

*How has “The Mechanisms of Market Efficiency”
fared over the past 20 years?*

MOME in Hindsight

BY RONALD J. GILSON, *Stanford Law School*
and REINIER KRAAKMAN, *Harvard Law School*

Two decades ago, the *Virginia Law Review* published our article “The Mechanisms of Market Efficiency” (MOME), in which we tried to discern the institutional underpinnings of financial market efficiency. We concluded that the level of market efficiency with respect to a particular fact depends on which of several market mechanisms — universally informed trading, professionally informed trading, derivatively informed trading, and uninformed trading (each of which we explain below) — operates to reflect that fact in market price. Which mechanism is operative, in turn, depends on how widely the fact is distributed among traders, which, in turn, depends on the cost structure of the market for information. Less costly information is distributed more widely, triggers a more effective efficiency mechanism, and is reflected more efficiently by market prices.

Revisiting our article is particularly appropriate today. A new framework for evaluating the efficiency of the stock market called “behavioral finance” and a growing number of empirical studies pose a serious challenge to the Efficient Markets Hypothesis. Michael Jensen’s 1978 statement that “there is no proposition in economics which has more solid empirical evidence supporting it than the Efficient Market Hypothesis” is now proffered with a tone somewhere between irony and condemnation.

The movement away from Jensen over the past few decades surely merits a reconsideration of the substance and policy implications of financial market efficiency. We remain convinced that the quickness and accuracy with which the stock market reflects information in the price of a security is a function of the performance of institutions. Twenty years, however, have made us appropriately more skeptical of the efficiency of those institutions.

Ronald J. Gilson is the Meyers Professor of Law and Business at Stanford Law School and the Stern Professor of Law and Business at Columbia Law School.

Reinier Kraakman is the Ezra Ripley Thayer Professor of Law at Harvard Law School.

This article is based on a longer paper that appeared in the 2003 *Journal of Corporate Law*.

THE RISE OF MODERN FINANCE

To place MOME in its proper context, we first need a capsule account of the development of modern finance. We will focus on three bodies of theory that sought to state rigorously how capital assets are priced, whether a corporation’s choice of debt and equity financing instruments affects its value, and whether the market price of freely traded securities reflects all available information concerning their value.

Those three familiar theories — the Capital Asset Pricing Model, the Miller-Modigliani Irrelevance Propositions, and the Efficient Capital Market Hypothesis — share an extensive set of perfect markets assumptions: rational investors, perfect information, and no transaction costs.

Start with the Capital Asset Pricing Model (CAPM). If one assumes that all unsystematic risk (risks that affect some assets but not others in the economy) can be eliminated through ownership of a diversified portfolio of investments, what else but systematic risk (risks that affect all assets simultaneously) could affect the price of capital assets? If investors need not bear unsystematic risk because they are diversified, then investors who do not bear it will require the lowest return (pay the highest price) for a capital asset, thereby setting the asset’s price. CAPM simply takes the next step and argues that the systematic risk that matters to investors is the covariance of an asset’s returns with those of the market — i.e., beta. Given those assumptions, CAPM is a tautology.

The Miller-Modigliani Irrelevance Propositions share the same conceptual structure. Like CAPM, the perfect capital market assumptions result in the Irrelevance Propositions appearing tautological. If debt or equity was mispriced, arbitrage would restore the proper relation so that increasing the amount of lower-cost debt would result in an offsetting increase in the cost of equity, and vice versa.

The Efficient Capital Market Hypothesis (ECMH) also builds on perfect market assumptions. As William Sharpe wrote, “Simply put, the thesis is this: that in a well-functioning securities market, the prices . . . of securities will reflect predictions based on all relevant and available information. This seems to

be trivially self-evident to most professional economists — so much so that testing seems almost silly.”

In addition to its prediction of the information content of stock prices, the ECMH also played a critical integrative role, providing the necessary link between asset pricing and capital structure choice through the medium of market prices. Both CAPM and the Modigliani-Miller propositions depend on an arbitrage mechanism for their proof: mispricing will be traded away. But for arbitrage to be triggered by mispricing, market prices must be reasonably informative. Thus, along this important dimension, the positive powers of the three theories rise and fall together.

THE MOME THESIS

Financial market efficiency, as we saw it, concerned how rapidly prices responded to information. By the early 1980s, a large body of empirical work demonstrated that price responded extremely rapidly to public and even “semi-public” information — too rapidly to permit arbitrage profits on most of that information. But how was this possible, given that most traders were likely to be uninformed about the content of much of this information?

We answered that question on two levels. On the level of the capital markets, MOME proposed that four mechanisms work to incorporate information in market prices with progressively decreasing relative efficiency. In “universally informed trading,” market prices immediately reflect information that all traders know simply because that information necessarily informs all trades, just as perfect market theorists assumed. In “professionally informed trading,” information that is less widely known but nonetheless public is incorporated into share prices almost as rapidly as information known to everyone through the trading of savvy professionals. In “derivatively informed trading,” inside information known to only a very few traders would find its way into prices more slowly, as uninformed traders learned about its content by observing tell-tale shifts in the activity of presumptively informed traders or unusual price and volume movements. Finally, in “uninformed trading,” information known to no one might be reflected, albeit slowly and imperfectly, in share prices that aggregate the forecasts of numerous market participants with heterogeneous information.

