
The Wisdom and Whims of the Collective

Michael J. Mauboussin
Chief Investment Strategist
Legg Mason Capital Management
Baltimore

The debate over whether financial markets are efficient has raged for decades. Proponents of market efficiency point to the difficulty of outperforming markets to support their view, whereas opponents note the frequency of mispricings to support theirs. The resolution to the dilemma may lie in the dynamics of group decision making.

It is my ambitious goal to unify two seemingly irreconcilable camps. The first camp consists of people who believe markets are efficient. This group includes many academics and investment industry luminaries. The practical recommendation from this camp is simply to buy an index fund, which is, indeed, very good advice for most people. The second camp includes people in the behavioral finance group, who argue that people behave suboptimally, creating so-called market anomalies. I am going to try to reconcile these two camps by proposing that the market is a complex adaptive system. This is a compelling view for a number of reasons:

- First, it allows us to relax the assumption of rationality either for all investors or even a subset of arbitrageurs.
- Second, it sets out conditions under which markets tend to be efficient and conditions under which markets are inefficient.
- Finally, it accommodates empirical realities, particularly fat-tailed price distributions.

I will first discuss the classic mean–variance and no-arbitrage-opportunity explanations for market efficiency. Then, I will explain the concept of markets as a complex adaptive system. Finally, I will wrap up with what I call the “whims of the collective,” which explains why markets periodically go to excess.

Some wonder whether we can understand the market by adding up the actions of the individual agents. In other words, are markets Newtonian (or reductionist) in the sense that the whole must equal the sum of the parts? The answer to that question, I believe, is a resounding no. Markets have properties

and characteristics that emerge from the interaction of investors that cannot be anticipated by looking at the underlying investors themselves. I like to think about this issue on three levels, each with its own unit of analysis.

We must first consider psychology, in which the level is the individual and the unit of analysis is the individual. Next, we must understand social psychology, in which the level of analysis is the group and the unit of analysis is the individual. By looking at these levels, we are attempting to gain an understanding of how individuals operate as parts of a group.

The final level is sociology, which studies groups and how groups behave. Here, the unit of analysis is the group. Notice, by the way, that these are all distinct disciplines, and we should be very careful about maintaining these distinctions. My view is that many market inefficiencies are largely sociological.

Market Efficiency

Let me start with a common sense observation about market efficiency: It is clearly hard for active managers to outperform the market, which has been true for as long as anyone has tracked the results of active managers. This may not mean that markets are efficient in any sort of technical sense, but it does suggest that they are practically efficient. That is, there are no simple, systematic ways to outperform the market.

In contrast, we obviously do see periodic, and sometimes dramatic, excesses. Efficient market proponents would be hard pressed to argue that the dot-com era of the late 1990s was a time when the markets were efficient in the sense that asset prices represented an unbiased estimate of asset values.

This presentation comes from the Efficient Market and Behavioral Finance conference held in Boston on 5–6 June 2007.

Finally, there is academic evidence that some investors do outperform the market over time. I personally believe these people exist, but I am not at all persuaded that their skill sets are transferable.

There are essentially three ways to explain market efficiency. The first approach stipulates the existence of rational market participants who understand their risk preferences and the distribution of asset prices and who rationally trade off risk and reward. The second approach relies on the ability of arbitrageurs to eliminate arbitrage opportunities. This is an important step in the right direction because now I can relax the assumption that all investors must be rational. Instead, all I really need to achieve market efficiency is for a subset of the population—arbitrageurs—to be rational. Arbitrageurs will identify price discrepancies and move to correct them by buying the undervalued assets and selling the expensive assets. The third approach relies on the wisdom of crowds. After all, it is really the aggregation of individuals who may or may not be rational that allows us to get to the right price. The wisdom-of-crowds approach does not exclude the notion of arbitrage, but it does not require it either.

The next question is whether the assumptions supporting each of these approaches are reasonable and empirically valid. In my opinion, the assumptions underlying the mean–variance approach are neither. In contrast, there seems to be at least mixed support for the no-arbitrage opportunity assumption. Finally, I would suggest that the wisdom-of-crowds approach is both a reasonable assumption and empirically valid.

