

Finance: Function Matters, Not Size

John H. Cochrane

The US economy spends \$170 billion a year on advertising, just to trick people into buying stuff they don't need. What a waste!

There are 2.2 people doing medical billing for every doctor that actually sees patients, costing \$360 billion—2.4 percent of GDP. Talk about an industry that is too big!

Wholesale and retail trade and transportation cost 14.6 percent of GDP, while all manufacturing is only 11.5 percent of GDP. We spend more to move goods around than to make them!

My wife asked me to look at light fixtures. Do you know how many thousands of different kinds of light fixtures there are? The excess complexity is insane. Ten ought to be plenty.

It's ridiculous how much people overpay for brand names when the generic is so much cheaper. People are pretty naive.

Business school finance professors are horribly overpaid. Ask an anthropologist! We get paid almost a half a million bucks, and work a grand total of 10 weeks a year, all to teach students that they can't make money trading in the stock market.

It's fun to pass judgment on waste, size, usefulness, complexity, naiveté, and excessive compensation, isn't it? But as economists, we have an analytical structure for thinking about these questions. We start with supply, demand, and competition,

■ *John H. Cochrane is the AQR Capital Management Distinguished Service Professor of Finance, University of Chicago Booth School of Business, Chicago, Illinois. He is also a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts; Senior Fellow, Hoover Institution, Stanford, California; and Adjunct Scholar, Cato Institute, Washington, DC. His email address is john.cochrane@chicagobooth.edu, and his website is <http://faculty.chicagobooth.edu/john.cochrane>.*

and with the suggestion of the first welfare theorem that these forces usually lead to socially beneficial arrangements. When outcomes seem puzzling using this analysis, we embark on a three-pronged investigation. First, we work harder to find how supply and demand might really operate, in the humble knowledge that initially puzzling institutions and outcomes have often taken us years to comprehend. Second, maybe there is a “market failure”—an externality, public good, natural monopoly, asymmetric information situation, or missing market—that explains our puzzle. Third, we often discover a “government failure,” that the puzzling aspect of our world is a consequence of laws or regulation, either unintended or the result of capture.

Only then can we begin to diagnose a divergence between reality and socially desirable outcomes, and only then can we start to think of how to improve reality. “I don’t understand it” doesn’t mean “it’s bad,” or “regulation will improve it.” And since that attitude pervades policy analysis in general and financial regulation in particular, economists do the world a disservice if we echo it.

I belabor this point, because I do not offer a competing black box. I don’t claim to estimate the socially optimal “size of finance” at, say, 8.267 percent of GDP. It’s just the wrong question. Hayek and the failure of planning should teach us a little modesty: Pronouncing on socially optimal industry size is a waste of time. Is the finance industry functioning well? Are there identifiable market or government distortions? Will proposed regulations help or make matters worse? These are useful questions.

With a rather catastrophic failure behind us and other crises bubbling on the back burner, it also seems a bit strange to be arguing whether 5 or 8 percent of GDP is the right “size” of finance, and whether it needs to be nudged to become larger or smaller. Many of us might happily accept an additional 3 percentage points of GDP in the financial sector in return for a financial system that is not prone to runs and crises. Our political system *has* accepted a big increase in resources devoted to financial regulation and compliance, and a potentially larger reduction in the efficiency, innovation, and competitiveness of financial institutions and markets, in the quest—misguided or not—for stability. The run-prone nature of the US financial system, together with its massive regulation, subsidies, government guarantees, and regulatory capture, looks to be a more fertile fishing ground for trying to understand market and government failures than does mere size.

Still, the size of finance represents a contentious issue, and my plea that we ask different questions isn’t going to silence the debate, so let us think about it. Let us use size as an organizing principle for studying function and dysfunction.

Greenwood and Scharfstein nicely review the key facts and ideas in their paper in this issue. Their most basic story is: quantity increased a lot, but prices didn’t fall. This description suggests a simple economic interpretation: The demand for financial services shifted out. People with scarce skills supplying such services made a lot of money. A system with proportional fees, which is a common structure in professional services, interacted with stock-price and home-price increases (a different surge in demand) to produce increased financial sector revenue. *Why*

demand shifted out, and why house and stock prices rose (temporarily, it turns out) are good questions—but they don't have much to do with the structure of the finance industry. This story also suggests that, like the weather, if you don't like the size of finance, just wait a while. Finance has contracted rather dramatically since 2007.