MOME’s second claim was that cost determines the distribution of information in the market. The cost of information, in turn, depends on the market institutions that produce, verify, and analyze information — institutions that range from the *Wall Street Journal* to the exhaustive research of the best professional investors.

THE CHALLENGE OF BEHAVIORAL FINANCE

Beginning in the 1980s, a growing literature challenged the empirical predictions of the 1960s’ perfect market theorems. That challenge, in turn, gave rise to a reassessment of the underlying theory of perfect markets.

EMPIRICAL ANOMOLIES A focus on imperfections in the market for information sparked a series of explanations of how capital structure (debt vs. equity) could matter if information was costly and asymmetrically distributed. If corporate managers



had private information concerning the corporation's future prospects, and if bankruptcy is costly to managers, then exposing the corporation to a greater risk of bankruptcy either by paying dividends or maintaining a higher debt-to-equity ratio could credibly signal that information to the market and thereby influence the price of the corporation's securities. Correspondingly, capital structure could also function as an incentive: an increased risk of bankruptcy resulting from a more leveraged capital structure could provide an incentive for managers (for whom bankruptcy would be costly) to work even harder.

The Capital Asset Pricing Model always had problems when attention shifted from theory to empirical testing. A security's beta does not predict its return very well, as two categories of

advantage of mispricing, thereby driving prices back to their proper level. Behavioral finance generalizes from the accretion of empirical discrepancies to argue that many investors are not rational in their financial decision-making, that there are observable directional biases resulting from departures from rational decision-making, and that significant barriers prevent professional traders from fully correcting the mistakes made by less-than-rational investors.

COGNITIVE BIASES The criticism of the rationality premise builds on an important literature that uses decision-making experiments to show how individuals' cognitive biases can lead them to systematically misassess an asset's value. The list of

Behavioral finance argues that significant barriers prevent professional traders from fully correcting the mistakes made by less-than-rational investors.

evidence demonstrate. First, studies show that asset-pricing models with multiple factors in addition to systematic risk do a better job of predicting prices than CAPM does alone. Eugene Fama and Kenneth French, for example, found that they could better predict market prices with a three-factor pricing model that includes company size and book-to-market ratio in addition to systematic risk.

Second, CAPM's empirical failures exhibit certain regularities. The literature identifies a number of what are styled "anomalies" — that is, persistent evidence of higher-than-predicted returns based on publicly available information. Such anomalies include the tendency of small companies to earn higher-than-predicted returns, the seeming existence of a "January effect" (in which much of the abnormal returns to smaller firms occurs during the first half of January), the "weekend effect" (in which stock returns are predictably negative over weekends), and the "value effect" (in which firms with high earnings-to-price ratios, high dividend yields, or high book-to-market ratios earn higher-than-predicted returns).

Various explanations have been offered for the empirical discrepancies. Some commentators attribute them to incorrect asset pricing models. Others note that the studies revealing the anomalies are sensitive to the particular empirical techniques used, or demonstrate that at least some of the anomalies disappear or are dramatically reduced in size following their announcement in the literature, thus suggesting that markets learn, although not necessarily quickly.

Nevertheless, the accumulation of anomalies has had an effect, despite efforts to explain those anomalies in finance-friendly ways. The core theories of modern finance assume that investors are fully rational (or that the market acts as if they are) and that markets are efficient and transaction costs small, so that professionally informed traders quickly notice and take

biases has grown impressively with time and includes overconfidence (the tendency of individuals to overestimate their skills), the endowment effect (the tendency of individuals to insist on a higher price to sell something they already own than to buy the same item if they do not already own it), loss aversion (the tendency for people to be risk averse for profit opportunities but willing to gamble to avoid a loss), anchoring (the tendency for people to make decisions based on an initial estimate that is later adjusted but not sufficiently to eliminate the influence of the initial estimate), framing (the tendency of people to make different choices based on how the decision is framed, such as whether it is framed in terms of the likelihood of a good outcome or in terms of the reciprocal likelihood of a bad outcome), and hindsight (the tendency of people to read the present into assessments of the past).

Individuals whose decisions are subject to one or more of those biases, referred to in the literature as "noise traders," make investment decisions that deviate from those that theory would predict of rational investors. Charles Lee, Andrei Shliefer, and Richard Thaler, in a 1991 *Journal of Finance* article, argued that the discount often associated with closed-end mutual funds, one of the longstanding phenomena that conflict with the ECMH, illustrates the potential for misguided investors to influence price efficiency. When an investor sells shares in a closed-end mutual fund, she receives whatever a buyer is willing to pay rather than a proportionate share of the fund's net asset value (as she would if she redeemed her interest in an open-end mutual fund). Because the net assets of a closed-end fund are observable, the ECMH predicts that the fund's stock price should reflect its net asset value. In fact, closed-end funds systematically (but not uniformly) trade at a discount from their underlying asset value, a serious problem for the claim that stock prices generally are the best estimate of a security's value.