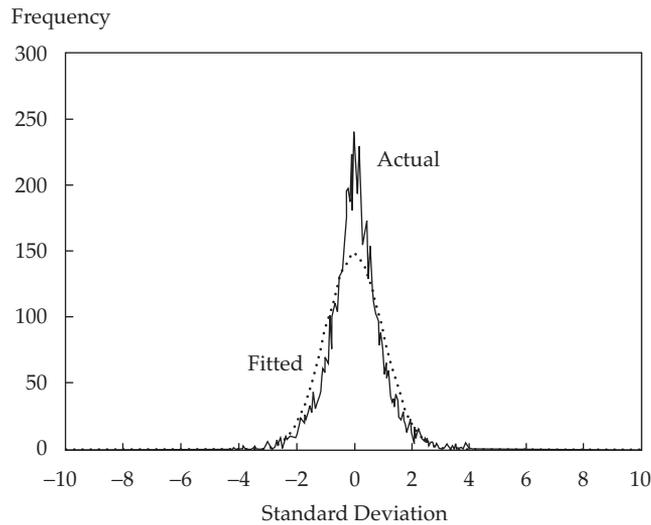
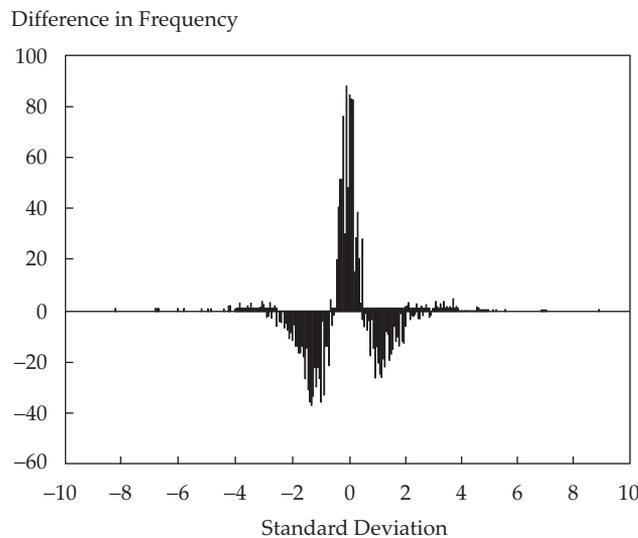
Tests of Mean–Variance/Market Efficiency Assumptions. Just how closely does the traditional model of how markets work conform to empirical reality? The answer is, unfortunately, not as close as we would like. Consider the capital asset pricing model (CAPM), which is an elegant theory that has been taught to MBA students for the past 30 years. But in a paper from the early 1990s, Eugene Fama and Kenneth French (1992) concluded that “tests do not support the central prediction of the SLB [Sharpe–Lintner–Black] model, that average stock returns are positively related to market β ” (p. 449). More recently, Fama and French (2004) added that “the empirical record of the [SLB] model is poor—poor enough to invalidate the way it is used in applications” (p. 25).

Importantly, one of the simplifying assumptions underlying the efficient market hypothesis is that stock price changes are normally distributed. Academics and practitioners have known for a long

time that this is not the case but have, to some degree, understated its significance. Panel A of **Figure 1** shows actual daily changes in the S&P 500 Index from January 1978 to March 2007. The solid line in Panel A is the actual distribution of price changes. The dotted line represents a fitted normal distribution based on the standard deviation of the actual returns. Looked at in this manner, the actual returns are deceptively normal. But to emphasize the extent to which the actual distribution is non-normal, in Panel B, we plotted the difference between the normal distribution and the actual outcomes. With this view, three observations become noteworthy. First, many more small-change days occurred than expected. Second, fewer medium-sized-change days occurred than normality suggests. Third, and perhaps most importantly, there were many more fat-tailed days than there should have been. As can be seen, there are 5, 6, and 7 sigma events represented here.

To give a sense of just how important the fat-tailed events are, I did a “knockout” study. There are roughly 7,300 observations in Figure 1. The return in the market during this period was about 9.5 percent (excluding dividends). In the actual distribution, if I knock out the 50 worst days (less than 7/10ths of 1 percent of the sample), the return goes from 9.5 percent to 18.2 percent. These few bad days can have an enormous impact. In contrast, if I knock out the 50 best days, the return is 1 percent. Just to give some standard of reference, by simulating a normal distribution based on these data, if I knock out the 50 worst days from my simulation, the return is 15.2 percent—up from 9.5 percent but much less dramatic than the change to 18.2 percent. Alternatively, if I knock out the 50 best days, return is 3.5 percent. The point that I want to emphasize is that fat-tailed events have a material impact on returns over time. The inescapable conclusion is that the assumption of normality inherent in the mean–variance framework represents a misleading model of how markets actually work.