Many puzzles remain, however, and the current academic literature paints an interesting and quite novel picture of how the finance industry functions—and maybe does not function.

The Controversy over Active Management Fees

Management fees are a big part of the “size of finance.” Fees aren't GDP, of course, but they are much more easily measured. The large overall rise in fee revenue reflects several offsetting trends. Individuals moved investments from direct holdings to mutual funds, and then to index funds or other passive funds. This trend continues. New investors in defined contribution plans invest almost exclusively in mutual funds or exchange-traded funds.

Mutual fund fee *rates* came down sharply, in part reflecting the slow shift to very low-fee index and semi-passive funds, and in part reflecting competitive pressure. French (2008) reports that the average actively managed equity mutual fund fee fell from 2.19 percent in 1980 to 1 percent in 2007. Greenwood and Scharfstein (2012) report that average bond fund fees fell from 2.04 to 0.75 percent. Some index funds charge as little as 0.07 percent. Fee-based advisers and wealth managers are lowering fees, and bundling larger arrays of services, including tax and estate planning.

Funds are far more efficient vehicles for individual investors than holding individual stocks. The measured GDP of the fund industry is at least in part a benefit rather than a cost, as it displaces inefficient and unmeasured home production of financial management services. Hiring a (legal) house cleaner also raises measured GDP.

Thus, mutual fund fee revenues reflect declining rates multiplied by a much larger share of assets under management. This market does reflect sensible forces, if one is willing to grant a rather long time span for those forces to affect industry structure. But after all, the moves to low-cost airlines and big-box retailers took a while too.

However, at the same time that individuals were moving to passive funds and those funds were expanding, high-wealth individuals and institutions (pensions, endowments, sovereign-wealth funds, and so forth) moved their investments to hedge funds, private equity, venture capital, and other even higher-fee and more-active investment vehicles. Hedge fund fee rates are reportedly stable over time, and surprisingly large: managers charge 1.5–2.5 percent of assets each year, and also 15–25 percent of profits. This part of the market offers the more puzzling behavior.

The Traditional View

High-fee active management and underlying active trading have been deplored by academic finance for a generation. French (2008) offers a comprehensive summary. French estimates that equity investors in aggregate, between 1980 and 2006, paid 0.67 percent per year in active management fees, whose present value he estimated to equal 10 percent of their investments. French eloquently conveys the view that these investors wasted their money.

The standard analysis divides investment returns into “alpha” and “beta.” We run a regression of a fund’s returns on the returns of a low-cost index, both returns in excess of the risk-free rate. Beta is the slope coefficient. Beta times the index return is the component of the fund’s return that is earned for passively shouldering systematic risk, and can be synthesized by the investor without paying fees. Alpha is the intercept in this regression, and gives the mean of that part of the fund’s return that cannot be easily replicated. Alpha is conventionally interpreted as the extra return that the fund earns, on average, from the manager’s talent or superior information, and therefore potentially worth paying a fee to obtain. Both alpha and beta are, conceptually, one’s best estimate of this decomposition of returns going forward, of course. Estimated alphas from past history contain a great deal of luck.

The average alpha of all equity mutual funds, before fees, is very nearly zero. This result follows almost by accounting, since the portfolio of equity mutual funds, taken as a whole, is almost exactly the value-weighted market portfolio.

The evidence on hedge fund, private equity, venture capital, and other returns is complicated by survivor bias (funds that perform badly tend to drop out of the data) and by difficulties of calculating benchmarks that appropriately reflect the risks, time horizons, and illiquidity of these investments. But the academic argument over whether such funds as a class provide substantial alpha ends up arguing over a few percentage points one way or the other—hardly the promised gold mines.

Mediocre average results for actively managed investments might not be surprising. Entry into the business is relatively free. The average artist isn’t that good, either.

But one might expect that, as in every other field of human endeavor, the good managers would be reliably good. Michael Jordan’s past performance was a good forecast of what would happen in the next game. Yet the nearly universal conclusion of the academic literature is that there are no reliably “good” managers.