In the one case where we can actually observe underlying asset value, stock price diverges from it.

Lee, Shleifer, and Thaler blame the divergence on noise traders, whose views about value — perhaps because of some or all of the cognitive biases reviewed above — plainly ignore the value of securities held by closed-end funds, as they are reflected in market prices. But two other elements are also necessary. First, the biases held by the noise traders must be more or less consistent; otherwise, at least some of the biases will regress out. Second, arbitrageurs must be unwilling to police the resulting price inaccuracies. Under perfect capital market assumptions, fully informed traders with unlimited access to capital immediately pounce on mispriced securities. If arbitrageurs were available to trade against the noise traders, then their action would suffice to return prices to their efficient level. In the case of closed-end mutual funds, however, the absence of institutional investors in this niche limits the extent of corrective arbitrage, and prices retain a rational component reflecting the risk of noise trader irrationality.

LIMITS ON ARBITRAGE Limited arbitrage is critical to the behavioral finance perspective, and the problem is more general than the simple case of closed-end mutual funds. Limits on arbitrage fall into four general categories:

- fundamental risk;
- noise trader risk;
- institutional limits, both regulatory and incentive; and
- the potential that even professional traders may be subject to cognitive biases.

The problem of fundamental risk arises because the arbitrageur, unless hedged, has a position in the stock of a particular company that is exposed to loss from a change in the company's fortunes. That risk can be counterbalanced by holding an offsetting position in a substitute security. However, substitutes may not be available, and in all events will be imperfect. Imagine an arbitrageur who believes that Ford is underpriced. For example, suppose an arbitrageur shorts General Motors to hedge the risk associated with purchasing Ford. This strategy works, but it only hedges against bad news in the automobile industry generally; it does not hedge against firm-specific bad news about Ford (and, to the extent that bad news for Ford is good news for General Motors, it may actually increase the arbitrageur's firm-specific risk). The arbitrageur must therefore expect a higher return to offset the risk that she cannot lay off, which in turn reduces arbitrage activity and lowers market efficiency. The result is much like Sanford Grossman and Joseph Stiglitz's now familiar point that informationally efficient markets are impossible because full efficiency eliminates the returns to the very activity that makes the market efficient, with the result of an "equilibrium degree of disequilibrium."

Noise trader risk similarly reduces arbitrage effectiveness because arbitrageurs bear the risk that noise traders will continue to be irrational, therefore maintaining, or even increasing, the mispricing. Because the arbitrageur will also have to be compensated for the risk that noise traders' continued confusion will adversely affect the value of her rational bets, the

required return goes up and level of activity goes down, resulting in a cost-driven level of market inefficiency.

Institutional limits on arbitrage prevent arbitrageurs from trading away information inefficiencies that result not from market risk, but from the structure of the institutions through which the arbitrageurs act. For our purposes, those limits fall into two categories: regulatory and market constraints on the mechanisms of arbitrage, and the structure of arbitrageurs' incentives. Each category operates to restrict the extent to which arbitrage can correct mispricing.

Regulatory restrictions on arbitrage are directed at short sales, undertaken by an arbitrageur when she believes the market price of a security is higher than its efficient price. In a short sale, the arbitrageur sells a security she does not own. To accomplish that, she must first find an existing owner of the overpriced security who is willing to lend the security to the arbitrageur. The borrowed stock is then sold, the arbitrageur betting that the price of the security will fall before the security must be purchased to repay the loan.

Securities Exchange Act rules 10a-1 and 10a-2 provide the basic regulatory framework. Rule 10a-1, the "uptick test," generally prohibits a short sale at a price below the security's last reported price, and Rule 10a-2 restricts activities by broker-dealers that could facilitate a violation of the uptick test. The idea behind those Depression-era prohibitions is to prevent "speculators" from driving down the price of a stock by selling it at prices below those that would prevail if all sales were "long" sales, i.e., sales made by traders who actually held the stock. The difficulty with the rule is simply the obverse of its asserted benefit. The clientele of long traders who hold a stock are self-selected optimists. Short selling, through its information-revealing properties, pushes stock prices to a lower, more efficient level; to the extent that the uptick rule actually succeeds in restricting arbitrage, the level of market efficiency suffers.

There is good news here. The Securities and Exchange Commission recently announced an ambitious experiment to test the effects of liberalizing the short-selling regulations. Over the coming months, the SEC will permit a large sample of stocks to be sold short, while its analysts gather data on how that practice affects the market without an uptick test.