The Role of “Sharks.” Most financial economists do not really believe the assumption that all market agents are rational. But at the same time, they do not think that needs to be the case to achieve market efficiency. For the most part, they agree that we need only a handful of these agents to be smart, which is obviously much more possible and plausible. The core idea here is that the act of arbitrage eliminates the very opportunities it seeks to exploit. Essentially, if two assets are obviously inappropriately priced relative to each other, a well-heeled arbitrageur will buy one and sell the other, thus

Figure 1. Frequency Distribution, January 1978–March 2007**A. Frequency Distribution of S&P 500 Daily Returns****B. Frequency Difference: Normal vs. Actual Daily Returns**

Source: FactSet.

eliminating the arbitrage opportunity. In an address to the European Mortgage Finance Agency, Stephen Ross expressed this concept well:

Neoclassical finance is a theory of sharks and not a theory of rational homo economicus. In liquid securities markets, though, profit opportunities bring about infinite discrepancies between demand and supply. Well financed arbitrageurs spot these opportunities [and] pile on, and by their actions they close aberrant price differentials. Rational finance . . . has worked very hard to rid the field of its sensitivity to the psychological vagaries of investors. (Ross 2001)

It is also worth noting the mean–variance framework and the elimination of arbitrage opportunities are fundamentally different approaches to market efficiency. The rational model relies on notions of general equilibrium and is an absolute-pricing model. It will allow one to achieve allocative efficiency in markets. In contrast, arbitrage is a relative-pricing model. It does not ask where prices come from. It only seeks to exploit relative price differentials. The notion that one is absolute and the other is relative is a very important distinction to keep in mind.

Behavioral finance has been correct to point out that there are limits to arbitrage, such as the inability to short, lack of appropriate substitutes, the fact that arbitrage is not always riskless, and implementation costs. We also know that arbitrageurs are not always there when we need them. One of the best examples in recent memory is what happened to Long-Term Capital Management (LTCM) in the late 1990s. One of the more thoughtful accounts of the LTCM events comes from economic sociologist Donald MacKenzie in his book *An Engine Not a Camera* (which I would highly recommend if you are interested in these topics). Here is what he writes in a paper included in the book:

As “spreads” widened [between off-the-run and on-the-run bonds], and thus arbitrage opportunities grew more attractive, arbitrageurs did not move into the market, narrowing spreads and restoring “normality.” Instead, potential arbitrageurs continued to flee, widening the spreads and intensifying the problems of those who remained, such as LTCM. (MacKenzie 2004, p. 30)

Mean–variance and arbitrage represent reasonable approximations of reality, but only partly. Both of these models rest on implausible assumptions, and both models fail at critical junctures or in critical ways. I would like to develop an understanding of markets that addresses these shortcomings and introduce the concept of markets as a complex adaptive system.

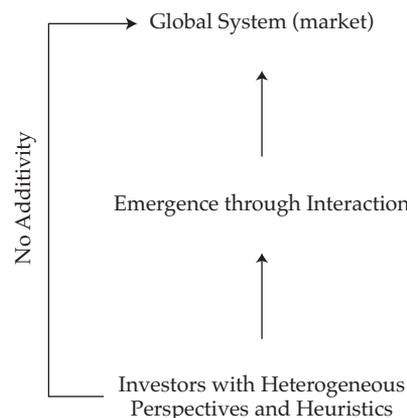
Complex Adaptive System

“Complex” simply means there is a lot of interaction. “Adaptive” means that the agents change and evolve. “System” means that the whole is greater than the sum of the parts.

Figure 2 shows a picture of how this might work. Investors interact with one another. As a result of this interaction, an emergent phenomenon develops that has properties and characteristics distinct from the underlying investors. Complex adaptive systems are ubiquitous in nature and society. Examples include ant colonies, consciousness through neuron activity, and the City of Boston.

One point I would like to stress is that there is no additivity in these systems. That is, the system is not the simple sum of the individual participants. Rather, the macro system has properties and characteristics distinct from the underlying participants. Just as you cannot understand how an ant colony works by interviewing ants, you cannot understand the workings of markets by observing individuals within a market. Note that a complex adaptive system allows for arbitrageurs, but at the same time, it does not require them.

