those assumptions seem sensible. If the assumptions underlying a particular intervention don't seem plausible, then the odds of that intervention succeeding are remote.

As one example of this process in action, consider the decision to implement forced-curve performance ranking, something advocated by Jack Welch, the former CEO of General Electric, among others, and a practice that is widely implemented in companies of all sizes (e.g. Novations Group 2004). Although there is actually a great deal of research on the effects of forced-curve ranking systems under different business circumstances such as the degree of interdependence among tasks, the frequency of feedback, and what happens to low and high performers (e.g. Blume et al. 2009), one doesn't actually need to even access that evidence to ascertain whether or not implementing a forced-curve ranking system will be helpful. That's because like all organizational interventions, this management practice has embedded within it a set of implicit assumptions about employees, managers, and organizational effectiveness. Some of those assumptions in the case of forced-curve ranking systems are: people can be objectively ranked against each other; people will respond positively with efforts to improve their performance when they know their position in the rankings; managers will provide objective and reasonably frequent feedback to their employees telling them where they stand; and the competitive dynamics that are an

**64 JEFFREY PFEFFER**

inherent property of systems that cause people to compete with each other will not adversely affect learning or organizational performance. Although such assumptions may be true in some cases, in general they probably do not hold—which is why there is little evidence that forced-curve ranking improves performance and some evidence that suggests that this management practice causes numerous problems (Novations Group 2004).

In addition to uncovering assumptions, companies can move to better capture and utilize the data they gather as a result of their ongoing activities. Even relatively small companies today gather lots of data as part of their operations, information ranging from sales to customer complaints and returns to product development time to data on turnover and employee recruitment. Information-gathering is more automated and the cost of computer memory has fallen to trivial levels. Although there is much data on operations and sales, often such data are used primarily for accounting purposes or in organizational subunits such as operations or human resources, but they are not brought together at the senior level to provide a foundation for comprehensive, data-based strategic action. This is not just a problem in small, new enterprises—David Larcker, an accounting professor at Stanford, has shared numerous examples of larger companies that do not understand the process by which they make money in that they don't know their most profitable customers or even products and often have inaccurate estimates of costs.

Finally, entrepreneurs and their funders would be well-served to recognize the inherent difficulties and biases in estimating their degree of expertise and in over-learning the apparent importance of persistence in the face of seemingly contradictory evidence about the prospects for success. Business success is inherently an uncertain process. The point of EBM, much like its counterparts in medicine and the policy sciences, is not to perfectly account for every single instance, but rather, by the systematic application of data and theory, improve the odds of making a better decision.

**HOW THINGS MIGHT BE DIFFERENT: EVIDENCE-BASED MANAGEMENT IN SMALL ENTREPRENEURIAL COMPANIES**

Businesses founded on or using the Internet automatically generate a great deal of data. These data have traditionally been used mostly for analyzing and designing marketing campaigns and, of course, for assessing the effectiveness of various advertising strategies. But it is possible to use such data to build truly evidence-based companies, and the model offered by some of these enterprises provides ideas that can be employed by any organization or start-up, not just those doing software development or focused on the Internet.

Because of the high failure rate and its associated waste of resources, recently, some venture capitalists—although not necessarily the largest or most well-known—and some software companies have begun to advocate a different way of doing business and managing. Sometimes called agile software development, the movement



began in the software space particularly for products oriented to the web. But the movement—which is what this set of ideas should be called since it has advocates and seeks to change how companies do their business—seems to be diffusing. Based on the “lean” principles of Toyota, the idea is to expend as few resources as possible while you learn what the customer wants from your product or service, doing rapid iterations of new releases—putting the rapid prototyping ideas from new product development into action. Because learning is an explicit goal, agile development and lean design necessitates gathering and analyzing information so that every new iteration can incorporate past experience as efficiently and effectively as possible.

One articulation of the idea of quickly learning from experience comes from the book and website, *Getting Real*. As the company behind the book and the website, 37 Signals, explains it:

- Getting Real is about skipping all the stuff that represents real (charts, graphs, boxes, arrows, schematics, wireframes, etc.) and actually building the real thing.
- Getting Real is less. Less mass, less software, less features, less paperwork, less of everything that’s not essential . . .
- Getting Real is staying small and being agile.
- Getting Real starts with the interface, the real screens that people are going to use. It begins with what the customer actually experiences and builds backwards from there. This lets you get the interface right before you get the software wrong.
- Getting Real is about iterations and lowering the cost of change. Getting Real is all about launching, tweaking, and constantly improving. (37 Signals 2010).