To evaluate this question, we must separate skill from luck. “Why did Warren Buffet earn so much money?” is not a productive question. The classic technique is to examine rules by which one might have chosen funds in the past, and then study the subsequent returns of all such funds. Study after study finds no reliable rule that one can use to identify funds that will perform well in the future, after controlling for betas. (Carhart 1997 is an excellent example.) Fama and French (2010) pursue a clever measurement that does not require one to hypothesize such a rule. They show that the distribution of estimated alpha across mutual funds is only very slightly wider than what one would expect if sample alphas were just due

to luck. Fama and French estimate (p. 1935) that the distribution of true alpha has a standard deviation of only 1.25 percent on an annual basis, meaning that only about one-sixth of funds have true alphas (gross, before fees) of 1.25 percent or greater—while another one-sixth have “true alphas” of negative 1.25 percent or worse. (True negative alpha is a bit of a puzzling concept. You should not be able to reliably underperform the market either, as all I have to do is short what you buy.) And all of this before fees.

A Supply-and-Demand View of Active Management and Its Fees

It seems the average investor should save 60 basis points a year and just buy a passive index such as Vanguard’s Total Stock Market Portfolio. It seems that the stock pickers should do something more productive, like drive cabs. Active management and its fees seem like a total private, and social, waste.

Yet his hallowed view—and its antithesis—do not completely make sense. After all, active management and fees have survived 40 years of efficient-market disdain. Economists who would dismiss “people are stupid” as an “explanation” for a pricing anomaly that lasts 40 years surely cannot use the same “explanation” for the persistence of active management. Economists who think the evidence favors lots of “inefficiencies” in the market are even less well placed to deplore active management. They should conclude that we need more, or at least better, active management to correct the market’s inefficiencies. Their puzzle is the inability of existing managers to pick low-hanging fruit.

Progress is being made at last. Berk and Green (2004) have created a supply-and-demand economic model that explains many of the basic facts of mutual fund performance, flows, and fees. (Berk 2005 offers a simple exposition.)

Suppose that some fund managers do have alpha. Alpha, however, has diminishing returns to scale. Traders report that many strategies apply only to smaller stocks (see evidence in Fama and French 2006) or that prices move against them if they try to execute trades that are too large. As an example, suppose that a manager can generate 10 percent risk-free alpha with \$10 million in assets under management. Suppose also that the manager’s fees are 1 percent of assets under management, and suppose that the market does not go up or down. Then, in his first year, the manager makes \$1 million abnormal return. The manager pockets \$100,000 and investors in the fund receive \$900,000.

Seeing these good results, investors rush in. But the manager’s idea cannot scale past \$10 million of assets, so the manager invests extra money in an index. With \$20 million under management, the manager generates \$1 million alpha on the first \$10 million and nothing on the rest. The manager again receives 1 percent of assets under management, which is now \$200,000. But investors still get \$800,000 alpha. More investors pour in.

The process stops when the manager has \$100 million under management. The manager still generates \$1 million alpha, but now he collects \$1 million in fees. His investors get exactly zero alpha, the competitive rate of return. Everyone is acting rationally.

Berk and Green's (2004) model is much more sophisticated than this simple example. They include uncertainty in returns and a signal extraction problem for investors, which give rise to interesting dynamics. A large literature has followed.

This model explains many puzzling facts: In equilibrium, returns to investors are the same in active and passively managed funds. Funds earn only enough alpha to cover their fees. Good past fund returns do not forecast good future returns. Investors chase managers with good past returns anyway, seemingly irrational behavior and thus one of the most famous puzzles in the mutual fund literature (for example, Chevalier and Ellison 1997). Returns to investors do not measure alpha. Fees do. Managers with good track records get paid a lot.

This model is the focus of the current debate. Fama and French (2010) complain that the average alpha before fees is nearly zero—and negative, not zero, after fees. Berk and Van Binsbergen (2012) answer that Fama and French's benchmarks are not tradable, and skill should be measured as alpha times assets under management, as 0.1 percent alpha on a billion dollars is a lot. Using these measures, they find investors just about breaking even, and a good deal of positive skill. (Fama and French's Table AI agrees.) The model needs to be brought to the data quantitatively: Does the magnitude of fund flows following performance follow the model's predictions? Does it describe fund exit, the persistence of negative alpha, and the shift to passive management? Like all models, one can explore deeper foundations. What is this alpha, anyway? Why are fees a flat percentage of assets under management? If the manager could simply charge a \$1 million fee to start with, the fund would not need to expand.