Figure 2. Common Features of Complex Adaptive Systems



The Wisdom of Crowds

Let me give some examples of the wisdom of crowds. In the class I teach at Columbia Business School, I pass around a jar of jellybeans, and there is a \$20 reward for the student who comes closest to guessing the actual number of jellybeans. In the most recent class, the jellybean jar had 1,116 beans and the consensus was 1,151—off by only about 3 percent. But significantly, only 2 of the 73 students did better than the consensus.

Here is another example I do with my class. About three weeks before the Academy Awards, I hand the students a form. On the front of the form are six major categories, such as best actress, best actor, and so on. On the back are six much more remote categories, such as best cinematography, best music—things that few people pay much attention to. Again, in my most recent class, the consensus correctly named 11 out of 12 winners, although the best student named only 9 out of 12. More interesting, though, the average student guessed only 5 out of 12 winners. So, the consensus was vastly better than the average, but more importantly, it was better than even the best individual.

It is worth noting that financial markets are very different in at least two ways from these examples. In markets, there is no right answer. No one stops trading at the NYSE and says, “By the way, IBM is truly worth \$120 a share; thank you very much for showing up today.” There is also no time horizon. The perpetuity dimension is very important. That said, I hope these examples give some sense of how the wisdom of crowds can work in market-related settings.

An interesting new book that addresses the “why” of the wisdom of crowds is Scott Page’s *The Difference* (2007). One of Page’s core ideas is the “diversity prediction theorem,” which says the *Collective error* = *Individual error* (the ability of the individual) – *Prediction diversity* (cognitive diversity).

The diversity prediction theorem has some important implications that can be applied to the previous examples. (1) The collective is smarter than the average person within the collective. Not sometimes, but always. This means that as an investor, you have to recognize that if you are buying or selling at the market price, unless you are above average, the market will be smarter than you are. (2) Collective error is determined in part by ability and in part by diversity. Diversity, therefore, plays a very important role in collective accuracy. (3) The collective is usually better than even the best of the individuals.

It is important to consider the conditions the wisdom of crowds requires in order to work. When one of these conditions is violated, that is where we will find breakdowns in market efficiency.

Conditions Needed. The first condition is diversity of the underlying agents. Many firms have diversity initiatives that focus on social identity diversity—race, ethnicity, age, and so on. These could be considered proxies for this condition, but they do not reflect it directly. In a market context, we are looking for cognitive diversity. The dimensions that really matter are things like perspectives on a problem, rules of thumb, and interpretation.

The second condition is an aggregation mechanism that brings information together. The financial markets are the aggregation mechanism we are most familiar with.

The third condition is the presence of incentives—rewards for being right and penalties for being wrong. In our society, monetary incentives are the most common. But incentives can also be other things, such as reputational.

Understanding Diversity Breakdowns. In my opinion, the key to resolving the market efficiency debate relies on recognizing that inefficiency prevails when one or more of these conditions are violated. Moreover, the most likely condition to be violated is diversity, which occurs when we start to act in a coordinated fashion.

There are two broad types of diversity breakdowns. The first is what I call “dramatic” diversity breakdowns, usually accompanied by abrupt price movements. The most famous examples of this type of breakdown are the South Sea bubble, tulip mania, and most recently, the internet bubble. These breakdowns tend to occur when positive feedback produces extremely optimistic or pessimistic sentiment about a stock or an asset class. Ultimately, the fundamental view is clearly overstretched, and it creates a set of expectations that are out of sync with the fundamentals.

The second type of diversity breakdown, time arbitrage, is much more subtle. It occurs when short-term prices reflect more noise than signal.

To illustrate a dramatic diversity breakdown, consider **Figure 3** as it relates to two *Wall Street Journal* articles about a barber in Cape Cod, Massachusetts.