Traditional software or, for that matter, almost any traditional product development process proceeds following what is sometimes called a waterfall or cascade process. First, an engineer or marketing person comes up with a product idea or change to an existing product. The product is then designed, often by engineers, and specifications developed. A prototype is made, and if it is a physical product, a bill of materials gets created and a manufacturing process is designed. If it is a software product, code gets written. Then the product goes through quality assurance to ensure compliance with the original design specifications, after which it is released to the market. At that point, marketing and sales tries to promote and sell something that has already been created, often with little to no end-user input.

Note that this traditional product or service development process takes a long time and entails significant investment *before* the company receives *any* market feedback. The agile process aims to short circuit this delay and cut the amount of investment by getting customer input early in the process and engaging in rapid, low cost product iterations, in each instance gathering data and learning as much as possible from such data.

One example of a company that assiduously adheres to an evidence-based management, agile approach is Rypple. Founded by some former senior leaders from the workforce scheduling company, Workbrain, Rypple’s aim was to overcome many of the limitations of the traditional performance management process. Although

## 66 JEFFREY PFEFFER

people want to improve and desire feedback, the traditional, hierarchical appraisal process, often tied to monetary compensation, sometimes requiring ranking people against each other, is almost universally disliked. Employees do not like getting appraisals and few managers enjoy doing them. Rypple's goal was to design a software system that made asking for and providing feedback anonymous, quick, and easy.

The company and its culture are very data driven. Even as a very small start-up, they hired an MIT business school graduate whose mission was data and analytics. And every step of the way, Rypple emphasized gathering information, learning from it, and then improving the product as well as the distribution process.

Daniel Debow, co-CEO and co-founder, noted that instead of thinking in terms of a "product launch cycle," it was more useful to think of a "customer discovery cycle." Right from the beginning, Rypple's people had the objective of not proving themselves right, but instead, proving themselves wrong, and along the way, to have humility in their inability to accurately predict the future. In the beginning, they built just paper (PowerPoint) prototypes of the Rypple product. As Debow noted, "before we even hired a developer, we just built PowerPoint mock-ups and put them in front of people to see how they would react." Based on potential user interest and comments, the company then built a very bare prototype—no security, for instance, or log-in. They showed this prototype to some more people who wanted to use it, so Rypple put this bare-bones prototype into a couple of companies to see what happened. Based on that initial user feedback, the cycle of iteration and learning continued.

Rypple could and did continuously gather data on how many people were using the various prototypes, how many people were responding to requests for feedback, what ways of requesting feedback generated the most and best responses, what the pattern of usage was, and so forth. They tweaked phrasing, the user interface design, every aspect of the service, and carefully monitored variations in user response. Over time, other features such as security were added. But the cycle time for new iterations was almost weekly, and the cycle time for learning and incorporating that learning into new generations of the product were just about as fast.

One of the reasons Rypple could "afford" to have its users help develop the product without upsetting those people is because the initial version of the product, not for companies but for individuals, was free. If people haven't committed money to some licensed software, they are less irritated if the product isn't perfect. And because of the agile and lean software development process, Rypple didn't need to generate as much money because its expenditures on the typical product development process had been drastically reduced—not only in resources but in time.

Debow was quite clear as to why he had not been able to implement a similar process in Workbrain and why there was resistance to a data-driven development cycle in many traditional companies: control. As he told me:

Everybody in software had been brought up on the rational IBM-like products and these very engineering-oriented processes and documentation. That was the way people were going to control for bugs or other issues and get things

right. It was all about control. What is at the core of a lot of this agile and lean thinking is that actually you can't claim anything. And that kind of killed people.

The lean movement in software development requires, everyone agrees, more than just analytics expertise. It entails an enormous mindset shift, about how a company does product development, about the amount of control exercised at each stage of the process, about listening to customers, about the role and treatment of employees, about the importance of speed—and most importantly, about being committed to hearing the truth, whatever that truth is. That is the biggest barrier to implementing evidence-based management: the shift required in how leaders think about their job and the process of getting work done. This is a barrier that exists in large, traditional companies but also persists in small, entrepreneurial ventures. Overcoming the traditional mindsets can, as in the case of agile software development, lower risk by getting better market data more quickly and lower the waste of resources through a leaner, more efficient process. As such, this evidence-based approach can provide a competitive advantage, but only to those people with the wisdom to use it.

## CONCLUSION

Two things seem to be true. First, evidence-based management could improve entrepreneurial decision-making, reducing risks, costs, and wasted time and effort—just as an evidence-based approach could benefit most if not all organizations and just as evidence-based medicine has improved medical practice and outcomes while saving money. Second, the mindset shift required to implement EBM is apparently large. Therefore, evidence-based approaches struggle even when they could demonstrably save vast sums of money and even lives.