And all that is how it should be. After 40 years, the research agenda is finally about how to fit the facts into a supply and demand framework. Arguing about benchmarks, calibration, and optimal contracts is a lot more productive than deploring the financial industry as folly, or declaring that if it survives, markets must be working. The answer will surely not end up all on one side or another: Surely some investors have overpaid for pointless trading. Surely there is some durable value in an industry that has lasted so long. Surely there are some understandable distortions. On this path, we may finally understand how this market works, and maybe, humbly, suggest some improvements. This is a great example of how the economic framework operates—and a sobering reminder of how long it often takes to see that a straightforward economic analysis is possible.

Is It Silly to Pay a Proportional Fee?

Much of the argument that "finance is too big" rests on the view that fees based on a proportion of assets under management are a suboptimal contract. Assets under management went up, fees went up, and managers laughed all the way to the bank. This is a big part of Greenwood and Scharfstein's story in this issue. On closer examination, this argument seems awfully strained.

First, we have seen in fact a substantial decline in many fees and migration to lower fee vehicles, in mutual funds, exchange traded funds, and many wealth management services. Competition does seem to be working, though more slowly than we may like.

Second, fee revenue is not a good measure of the “size” of finance. Fees are a transfer, like gambling losses, not a measure of resources consumed or output produced. Policy may and obviously does care a lot about transfers, but that is a conceptually different question than worrying about wasted resources. Moreover, fees vary based on outcomes. If the fund gains or loses money, fee income rises and falls as well. Hedge fund fees, usually 2 percent of assets and 20 percent of profits, vary enormously. The same fees that were puzzlingly high in 2006 were a lot lower in 2008. Fees have much of the character of a risk-sharing arrangement among co-investors, rather than an expense for professional services.

Third, if the fund doubles in value because everything else in the economy doubles—capital stock, earnings, and so on—then surely by constant returns to scale, the value of investment management (whatever that is) also doubles.

And finally, I’d like to see a specific claim as to what the alternative, realistic, and privately or socially optimal contract is. Funds cannot bill by the hour, passing on “cost” as lawyers do (or rather, used to do), for obvious monitoring and principal-agent reasons. Should we agree to pay a fraction of initial investment, regardless of subsequent performance? It’s obvious why we don’t do that. Accounting for different vintages of investment would be a nightmare. It would also violate the regulatory principle that all investors must be treated equally.

Proportional fees seem almost inescapable in funds that allow investors to withdraw money and invest freely. Suppose funds charge 1 percent for new money, but do not lower dollar fees after losses. Then after a fund has lost half its value, its investors face 2 percent fees going forward. They will quickly withdraw their money and give it to a new fund. Funds that lost money would quickly spiral out of existence, or investors would undermine the fee by withdrawing and reinvesting the next day as new money. Venture capital, private equity, and some hedge funds do not allow free withdrawal so for them, this argument does not apply as strongly—and they have more complex fee structures.

Percentage fees pervade professional services. Real estate agents charge percentage fees, and do better when house prices rise. Architects charge percentage fees. Contingency-fee lawyers take a percentage of winnings. Salesmen get percentage commissions. Even corrupt officials often take percentage bribes.

Perhaps the argument boils down to the claim that there is no alpha, so nobody should pay any fees at all for active management. That’s a different question. If there is alpha or some other function of active management, its optimal contract is a difficult (and much-studied, though I do not review it here) principal-agent problem. Skill is hard to measure, and a fund’s actions are hard to monitor. It seems a big jump to conclude that percentage fees came into existence and have persisted for decades, across a wide range of industries, while inflicting important private and social costs, just because people are naive or irrational in some unspecified way.

Are Fee-Payers Naive?

Delegating active management and paying large fees is common and increasing among large, completely unconstrained, and very sophisticated investors. For

example, the Harvard endowment was in 2012 about two-thirds externally managed by fee investors and was 30 percent invested in “private equity” and “absolute return,” largely meaning hedge funds.¹ The University of Chicago endowment is similarly invested² in private equity and “absolute return.” Apparently, whatever qualms some of its curmudgeonly faculty express about alphas, fees, and active management are not shared by the endowment. Its most recent annual report states: “The majority of TRIP’s [Total Return Investment Portfolio] assets are managed by external managers specializing in a specific asset class, geography, or strategy. These asset managers outperformed their respective benchmarks in every asset class, adding over 500 basis points of performance versus the strategic benchmark.” Five hundred basis points! Put that in your pipe and smoke it, efficient marketers. At least we know one active manager’s perception of what they get for their fees.