Figure 3. NASDAQ Performance, 1995–2005



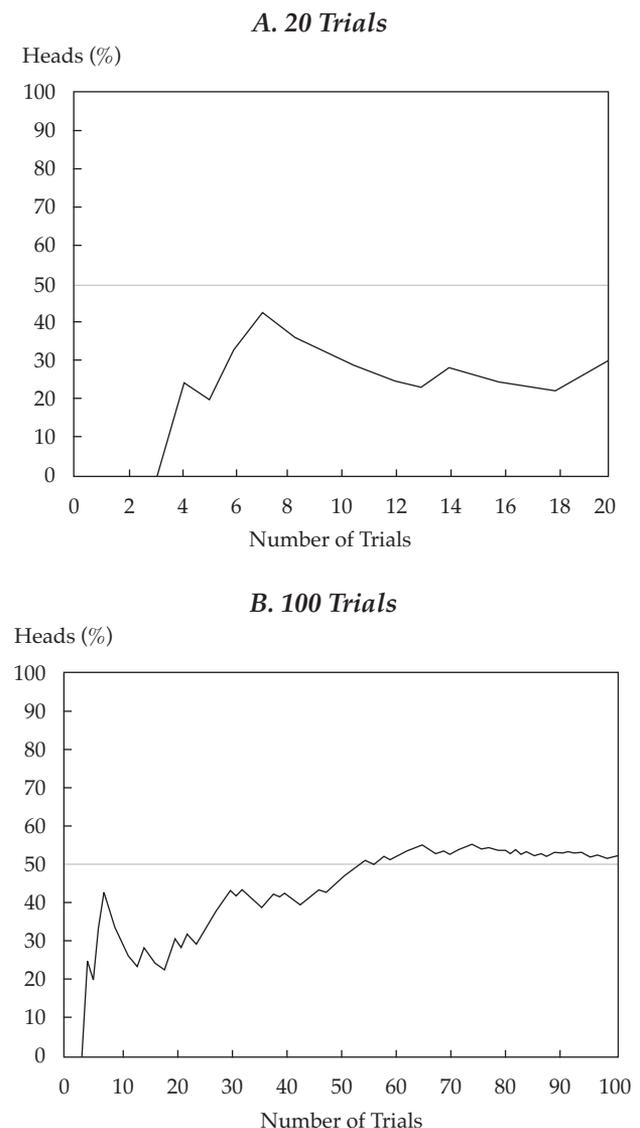
On 13 March 2000, the *Wall Street Journal* published the article “Stock Chit-Chat Enriches Many Cape Cod Locals” about how the barber had become a local celebrity because he had put many of his clients into NASDAQ stocks that had done spectacularly well. He had made hundreds of thousands of dollars in the market. CNBC was on television whenever he was cutting hair, and he was anticipating an early retirement. Here’s a quote from the article, “. . . I don’t think anything could shake my confidence in this market . . . even if we go down 30 percent, we’ll just come right back.” Unfortunately, this extreme confidence came right as the NASDAQ was peaking and could only lead to a bad outcome. When the NASDAQ was near its ultimate low, the *Wall Street Journal* went back to the same barbershop. On 8 July 2002, the story “Cape Cod Barber Shop, Slumping Stocks Clip Buzz” told about how the barber was doing. Needless to say, he was despondent, as were his clients. He would have to cut hair for a lot longer because he could not retire early. CNBC had been turned off, and he summed up the complete change in his attitude by saying, “All they ever say is buy, buy, buy, all the way down from \$100 a share to bankruptcy.” When they then asked if he was in the market anymore, he said he had given up on the market altogether and was now going to the casino for his investments.

A similar, albeit less dramatic, example of the same phenomenon can be found in the March/April 2007 issue of the *Financial Analysts Journal*, which featured an article on the predictive value of magazine covers (Arnold, Earl, and North 2007). The authors reviewed cover stories in *BusinessWeek*, *Forbes*, and *Fortune* from the last 20 years and categorized the covers from the most bullish (optimistic spin) to the most bearish (pessimistic spin). They found that if they looked back two years prior to the date of covers in the first category (the bullish outlook), stocks had outperformed their peers by 42 percentage points (pps). In contrast, years prior to covers in the bearish category showed that stocks had underperformed their peers by 35 pps. So, there was massive overperformance and massive underperformance leading up to the featured magazine articles. But over the next two years, the stocks of the companies in the bearish articles outperformed those in the bullish articles by a 3-to-1 margin.

Turning our attention to the other type of diversity breakdown, time arbitrage is an important concept in the context of financial markets. Understanding time arbitrage requires recognizing that any probabilistic system has a random component (noise) and a signal. One of the easiest illustrations of this is a coin toss. If you flip a fair coin

countless times, you would expect to see 50:50 as a signal. But if you flipped it only 10 times on three occasions, you would see vastly different results. I believe that at times, the market extrapolates the noise as if it were the real signal, which creates a time arbitrage opportunity. Consider the following simple illustration, shown in **Figure 4**. Panel A shows the plot of 20 trials of a coin toss, and the percentage of heads is only about 35 percent. Panel B is the same trial extended to 100 tosses, and it settled not too far from 50:50. In a market context, if investors are being short-term oriented and see the first 20 tosses, they may conclude that it is a tails-biased coin and price the asset accordingly. But if investors understand what the true signal is, then it would represent a

Figure 4. Illustrating Noise vs. Signal



Source: Mauboussin (2005).

time arbitrage opportunity and they should buy the asset. Although this sounds simple enough, be aware that two significant things have to happen for it to work. First, you must recognize the true long-term return signal for the business, and second, the signal must reveal itself. Provided these two things happen, time arbitrage represents an important source of potential diversity breakdown.