But as the recent history of the lean or agile software movement illustrates, the competitive advantages from listening to the data are substantial. And in the end, much like medicine and various branches of public policy, particularly in countries other than the United States, the implementation of an evidence-based approach will gain traction. It is just a matter of time. In the meantime, however, those entrepreneurs and suppliers of risk capital who avail themselves of evidence-based thinking will be in a much stronger competitive position.

## REFERENCES

- Blume, Brian D., Baldwin, Timothy T. and Rubin, Robert S. (2009) "Reactions to Different Types of Forced Distribution Performance Evaluation Systems," *Journal of Business and Psychology*, 24: 77–91.
- Brown, Tim (2008) "Design Thinking," *Harvard Business Review*, 86 (June): 84–92.
- Chambers, John R., and Windschitl, Paul D. (2004) "Biases in Social Comparative Judgments: The Role of Nonmotivated Factors in Above-average and Comparative-optimism Effects," *Psychological Bulletin*, 130: 813–838.

## 68 JEFFREY PFEFFER

- Cochrane, John H. (2005) "The Risk and Return of Venture Capital," *Journal of Financial Economics*, 75: 3–52.
- Domurad, Frank (2005) "Doing Evidence-based Policy and Practices ain't for Sissies," *Community Corrections Report on Law and Corrections Practice*, 12: 49–50ff.
- Heath, Chip (1995) "Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting," *Organizational Behavior and Human Decision Processes*, 62: 38–54.
- Kaplan, Steven N. and Schoar, Antoinette (2005) "Private Equity Performance: Returns, Persistence, and Capital Flows," *Journal of Finance*, 60: 1791–1823.
- Kelley, Tom (2001) *The Art of Innovation*. New York: Currency.
- Kruger, Justin (1999) "Lake Wobegon be Gone! The 'Below-Average Effect' and the Egocentric Nature of Comparative Ability Judgments," *Journal of Personality and Social Psychology*, 77: 221–232.
- Levinthal, Daniel A. and March, James G. (1993) "The Myopia of Learning," *Strategic Management Journal*, 14: 95–112.
- Levitt, Barbara and March, James G. (1988) "Organizational Learning," *Annual Review of Sociology*, 14: 319–340.
- Mackintosh, James (2009) "Hedge Fund Investors Have a Great Chance to Cut Fees," Ft.com, February 6.
- Martin, Roger (2009) *The Design of Business*. Boston: Harvard Business Press.
- Mason, Colin M. and Harrison, Richard T. (2002) "Is it Worth it? The Rates of Return from Informal Venture Capital Investments," *Journal of Business Venturing*, 17: 211–236.
- Meacham, John A. (1983) "Wisdom and the Context of Knowledge: Knowing What One Doesn't Know One Doesn't Know," in D. Huhn and J. A. Meacham (Eds.), *On The Development of Developmental Psychology*. Basel: Krager, 111–134.
- Novations Group (2004) "Uncovering the growing disenchantment with forced ranking performance management systems." Boston: Novations Group, White Paper.
- Pfeffer, Jeffrey and Sutton, Robert I. (1980) *The Knowing-Doing Gap: How Smart Companies Turn Knowledge into Action*. Boston: Harvard Business School Press.
- Pfeffer, Jeffrey and Sutton, Robert I. (2006) *Hard Facts, Dangerous Half-Truths, and Total Nonsense: Profiting from Evidence-Based Management*. Boston: Harvard Business School Press.
- Rosenzweig, Phil (2007) *The Halo Effect*. New York: Free Press.
- Shane, Scott (2008) Startup failure rates—the REAL numbers. Available at: <http://smallbiztrends.com/2008/04/startup-failure-rates.html>.
- Stark, Antony (2001) SME support in Britain. Available at: [www.adb.org/Documents/Reports/PRC-SME/App5-UK.pdf](http://www.adb.org/Documents/Reports/PRC-SME/App5-UK.pdf).
- Staw, Barry M. (1975) "Attribution of the 'Causes' of Performance: A General Alternative Interpretation of Cross-sectional Research on Organizations," *Organizational Behavior and Human Performance*, 13: 414–432.
- Staw, Barry M. (1976) "Knee Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action," *Organizational Behavior and Human Performance*, 16: 27–44.
- Sternberg, Robert J. (1985) "Implicit Theories of Intelligence, Creativity, and Wisdom," *Journal of Personality and Social Psychology*, 49: 607–627.
- Tetlock, Phillip E. (2000) "Cognitive Biases and Organizational Correctives: Do Both Disease and Cure Depend on the Politics of the Beholder?" *Administrative Science Quarterly*, 46: 293–326.