Market restrictions on short selling involve limits on both the demand side and the supply side — the costs and availability of shares to borrow to affect a short sale. While Securities Exchange Act Section 316(c) restricts short selling by officers, directors, and large shareholders of publicly traded companies, the more serious demand constraint is voluntary; a recent SEC study reports that only some 43 percent of mutual funds are authorized by their charters to sell short. During the six-month period ending April 30, 2003, only about 2.5 percent of registered investment companies (236 out of some 9,000) actually engaged in short selling. Because 79 percent of mutual funds report that they do not use derivatives, it is unlikely that the charter restrictions are being avoided through the use of synthetic securities.

Market restrictions on the supply side relate to the lending market for the securities that must be borrowed in order to make a short sale. Preparation for a short sale begins with a request

that the arbitrageur's broker find a lender for the shares that are to be sold. The universe for potential lenders includes the broker itself (if it has an inventory of the desired stock) or institutional investors (including pension funds, insurance companies, and index funds, all of whom have long-term strategies that are unlikely to be negatively affected by liquidity constraints resulting from securities lending). The arbitrageur transfers collateral (typically cash) to the lender in the amount of 102 percent of the value of the borrowed securities. The lender then pays interest to the arbitrageur on the cash collateral, termed the rebate rate, and has the right to call in the loan at any time. If the loan is called at a time when the shares have risen in value, the arbitrageur will be forced to close her position at a loss unless another lender is found. Additionally, SEC Regulation T requires that the arbitrageur post a margin of 50 percent of the borrowed securities' value in additional collateral.

In general, the lending market available to short sellers for large-issuer securities is broad and deep. Large-cap stocks usually are easy and cheap to borrow, with the great majority requiring loan fees of less than one percent per year. In contrast, borrowing small-cap stocks with little institutional ownership may be difficult and expensive. As much as 16 percent of the stocks in the Center for Research in Security Prices file may be impossible to borrow.

Recent theoretical and empirical work suggests, however, that it becomes more costly to borrow a stock as the divergence of opinion about the security's value increases. The logic reflects the fact that those who do not lend the security forgo the price they would have received for its loan. Thus, those holding a stock must value it more highly than those who lend it by an amount in excess of the loan fee. The greater the divergence of opinion concerning the stock's value, the higher the loan fees, yielding the perverse result that the transaction costs of arbitrage increase in precisely the circumstance when the activity is most important.

Consistent with significant market limits on arbitrage, short interest in securities is generally quite small. A recent study reports that over the period 1976–1993, more than 80 percent of listed firms had short interests of less than 0.5 percent of outstanding shares, and more than 98 percent had short interests of less than 5 percent, a level consistent in magnitude with earlier assessments. And as one might expect from strong limits on arbitrage, the empirical evidence “is broadly consistent with the idea that short-sale constraints matter for equilibrium stock prices and expected returns.”

The problem, however, is with the magnitude of the costs. If all but small, non-institutional stock is readily available for borrowing, the regulatory and market-imposed transaction costs of short-selling seem too small to account for the limited amount of short-selling we observe and for its impact on pricing. A recent study by Joseph Chen, Harrison Hong, and Jeremy Stein of the impact of short selling constraints concluded that “an interesting question that our work raises, but does not answer, is this: why do short-sale constraints seem to be so strongly binding? Or said slightly differently: why, in spite of the high apparent risk-adjusted returns to strategies involving shorting, is there so little aggregate short interest in virtu-

ally all stocks? . . . We are skeptical that all, or even most, of the answer has to do with . . . specific transaction costs.”

The structure of arbitrageurs' incentives may provide the identity of the dark matter of the short sale universe — the source of constraints that the transaction costs of short selling do not explain. Recent work highlights a number of incentive problems, including a more realistic account of arbitrageurs' goals and the agency costs of arbitrage.

The first problem is that we have treated arbitrageurs as a kind of market maker whose role is to police the efficiency of prices and whose efforts will be compromised to the extent that regulatory and transaction costs make short-selling costly. In fact, however, arbitrageurs have a quite different goal: to make money. That, in turn, suggests that arbitrageurs act not only on a difference between a stock's market price and its fundamental value, but also on a difference between a stock's current market price and its future market price regardless of the relation between its future market price and its fundamental value.

If overly optimistic noise traders are in the market, shorting the stock is not the only way to make money. Instead, one can profit by anticipating the direction of the noise traders' valuation error, and taking advantage of that error through long, not short, positions with the goal of selling the shares to noise traders at a higher future price. The result may be to drive up the price of already overvalued stocks, and to prolong the length and increase the extent of bubbles.

The second problem is the agency costs of arbitrage. Keep in mind that arbitrage positions are made based on *ex ante* expectations, but the gain realized depends on *ex post* outcomes. The two may differ because of either the arbitrageurs' skill in identifying mispricing or because of fundamental or noise trader risk; that is, an investment may fail either because of bad judgment or because of bad luck.

