

EFFICIENT MARKETS AND EMPIRICAL FINANCE

. . .

John H. Cochrane and Tobias J. Moskowitz

In 1970, Gene Fama defined a market to be “informationally efficient” if prices at each moment incorporate available information about future values.

A market in which prices always “fully reflect” available information is called “efficient.”—Fama (1970)

Informational efficiency is a natural consequence of competition, relatively free entry, and low information costs. If there is a signal that an asset’s future value will be higher than its current value, competitive traders will try to buy the asset today. They bid prices up, until prices reflect the new information. We illustrate the process in the picture below.

As an immediate consequence, in a competitive asset market, price *changes* should not be predictable. This is a powerful result, which finance practitioners—whose job is to predict price changes—have great trouble digesting even as a matter of theory.

Like all good theories, it only seems simple in retrospect. It was a much harder and deeper intellectual achievement than our story suggests. It took

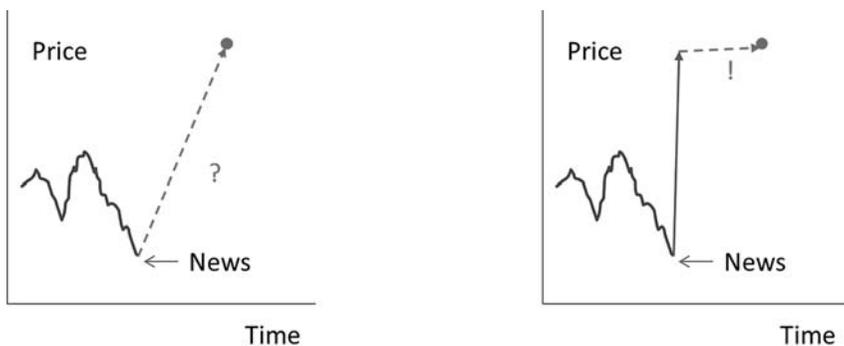


Figure 1

nearly a century to figure out the *basic predictions* of market efficiency, from Bachelier's random walk, Samuelson's proof that price changes should not be predictable, through Fama's efficient market essay, to the stochastic discount factor expression—price equals expected discounted payoff—that constitutes today's distillation of the theory.

EFFICIENCY AND EMPIRICAL FINANCE

The famous 1970 essay with which we start this volume is usually considered important as a statement of theory. The textbook treatment usually runs through Gene's catalog of "strong" to "weak" form predictions. But that is, in our view, not its key contribution. Rather, Gene showed how the simple idea of informational efficiency can organize and give purpose to a vast empirical project. Gene's contributions are more like Darwin's than Einstein's. Evolution by natural selection is a simple-sounding principle, with a lot of hard thinking needed to make it useful, that organized a vast empirical project in biology. Without evolution, natural history would have just been a collection of curious facts about plants and animals. Efficient markets is a simple-sounding principle, with a lot of hard thinking needed to make it useful, that organized a vast empirical project in financial economics. Without the efficient markets hypothesis, empirical finance would have just been a collection of Wall-Street anecdotes, how-I-got-rich stories, and technical-trading newsheets.

That empirical work consists, fundamentally, of applying scientific method to financial markets. Modern medicine doesn't ask old people for their health secrets. It does double-blind clinical trials. To this, we owe our ability to cure and prevent many diseases. Modern empirical finance doesn't ask Warren Buffet to share his pearls of investment wisdom, as the media does. We study a survivor-bias-free sample of funds sorted on some *ex ante* visible characteristic to separate skill from luck, and we correct for exposure to systematic risk. To this we owe our wisdom, and maybe, as a society, a lot of wealth as well.

Empirical work is not easy, either. The efficient markets hypothesis doesn't just hand you "predictions." The empirical implications of the efficient markets hypothesis are subtle and often deeply counterintuitive. Empirical work often tackles tough anomalies, each of which looks superficially like a glaring violation of efficiency, and each endorsed by a cheering crowd of rich (or perhaps lucky?) traders and hoping-to-be-rich current students. Like any nonexperimental science, distinguishing cause from effect is tough. There are always at least two stories for every correlation. It took hard work and great insight to

account for risk premiums, selection biases, reverse causality, and endogenous variables, and to develop the associated statistical procedures. It took genius to sort through the mountains of charts and graphs that computers can spit out, to see the basic clear picture and a framework to organize it.

For example, efficiency implies that trading rules—“buy when the market went up yesterday”—should not work. The surprising result is that, when examined scientifically—and especially when accounting for the selection bias that only successful traders leave a paper trail—trading rules, technical systems, market newsletters, and so on have essentially no power beyond that of luck to forecast stock prices. That was true in 1970 and remains true today. This statement is not a theorem, an axiom, a philosophy, a tautology, or a religion: it is an empirical prediction that could easily have come out the other way.