These endowments’ approach to portfolio management is pretty much standard at endowments, nonprofits, sovereign wealth funds, family offices, pension funds, and so forth—anywhere there is a big pot of money to invest. These investors pay a lot of attention to allocation among name-based buckets, as represented in the pie charts, “domestic equity,” “international equity,” “fixed income,” “absolute return,” “private equity,” and the like. Then, they allocate funds in the buckets to groups of fee-based active managers.

This approach bears no resemblance to standard portfolio theory, in which an investor pays attention only to means and covariances, not buckets. And don’t even ask how often hedge fund manager A is shorting what B is buying; what happens to fees when you give a portfolio of managers 2+20 compensation and half of them win and half lose; or why one would pay the manager of a growth-oriented fund to buy the same stock that the manager of the value-oriented fund just sold.

Why have these decision procedures become standard practice? Vague reference to “agency problems” and “naiveté” seem unpersuasive. Harvard’s endowment was overseen by a high-powered board, including its president Larry Summers, possibly the least naive investor on the planet. The picture that Summers and his board, or the high-powered talent on Chicago’s Investment Committee are simply too naive to demand passive investing, or that they really want the endowments to be invested in the Vanguard total market index, but some “agency problem” with the managers they hire and fire with alacrity prevents that outcome from happening, simply does not wash. (Yes, delegated portfolio management is a classic principal-agent problem. But no, it’s hard to conceive that it produces this result.) Perhaps instead, we should admit that standard portfolio theory is not much help in situations of any real-world complexity, try to understand what these rough and ready procedures achieve, and offer more helpful advice.

As for “excessive” compensation, in the first layer of fees (fees to the manager who pays fees to the other managers) Harvard endowment’s CIO Jane Mendillo

¹ See the Harvard Management Company website: <http://www.hmc.harvard.edu/investment-management>.

² See the University of Chicago’s *Annual Report*, “The Endowment”: <http://annualreport.uchicago.edu/page/endowment>.

was paid \$4.7 million, most of which was straight salary.³ The University of Chicago's Mark Schmid gets only \$1.8 million, though our measly \$5.6 billion assets under management relative to Harvard's \$27.6 billion may have something to do with it. If major nonprofit university endowments are paying this much, is it really a puzzle that pension funds do the same thing?

Finding Alpha? Implications for Active Trading

To justify fees for active management, one must explain why active trading is worthwhile. The *average investor theorem* is an important benchmark: The average investor must hold the value-weighted market portfolio. Alpha, relative to the market portfolio, is by definition a zero-sum game. For every investor who overweight a security or invests in a fund that earns positive alpha, some other investor must underweight the same security and earn the same negative alpha. Collectively, we cannot even rebalance. And each of us can protect ourselves from being the negative-alpha mark with a simple strategy: hold the market portfolio, buy or sell only the portfolio in its entirety, and refuse to trade away from its weights, no matter what price is offered. If every uninformed trader followed this strategy, informed traders could never profit at our expense.

Alphas and Multiple Factors

Alpha seems a dicey proposition. But the last 20 years of finance research is as clear as empirical research in economics can be: There is alpha relative to the market portfolio—there are strategies that deliver average returns larger than the covariation of their returns with the market portfolio justifies—lots of it, and all over the place. In Cochrane (2011), I provide a summary of this huge literature; I won't provide a separate citation for each fact here.

Examples of such strategies include value (stocks with low market value relative to accounting book value), momentum (stocks that have risen in the previous year), stocks of companies that repurchase shares, stocks of companies with accounting measures of high expected earnings, and stocks with low betas. The “carry trade” in maturities, currencies and, credit—buy high-yield securities, sell low-yield securities—and writing options, especially the “disaster insurance” of out-of-the-money put options, all generate alpha. Expected returns on the market and most of the anomaly strategies vary predictably over time, implying profitable dynamic trading strategies.

Many of these anomalies lead to new “factors,” new dimensions of “systematic” risk and rewards. For example, if one buys a large portfolio of “value” (low-price) stocks, engineered to have zero correlation with the market, thinking that one will reap the value-stock alpha and diversify away the risks, one soon discovers the

³ See “Chart: Top Paid CIOs of Tax-Exempt Institutions,” <http://www.pionline.com/article/20111107/CHART04/111109905>.

tendency of all value stocks to rise and fall together. The portfolio remains risky no matter how many stocks one adds. In this way, pursuing the “value” alpha requires one to take on this additional dimension of undiversifiable risk.