For an arbitrageur trading for her own account, we can presume the explanation for a failed investment is observable. But now assume that the arbitrageur is instead an investment professional whose capital is raised from institutional investors and who receives a portion of the profits — the arbitrageur runs a hedge fund. Because the initial *ex ante* assessment of the portfolio investment is not observable to the fund investor, the investor then may use the investment's *ex post* outcome as a proxy of the arbitrageur's skill, with the effect of exposing the arbitrageur's human capital to both fundamental and noise trader risk because the fund investor may mistakenly treat a loss that really results from bad luck as evidence of bad judgment. Arbitrageurs thought to have “bad judgment” will have difficulty raising new funds. That, in turn, will cause the arbitrageur to reduce her risk by taking more conservative positions; that is, taking positions more like those of everyone else. Importantly, the personal risk to the arbitrageur increases as the importance of arbitrage as a means to correct market price increases. The greater the disagreement about a stock's price, the greater the bad luck risk that the arbitrage position turns out badly and, hence, the greater risk to the arbitrageur's human capital.

This interaction between noise trader risk and the agency costs of arbitrage can plausibly lead to bubble-like conditions. Once noise traders enter the market in large numbers, the risk

to arbitrage increases, which in turn results in an independent reduction in the level of arbitrage. This reduction, one might imagine, is more or less linear. More important, the presence of a market driven by noise traders has the potential to create a kink in the arbitrage supply curve, when the potential profits from momentum trading exceed the potential profits from short selling. From this perspective, one might consider a sharp increase in the number of individual investors in the market as a pre-bust signal of a bubble. That assessment turns on its head the familiar anecdotal observation that when individuals get into the market, the professionals get out: when individuals enter the market in large numbers, professionals find something to sell them.

A final potential limit on arbitrage looks back to the psychological biases that may underlie the noise trader phenomenon. To this point, we have treated arbitrageurs as if they still met the perfect rationality assumption of traditional theory — even if they are responding to the presence of noise traders or frictions in the incentive structure they face, they do so rationally. However, even professional traders are people. Maybe they are subject to cognitive biases as well; that is, the existence of irrational professional traders may be a limit to arbitrage.

The issue of whether some or all of the cognitive biases are hardwired or can be diminished by education or experience is a contested subject whose review is far beyond our ambition here. We note only that when the studies place individuals in a position where the goal is to make money, the cognitive biases seem to disappear quickly. And because the organization has the capacity to shape the traders' incentives so that the goal is clear, the potential for learning to occur and be reinforced is significant. Thus, we will treat professional traders as rational actors in responding to the incentives that they face.

A TENTATIVE ASSESSMENT OF THE BEHAVIORAL FINANCE PRINCIPLES

Despite the body of experimental evidence supporting persistent decision-making biases in some portion of the population, we are skeptical that this phenomenon will be found generally to play a significant role in setting aggregate price levels because the biases of different individuals to some important extent offset each other. Investor irrationality should be a matter of real concern when a single bias affects most noise traders, leaving a much heavier burden on arbitrage. And the problem will increase more than monotonically as the number of infected noise traders increases. As the volume of irrational trades increases, a point is reached at which arbitrageurs' most profitable strategy shifts from betting against the noise traders to buying in front of them, with the goal of exploiting the noise traders' mistake by selling overvalued stock to them.

A sharp increase in the participation of individual investors is a powerful indication that they share a common bias — the likelihood that a coincidence of different biases all leading to increasing participation at the same time seems small. Thus, a spike in individual trading — Lee, Shleifer, and Thaler's proxy for noise trading — may serve as a limited predictor of price bubbles. On those occasions, arbitrage constraints on price are relaxed and the effects of cognitive bias-

es on prices are likely to be of significantly greater magnitude than cost-based deviations from perfect market conditions.

Except for those situations in which the interactions of noise traders and market professionals create bubbles, the behavioral finance bias literature will have its greatest impact on circumstances when the concern is not with aggregate price effects, but with the behavior of individual investors. We may care a great deal if individuals systematically make poor investment decisions with respect to their retirement savings, especially with the growing shift from defined benefit to defined contribution pension plans, even if their mistakes do not affect price levels at all.

In contrast to our skepticism that cognitive biases affect market efficiency only episodically (when the number of individual investors spikes and their biases therefore likely coincide), we are quite sympathetic to concerns that agency and incentive problems constrain the professionally informed trading mechanism continuously, even in times of normal trading. MOME's relative efficiency concept, following Grossman and Stiglitz, builds on the idea that the cost of information limits the effectiveness of professionally informed trading — it has to pay to be informed. Agency and incentive problems between, for example, hedge funds and their investors and between hedge funds and their portfolio managers pose the same kind of tradeoff — it has to pay to reduce those costs.

MOME AND BEHAVIORAL FINANCE

How well does MOME's focus on the distribution and cost of information stand up to behavioral finance today? The answer, we believe, is mixed. The good news is that the central categories of MOME, including the market mechanisms and the concept of relative efficiency, are consistent not only with the established empirical findings of behavioral finance but with some of its more promising models as well. The bad news is that back in the early 1980s, we greatly underestimated the institutional obstacles to the production and rapid reflection of information in share prices.