If markets are efficient, then professional managers should do no better than monkeys with darts. This prediction too bore out remarkably well in the data in 1970 and still does to a large extent. It too could have come out the other way.

As Ray Ball points out, stock prices move on all sorts of information, which we as students of the market do not observe. How can you tell that prices are reacting correctly? It took genius to invert the question and to look instead at the average response to known pieces of information. And there, to see almost perfect stair-step reactions to information as predicted by the efficient markets view.

Already in 1970, Gene found that prices are not always efficient. For example, prices rise on the release of inside information, so that information, though known by someone, was not reflected in the original price. Perfect efficiency, like perfect competition, is always an ideal, toward which markets may be closer or further away.

This is *great* news. Only a theory that can be proved wrong has any content at all. Theories that can “explain” anything are as useless as “prices went down because the Gods are angry.” Still, Gene found in 1970 that asset returns displayed patterns surprisingly close to efficient markets predictions, and surprisingly far from the assertions of most practitioners.

RISK, RETURN, AND MARKET EQUILIBRIUM

In a competitive market, investments that carry more risk must offer better returns on average. Therefore, if one finds a high expected return (i.e., an “underpriced” security), that fact does not necessarily imply inefficiency. It

might represent a risk premium. The simple idea that price changes should not be predictable obviously needs amendment for allowable variation in risk premiums.

The joint hypothesis theorem in Gene's 1970 article is at least as important as the definition and characterization of efficiency. How much price or expected return variation can risk premiums account for?

The theory only has empirical content, however, within the context of a more specific model of market equilibrium . . . (Fama 1970)

And, without specifying a model of market equilibrium, the theory (or its antitheses) therefore has *no* empirical content.

In more modern language, absent arbitrage opportunities we can write price as the expected discounted payoff,

$$\Lambda_t p_t = E_t [\Lambda_{t+1} (p_{t+1} + d_{t+1})] = \sum_s \pi(s) \Lambda_{t+1}(s) (p_{t+1}(s) + d_{t+1}(s))$$

where Λ_t is the stochastic discount factor, for example $\Lambda_t = \beta^t u'(c_t)$ in the canonical consumption based model, d_t are dividends, and s denotes states of nature. Really, *discounted* price changes, after correction for dividends, are unpredictable.

Informational efficiency says that the probabilities $\pi(s)$ forming the expectation E_t incorporate all available information. Risk aversion, risk premiums, and a "model of market equilibrium" are encoded in the discount factor Λ_t .

But, as the equation makes perfectly clear, $\pi(s)$ and $\Lambda_{t+1}(s)$ always enter together. You can't say anything about probabilities without saying something about discount factors, and vice versa. And, as seen most clearly in the utility formulation, the discount factor depends on the general equilibrium allocation of consumption c , as well as preferences and risk aversion encoded in u .

This theorem is a vital guiding light for empirical work. To this day, paper after paper looks at patterns in prices and declares them to be "irrational" or a "bubble" or otherwise "inefficient." But Fama proved in 1970 that such statements are simply empty, devoid of "empirical content." There is always some "model of market equilibrium" that will "explain" any price. At a minimum, one must state what kinds of models one finds plausible and verify that they do not generate risk premiums that, quantitatively, can explain a proposed anomaly.

In our view, the joint-hypothesis problem associated with efficient markets is as important as the theory itself. It has been key in organizing an ever-

expanding set of pricing facts, and it is the more important organizing principle for today's active empirical work.

EFFICIENT MARKETS II

So, nearly a half-century later, how is the efficient markets hypothesis—and its organizing principle for empirical work—doing?

We included Gene's "Efficient Capital Markets II" in this volume as his progress report. It was written halfway through the second revolution in empirical finance—the understanding that risk premiums ("rational" or not) vary through time and across assets in large and unexpected ways. Also, "Efficient Capital Markets II" is one of those papers that, though less influential as measured by citations and influence on the profession, is perhaps a better summary and distillation for readers new to the field than was the original.

Between 1970 and 1990, evidence that returns are predictable mounted, and we cover several key Fama, and Fama and French, papers in that discovery later. As Gene emphasizes here, the evidence lies right where macroeconomic risk premiums lurk. Time-varying expected returns are most visible at long horizons, correlated across markets, and strongly correlated with business conditions where one might expect risk premiums to affect prices.

Event studies are perhaps the cleanest grounds to watch the incorporation of information into prices without joint-hypothesis problems, and here efficiency has stood up well. The surprising inability of professional managers to make money has continued in the vast literature Gene covers here.