Now that we have taken this journey, I come back to my original goal of reconciling efficient markets with behavioral finance. I think both of these theories can sit very comfortably inside the tent of complex adaptive systems. Briefly, markets are efficient until diversity breaks down. But taking advantage of diversity breakdowns can be extremely difficult for two reasons. First, we must be prepared to run counter to the group, which is very difficult psychologically. Being part of the group has clear and profound evolutionary foundations. The second reason is agency problems. It is far safer for us to run our money management firms as businesses gathering assets and collecting fees than necessarily trying to deliver excess returns for our clients.

Taking Advantage of Diversity Breakdowns.

Assuming that this is the right way of thinking about the world, how does an active manager take

advantage of it to try to outperform the market? Following are some thoughts:

- Look for flags or signatures of diversity breakdowns, such as extreme bullishness or extreme bearishness. Also, look for simple indicators that stocks are trading on the cheap or expensive side. If a stock or an asset class is trading way out of its historical range, that might be an interesting indicator.
- Further analyze these diversity breakdown candidates by considering what expectations the stocks or financial assets reflect. Recognize that there is an important distinction between fundamentals, which is the value of that asset (present value of future cash flows), and expectations, which is the price. Try to keep those things separate in your own mind, and try to think probabilistically.
- Try to create a proper context for decision making within your organization. The hardest part is avoiding the psychological traps that organizations often compound rather than alleviate.
- Finally, be receptive to novel approaches using multidisciplinary techniques to help determine where excesses might exist or where diversity breakdowns may occur.

This article qualifies for 0.5 CE credits.

REFERENCES

Arnold, Tom, John H. Earl, Jr., and David S. North. 2007. "Are Cover Stories Effective Contrarian Indicators?" *Financial Analysts Journal*, vol. 63, no. 2 (March/April):70–75.

Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, vol. 47, no. 2 (June):427–465.

———. 2004. "The Capital Asset Pricing Model: Theory and Evidence." *Journal of Economic Perspectives*, vol. 18, no. 3 (Summer):25–46.

MacKenzie, Donald. 2004. "Models of Markets: Finance Theory and the Historical Sociology of Arbitrage." Working paper (January).

Mauboussin, Michael J. 2005. "Capital Ideas Revisited—Part 2." Mauboussin on Strategy (20 May): www.leggmasoncapgmt.com/pdf/CIR2.pdf.

Page, Scott. 2007. *The Difference*. Princeton, NJ: Princeton University Press.

Ross, Stephen. 2001. "Neoclassical and Alternative Finance." European Finance Agency Meetings, keynote address.

Question and Answer Session

Michael J. Mauboussin

Question: How can organization leaders create smart collectives in the form of work groups and teams within an organization?

Mauboussin: There are a couple of things leaders can do. First, they should collect a group that is diverse, recognizing that diverse does not mean that they look different but that they think differently or have different perspectives. Second, leaders must effectively solicit the opinions and views of everybody. Finally, leaders should not impose their view at any point, until perhaps the very end.

It is also important to recognize the type of problem you are trying to solve. Straightforward problems, such as fixing a leaky

faucet, are best taken care of by relying on experts. In contrast, collectives work best on unstructured, complicated problems.

Question: Will the increased connectivity caused by technological advances, such as the Blackberry, lead to more frequent diversity breakdowns?

Mauboussin: It creates a tug-of-war. On the one hand, more people have more access to capital markets than ever before, and all things being equal, that should inject diversity. On the other hand, we all have access to the same popular media, such as CNBC, and technology, such as Blackberrys. These are potentially diversity detractors. My

own sense as an investor is that diversity prevails.

Question: Is regression to the mean a precept we can rely on to produce alpha?

Mauboussin: It depends. I would say that microeconomics works, which means that, over time, competitive forces drive returns down to the cost of capital. I have not seen an instance where that has not happened sooner or later for every corporation. So, certainly on a micro level, I would think there is some truth to that. But we have seen evidence, at least for periods of time, of firms that have succeeded in increasing returns.