THE GOOD NEWS The good news about MOME extends to both fact and theory. On the empirical side, proponents of both rational markets and behavioral finance agree that many of the long-term pricing anomalies that cut against the efficiency of market prices largely disappear when analysts control for company size. Disappearing anomalies include, for example, the underpricing of initial public offerings and seasoned equity offerings. The size-related character of those anomalies is good news because it is precisely what MOME would predict on the assumption that the size of the float is a critical determinant of the amount and quality of information about issuers and the relative efficiency with which that information is reflected in market prices. The reasoning is simple. Small issuers have a limited following among analysts and other professional investors, in part because there is little profit in researching issuers whose size restricts the potential gains. As a result, less information is produced, verification of information is more costly, and net returns available to investors and securities traders are correspondingly lower.

Size, analyst coverage, and the attendant availability can account for pricing anomalies of other sorts as well. On the the-

ory side, an important model developed by Harrison Hong and Jeremy Stein explains momentum trading and skewness in stock prices on the basis of the slow diffusion of private information through the economy. Traders without access to private information rationally treat price movements as a proxy for the injection of new information, which explains momentum trading as well as sudden reversals in price when traders discover they have already overshot share value. In support of this model, Hong, Stein, and Terence Lim present evidence that momentum trading in shares is particularly strong among small firms and firms that attract little interest among analysts.

THE BAD NEWS If recent models of the production and diffusion of information confirm the continuing relevance of MOME's analysis, our original account of market mechanisms and the institutional production of information suffered from what might be termed "naiveté bias." For example, we implic-

the institutions that employ them. Or, put differently, an investment bank's reputation among issuers is likely to matter more to it than its reputation among the lay investors who rely on its analysts' reports.

In sum, on every dimension of information costs — the costs of producing, verifying, and processing valuation data — we confess error, not about the roles of the institutions that supply information to the market, but about how well they perform their roles. The point is perhaps too obvious today to merit elaboration, but the market cannot be more efficient than the institutions that fix quality and cost of valuation information permit it to be. That, after all, was MOME's principal point.

POLICY IMPLICATIONS

We have argued that the binding constraints on market efficiency arise either from institutional limitations or the interaction of the arbitrage mechanism with cognitive biases — not from the

The market cannot be more efficient than the institutions that fix quality and cost of valuation permit it to be.

itly underestimated the institutional complexities that attend the production, processing, and verification of market information, as well as its reflection in share prices. Some aspects of our naiveté were discussed earlier in this essay: in particular, the legal and institutional limitations on arbitrage, including the agency problems that afflict institutional investors — such as the role of incentive structures in encouraging herding behavior by fund managers at the expense of fund investors.

But even more important than underestimating the limits on the arbitrage mechanism, we failed to appreciate the magnitude of the incentive problems in the core market institutions that produce, verify, and process information about corporate issuers. As the Enron cohort of financial scandals demonstrated, lucrative equity compensation has had the side effect of creating powerful incentives for managers to increase share prices. We suppose that managers usually respond by creating additional value for shareholders. But sometimes they respond by feeding distorted information to the market — or even by lying outright, as in such recent cases as WorldCom and HealthSouth. Recent scandals also demonstrate that we also were too sanguine about the role of the institutions that we termed "reputational intermediaries" — the established investment banks, commercial banks, accounting firms, and law firms that use their reputations to vouch for the representations of unknown issuers, and so reduce the information costs of investors. Finally, we were naïve about the role of security analysts, and particularly those employed by the investment banks on the "sell side" of the market. Those analysts, it appears, often acted as selling agents for the client-issuers of

widespread existence of cognitive biases alone. There are implications of this view for the formulation of regulatory policy.

BEHAVIOR FINANCE CAN GUIDE REFORM We see two principal areas where behavioral finance is likely to have policy implications in the near term. One lies on the institutional side. Given the importance of limitations on the arbitrage mechanism that we have emphasized thus far, regulators should clearly seek to reduce legal and institutional barriers to arbitrage. Thus, the SEC should continue to explore removing the uptick rule and margin requirements that burden short selling, as well as campaigning against the lingering taint that makes institutional investors such as mutual funds reluctant to pursue short-selling strategies. Far from being suspect, short selling actually confers a positive externality on the entire market by speeding the reflection of unfavorable information in share prices. In addition, behavioral finance may support temporary interventions in the market, such as trading halts, when market behavior suggests a surge of biased trading that threatens to destabilize arbitrage. We hesitate to make this prediction too forcefully, however, as there is still much work to be done in parsing out the psychological and institutional roots of market frenzies.

We are far more confident about a second area in which behavioral finance might eventually inform regulatory policy: the protection of individual investors. The possible consequences for policy involve paternalistic responses to cognitive bias. As we argued above, three conditions must be met for psychological distortions to affect share prices:

- cognitive biases must be pervasive (as most com-

- mentators believe they are);
- they must be correlated (because they would otherwise be offsetting); and,
- the arbitrage mechanism must fail with respect to their effects.