BEHAVIORAL FINANCE

In 1970, there really wasn't a coherent alternative to efficient markets. "It takes a model to beat a model," Gene reminds us frequently. "People are dumb," while perhaps a trenchant observation on the human condition, or perhaps just an illusion of superiority, is not a theory. Or it's too good a theory—it can "explain" anything.

Starting in the late 1970s, a large literature that came to be known as "behavioral finance" worked to construct a coherent alternative model in which prices don't "efficiently" reflect information—people's probability assessments remain biased even though they may receive lots of information. These authors appealed to social psychology and the findings about how people misperceive probabilities in systematic ways in laboratory settings. By doing so, they hoped to avoid the charge of residual-naming. These insights from psychology were

then applied to financial decision making and markets, providing both an alternative view of markets as well as some unique insights.

Gene always carefully read the empirical literature, even “behavioral,” and here took on the evidence and its behavioral interpretation. A possible summary of his judgment is that the literature really didn’t escape residual-naming and is not yet constrained by its theory. He finds that “under-reaction” is as frequent as “over-reaction.”

Behavioral finance has, in our minds, had an important beneficial effect on the finance profession. In the 1970s, journals demanded that one have a “hypothesis” to “test.” Facts for facts’ sake didn’t get published. This attitude was a valiant defense against fishing, but in the end was unproductive in that effort and stifled the production of facts that challenge theory. By adducing a behavioral “theory” that an empirical investigation was testing, the range of facts that empirical finance got to chew on grew dramatically. Yes, quite often one could transplant a “risk premium” introduction and a “behavioral” introduction on the same set of tables—the value effect comes to mind here—but the result is we all got to digest a vastly expanded set of facts. Hard sciences develop theories after uncomfortable facts, too. In the end, finding the truth must come from both sides—theory that the data hopes to verify, and facts in the data that theory needs to explain. Finance stands out in economics for the deep interplay between theory and often uncomfortable facts, a point Gene has made often.

STILL CONTENTIOUS

Forty-four years later, “efficiency” remains contentious.

Some of that contention reflects a simple misunderstanding of what social scientists do—and what they should do. What *about* Warren Buffet? What about Joe here, who predicted the market crash in his blog? Well, “data” is not the plural of “anecdote.” These are no more useful questions to social science than “How did Grandpa get to be so old, even though he smokes?” is to medicine. Empirical finance looks at *all* the managers and all their predictions, tries to separate luck from ex ante measures of skill, and collects clean data.

Another part of that contention reflects simple ignorance of the definition of informational “efficiency.” Every field of scholarly inquiry develops a technical terminology, often appropriating common words and giving them meanings quite different from the colloquial ones. “Efficient” estimators in statistics and “Pareto-Efficient” allocations in economics also have precise definitions, little related to the colloquial meaning of “efficiency.” But people who don’t know those definitions can say and write nonsense about the academic work.

“Efficiency” in finance means information, and only information. An *informationally* efficient market can suffer *economically inefficient* runs and crashes—so long as those crashes are not predictable. An informationally efficient market can have very badly regulated banks. An informationally efficient market need not process orders quickly or “efficiently.”

People who say “the huge size of the 2008 (or 1987, or 1929) crash proves markets are inefficient” or “the fact that finance experts didn’t foresee the crash proves that markets are inefficient” simply don’t know what the word “efficiency” means. The main prediction of efficient markets is exactly that price movements should be unpredictable! And that returns must be balanced by risk. Steady profits without risk *would* be a clear rejection of efficient markets. The efficient markets hypothesis says nothing about how volatile or non-normal returns can be.

Cochrane once told a reporter that he thought markets were pretty “efficient.” The reporter misquoted him as saying that markets are “self-regulating.” Sadly, even famous academics say silly things like this all too frequently.

Part of the contention over efficiency reflects a misunderstanding of what financial economics, as a social science, tries to do. We’re not here to “explain” the latest market gyration or provide market commentary. “The will of the gods” provided a perfectly good ex post story. The central idea of efficient markets, and markets as aggregators of information in the Hayekian tradition, is exactly that nobody can explain ex post why most prices move, nor predict movements ex ante.

Efficient markets taught us to evaluate theories by their rejectable predictions and by the numbers, and to do real, scientific, empirical work, not to read newspapers and tell stories. When someone asks, “Yes, but why *did* Warren Buffet get so rich?” the real answer is just “sir, that’s a poorly posed question.” Like “Why did uncle Joe not get cancer even though he smoked his whole life?” Medicine is useful, though it does not answer such questions. We recognize that this is unsatisfactory to journalists. Sadly, many academics still offer ex post stories for market movements on morning TV shows. Gene isn’t one of them.

When Gene labeled the swift incorporation of information into prices as “efficiency,” it was, perhaps, a bit of marketing genius. There is a fascinating story here, worth study by historians and philosophers of science and its rhetoric. What would have happened had Gene used another word? What if he had called it the “reflective” markets hypothesis, stating that prices “reflect” information? Would we still be arguing at all?

WHERE ARE WE TODAY?

With 40 years' hindsight, *are* markets efficient? Not always, and Gene said so in 1970. More recently, in our view, we have seen evidence that short-sales constraints and other frictions can lead to prices that are clearly informationally inefficient. The empirical question has always been to what degree a given phenomenon approaches an unattainable ideal.

Still, the answer today is much closer to “yes” than to “no” in the vast majority of serious empirical investigations. It certainly is a lot closer to “yes” than anyone expected in the 1960s or than the vast majority of practitioners believe today. There are strange fish in the water, but most are surprisingly small fry. And having conquered 157 anomalies with patient hard work, many of us can be excused for suspecting that just a little more work will make sense of the 158th.

Admittedly, how one measures “close” is contentious. A small, persistent expected return anomaly can add up to a large difference in price. As we survey below, small prediction R^2 can add up to large 1–10 portfolio mean spreads. A tenth of a percent of arbitrage is worth billions if you can lever it up enough.

Today, empirical finance is less devoted to debating whether markets are efficient and more to assessing how efficient markets are. The vast majority of academic work agrees that market prices quickly reflect most visible pieces of information. Current efficiency debates include whether there is some slow diffusion of information on top of a large immediate reaction, or extended overreaction to an initial price reaction, or even to what extent prices move when there is no information and why. The current broad consensus that fully efficient markets are a theoretical ideal, real markets are always somewhat inefficient, the goal of research is to quantify how much, and the answer is “a lot better than you thought,” pretty much sums up Gene’s 1970 essay. In that sense, efficient markets won.

Moreover, most debates apparently about “efficiency” today are really about the nature of the “model of market equilibrium.” Most current research really explores the amazing variety and subtle economics of risk premiums—focusing on the “joint hypothesis” rather than the “informational efficiency” part of Gene’s 1970 essay. The discovery that risk premiums vary over time and across assets far more than anyone expected in 1970 counts as the major finding of the second revolution in empirical finance. Since the financial crisis of 2008, the major methodological debate has been between models that see macroeconomic sources of these risk premiums versus models that see

risk premiums driven by “institutional finance” agency frictions. In this latter view, prices are “wrong” because markets are temporarily segmented and only leveraged intermediaries are “marginal.” They require higher risk premiums than a fully integrated market would produce. Even much behavioral finance is not about mishandling of *information*, it’s about misperceptions of *risk*. In all this work, slow diffusion of information into prices is just not an interesting mechanism for understanding the puzzles before us. But “models of market equilibrium” are much more important than we thought. The debate largely shifted from “do prices incorporate information efficiently?” to “what’s the right pricing model?”

This is also great news. Healthy fields settle debates with evidence and move on to new discoveries. But don’t conclude that efficient markets are passé. As evolution lies quietly behind the explosion in modern genetics, markets that are broadly efficient, in which prices quickly reflect information, quietly underlie all the interesting things we do today. This is the best fate any theory can aspire to.

Efficient markets are also important to the world at large, in ways that we can only begin to touch on here. The assurance that market prices are in some sense basically “right” lies behind many of the enormous changes we have seen in the financial and related worlds, from index funds, which have allowed for wide sharing of the risks and rewards of the stock market, to mark-to-market accounting, quantitative portfolio evaluation and benchmarking, modern risk management, and law and regulation.

The efficient markets hypothesis also is a close cousin, or perhaps uncle, to rational expectations and more deeply to the revolution that brought people and time into macroeconomics. This revolution became a hallmark of the University of Chicago economics department later in the 1970s.

Empirical financial economics, and the part that thinks about how information is incorporated into prices, is a live field, asking all sorts of interesting and important questions, with great discoveries left to be made. Is the finance industry too large or too small? Why do people continue to pay active managers so much? What accounts for the monstrous amount of trading? How is it, exactly, that information becomes reflected in prices through the trading process? Do millisecond traders help or hurt? How prevalent are runs? How many of the hundreds of return forecasting variables in the current literature are real? How many of the dozens of current risk factors do we really need? The ideas, facts, and empirical methods of informational efficiency continue to guide these important investigations.

The empirical focus of finance, inaugurated by Gene Fama in the 1970 essay, has been one of its greatest strengths, and largest influence. Economics has followed finance and become much more empirical in the last few decades.

Gene's bottom line is always: Look at the facts. Collect the data. Test the theory. Every time we look, the world surprises us totally. And it will again.