Notice, however, that cognitive bias can injure investors even if it has no effect whatsoever on share prices, i.e., the second and third conditions above are not met. Perhaps the best example is the employee who, as a result of limited knowledge or cognitive bias, misallocates investment in a 401(k) plan by failing to diversify her investments, or assumes a level of risk inappropriate to her age and retirement aspirations.

The rise of defined contribution and voluntary investment plans has shifted discretion over retirement savings from professional traders to individual “lay” investors who often are noise traders. It might well be, then, that we would be wise to limit the investment discretion of employee-investors, precisely in order to prevent them from harming themselves. Such limitations might be mandatory for government-sponsored or tax-favored retirement plans: for example, an inflexible diversification requirement. Alternatively, the limitations might take the form of what has been termed “asymmetric paternalism,” i.e., default rules that sophisticated investors can avoid but that are binding on unsophisticated investors who are more likely to make costly errors as a result of cognitive bias or bounded rationality.

THE TAKEOVER DEBATE Once we leave the easy cases of short-selling restrictions, obvious market frenzies, and undiversified retirement savings, the legal implications of behavioral finance for corporate and securities law become much murkier for the simple reason that we know little about both the extent and nature of cognitive bias among traders or the interaction of cognitive bias with the institutions that generate information and the mechanisms that reflect it in price (including, above all, the arbitrage activity of sophisticated investors). We therefore find ourselves largely in agreement with Donald Langevoort’s assessment of the implications of behavioral finance for securities regulation, which, no doubt over-simplifying, we would summarize as, “not much so far, although lawmakers should stay tuned to current research and keep an open mind.” Indeed, we would go one step further to caution against the use of behavioral finance to advance policy agendas that it cannot possibly support. We close this essay with the cautionary example of a policy debate in which behavioral finance is sometimes said to have important implications when in fact it does not.

The example we have in mind is the claim that is sometimes made in debates over takeovers that investor irrationality demonstrates the wisdom of vesting discretion over the decision to defend against hostile takeovers in the hands of managers rather than shareholders.

We find this claim unpersuasive for several reasons that nicely illustrate the limits of cognitive psychology in setting basic corporate policy. In the first place, market efficiency has a limited role in the takeover debate. The primary policy trade-off is between the absence of strong-form efficiency (i.e., the possibility that managers have information about the corporation’s value that the market lacks, which is the reason for giv-

ing management discretion to defend) and the possibility of managerial agency cost (i.e., the reason for giving the decision to shareholders). This one comes out in favor of shareholder decision-making because target management can always ameliorate the failure of strong-form efficiency by disclosing its information if takeover decision-making is allocated to shareholders, while allocating authority to management does nothing to ameliorate the agency cost problem.

It is at this point that the cognitive bias component of behavioral finance comes into play: The balance may shift if, despite disclosure, shareholders will predictably reject target managers’ advice because of one or another cognitive bias. Of course, given the range of cognitive biases, one cannot entirely reject this possibility. Some biases predict that shareholders will tender too readily while others predict an unwarranted reluctance to tender. In the context of the allocation of takeover decision-making between managers and shareholders, the critical point is that cognitive bias analysis be applied on a bilateral or comparative basis.

This concern grows out of the fact that the experimental literature is largely unilateral in its focus. The experiments are concerned only with whether a particular decision-maker is subject to a cognitive bias, not whether one competing decision-maker is more impaired than another. But when cognitive bias is invoked to allocate authority among competing decision-makers, the analysis must be bilateral: the potential biases of the decision-makers must be compared. The question is whether managers or shareholders’ decisions are likely to be more distorted.

The comparison seems to us to favor allocating decision-making authority to shareholders. First, it is simply unclear which, if any, biases are likely to apply to individual shareholders when they must choose whether to accept a hostile offer. Moreover, the outcome of the takeover is likely to be determined by the decisions of institutional investors, who are less likely to be subject to cognitive biases (but may be subject to institutional influences) — the shareholders critical to the outcome of a hostile takeover look little like the noise trader clientele of closed-end mutual funds. Finally, the market for corporate control operates to an extent as a backstop in case cognitive biases cause target shareholders to tender into too low an offer. The ubiquity of competing bidders emerging in response to an underpriced offer can save the shareholders from their biases.

On the other side, one can imagine a range of biases that may influence target managers to resist a hostile takeover even when the transaction is in the shareholders’ best interests. A reaction to cognitive dissonance may cause managers to respond to an offer that calls into question their performance and competence by deriding the bidder’s motives and promising a brighter future if only the shareholders have patience. Managers may genuinely believe their claims, but behavioral finance suggests that their assessment may be driven by a cognitive bias. This effort at dissonance reduction may, in turn, be exacerbated by the overconfidence bias — managers’ vigorous defense may be encouraged by a biased assessment of their own skills. Other examples are possible, but the point by now should be clear: When cognitive bias analysis is invoked to illuminate the choice between two decision-makers, its application must be bilateral.

We conclude that the cognitive bias element of behavioral finance is unlikely to change the tradeoff between agency costs and strong-form market inefficiency that we believe supports allocating to shareholders the choice of whether a hostile takeover goes forward. To be sure, by highlighting the possibility of good faith but systematically misguided defensive action, the cognitive bias analysis does serve to give richness to the explanation for target managers' behavior that agency theory's simple self-interest paradigm lacks. But this useful insight reinforces, rather than undercuts, an allocation of decision-making authority to shareholders.

CONCLUSION

So where does our retrospective leave us? Twenty years further, we think, along the road leading from elegant models of the workings of the capital market in a frictionless world, to an understanding of how the market operates in a world where information is costly and unevenly distributed, agents are self-interested, transactions costs are pervasive, and noise traders are common. The nature of this more realistic understanding is beginning to take shape, and it can be described in a single word: messy. There are a lot more moving parts and, as a result, a much larger number of interactions to understand. Models will be necessarily partial, illuminating particular interactions, but will fall far short (and without the ambition) of a unified field theory.

That said, we come away with some confidence in a number of themes, some that were explicit in MOME, some that we missed, and others that reflect an assessment of the likely contribution of cognitive psychology to our understanding of how the capital market functions.

First, as was explicit in MOME, we believe that understanding the structure of institutions is central to understanding the operation of the capital market. MOME's shortcoming was the failure to drill deeply enough into the incentive and agency structure of important market institutions like those through which arbitrage is carried out. To the large extent that behavioral finance is composed of applying agency, information, and incentive analysis to capital market institutions, it promises to deepen our understanding of how the capital market operates in the real world.

Second, we are skeptical that the new focus on cognitive biases in the end will explain very much about price formation except in circumstances in which investor biases both coincide and give rise to increased participation. Thus, we expect that this component of behavioral finance will have a limited role in the market efficiency debate. In contrast, the literature can be quite important under circumstances in which we care about the consequences of biased decision-making on the decision-makers themselves, independent of whether aggregate price levels are affected. Reform efforts directed at individuals' decisions with respect to pension investments, as with 401(k), provide a good example.

Our final theme is one of balance. When cognitive psychology is used to analyze issues relating to the allocation of decision-making between competing parties, the application must be bilateral and comparative. It is insufficient merely to

demonstrate that one party may exhibit cognitive biases. Identifying a bias in one party begins the analysis; it is completed only when that impairment is compared to those of the other party. As we suggested in our analysis of the application of cognitive bias analysis to tender offers, the fact that shareholders may have a bias in deciding whether to tender does not demonstrate that managers should have the power to block an offer. Rather, the shareholders' bias must be compared with those biases that affect management.

Twenty years after publication, we remain comfortable with the analytic framework that animates MOME. We should have been more skeptical of market institutions then, but skepticism grows with age. R

READINGS

- "Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies," by Harrison Hong, Terence Lim, and Jeremy C. Stein. *Journal of Finance*, Vol. 55 (2000).
- "Breadth of Ownership and Stock Returns," by Joseph Chen, Harrison Hong, and Jeremy Stein. *Journal of Financial Economics*, Vol. 66 (2002).
- "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," by William Sharpe. *Journal of Finance*, Vol. 19 (1964).
- "The Cost of Capital, Corporate Finance, and the Theory of Investment," by Franco Modigliani and Merton H. Miller. *American Economic Review*, Vol. 48 (1958).
- "Dividend Policy, Growth, and the Valuation of Shares," by Merton H. Miller and Franco Modigliani. *Journal of Business*, Vol. 34 (1961).
- "Efficient Capital Markets: A Review of Theory and Empirical Work," by Eugene Fama. *Journal of Finance*, Vol. 25 (1970).
- "Investment Sentiment and the Closed-End Fund Puzzle," by Charles Lee, Andrei Shleifer, and Richard Thaler. *Journal of Finance*, Vol. 46 (1991).
- "The Mechanisms of Market Efficiency," by Ronald J. Gilson and Reinier Kraakman. *Virginia Law Review*, Vol. 70 (1984).
- "On the Impossibility of Informationally Efficient Markets," by Sanford Grossman and Joseph Stiglitz. *American Economic Review*, Vol. 70 (1980).
- "Regulation for Conservatives: Behavioral Economics and the Case for Asymmetric Paternalism," by Colin Camerer, Samuel Issacharoff, George Loewenstein, Ted O'Donoghue, and Matthew Rabin. *University of Pennsylvania Law Review*, Vol. 151 (2003).
- "Some Anomalous Evidence Regarding Market Efficiency," by Michael Jensen. *Journal of Financial Economics*, Vol. 95 (1978).
- "Taming the Animal Spirits of the Stock Market: A Behavioral Approach to Securities Regulation," by Donald C. Langevoort. *Northwestern University Law Review*, Vol. 97 (2002).
- "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," by Harrison Hong and Jeremy C. Stein. *Journal of Finance*, Vol. 54 (1999).