As formalized in Fama and French’s (1996) three-factor model and its larger successors, the world appears to have many such “factors,” acting as the market return factor did in our early understanding, each offering orthogonal dimensions of risk and a return premium to those investors who are willing to take the risks. Those “factor premiums” capture most of observed “alpha” relative to the market portfolio.

Large risk premiums opened up in the recent financial crisis, as prices of very nearly identical securities diverged. For example, corporate bonds traded at lower prices than their synthetic replication by a Treasury bond and a credit default swap. The “covered interest parity” condition failed: You could earn money by borrowing dollars, buying euros, investing in European money markets, and converting back to dollars in the futures markets. If you could borrow dollars! These events and other price movements in the crisis suggest to the researchers studying them “fire sales,” “financial constraints,” “financial frictions,” “price pressure,” and “limits to arbitrage”—all of which are ways of saying that the active managers of the time were insufficient to equalize prices of nearly identical securities, and active traders could have made alphas. Similar pricing divergences and insufficient arbitrage appeared in the trading frenzies of the Internet boom (for example, Lamont and Thaler 2003; Cochrane 2003).

There are multiple dimensions of risk, and bearing these risks generates expected-return rewards, rewards that change over time. These facts are not really under debate. Their interpretation is. These alphas might represent imperfect risk sharing and (often temporary) market segmentation, or “sentiment,” irrational attachment or aversion to broad categories of securities. They might also reflect a multidimensional and time-varying nature of risk premiums in a fully-integrated and informationally efficient market. They certainly look less and less like “information” about individual securities that is somehow improperly reflected in prices.

These facts and interpretations lend a quite new color to our central questions: Is the financial sector too large or too small? How should investors behave in a world with multiple dimensions of systemic risk? What is the economic function of active management, and the economic value of management fees?

Multidimensional Risk-Sharing

The conventional disdain for active financial management is based on a conventional perspective: The market portfolio is the one and only source of “systematic” risk which generates a premium. It is accessible through low-cost passive investments. The investor understands this opportunity and knows how much market risk he or she wishes to take. Alpha represents the trader’s knowledge of information not reflected in market prices.

But the dozens of semi-passive strategies, each of which produce alpha (relative to the market), each of which exposes the investor to new dimensions of undiversifiable risk, and many of which are poorly understood, changes the picture completely.

Each investor needs to decide which of the many sources of risk he or she is best able to bear, or needs to avoid despite their attractive premiums.

Investors need to consider the even larger set of asset market risks that do not bear premiums. Before chasing alphas, investors should hedge the risks of their jobs, businesses, outside income streams, real estate, or peculiar liability streams by setting up portfolios of assets whose returns are negatively correlated with those risks. You should want a portfolio that rises when there is bad news about your future income. Curiously, academic finance has done little to characterize these nonpriced risks and prescribe hedging strategies.

One can see this process beginning. Many pension funds are moving towards bond-like investments to match their liabilities. University endowments are beginning to recognize how their liability streams affect investments. They thought of themselves as “long term” investors able to reap the premiums of illiquid investments, and able to wait patiently through market downturns, until many in the crisis realized they were supporting a bond-like liability stream in salaries of tenured professors, and were leveraged by bond-financed construction. They found themselves trying to sell illiquid assets at the bottom like everyone else. Now, they are thinking about matching endowment funding to projects that can bear risk and adapting portfolios to their cash flows, including the implicit beta that alumni donations rise when the stock market goes up. Endowments are recognizing that their objectives include an important tournament relative to other universities (Goetzmann and Oster 2012). The wealth-management arms of big banks help to set up hedge portfolios for executives who have large unsaleable stock or option positions, to help them come as close to shorting their own business as possible. Websites available to individual investors are starting to emphasize intelligent and individual-specific choice of “style” rather than promise generic “alpha.”

But none of this is easy. Merton (1971) described state-variable hedging demands 40 years ago. Yet, with thousands of following papers, academic portfolio theory still really does not offer clear-cut real-world advice (Cochrane forthcoming).

The nature and amount of multidimensional systematic risk one should take is also much more nebulous and difficult to assess than the traditional question of how much market risk one should take. Should you write put options, to earn the premium? Or maybe you should buy put options as disaster insurance? Are you positioned to buy value stocks? To take on the credit risks of default? To take the risk that high-interest rate foreign currencies depreciate against the dollar? Do the alpha premiums these strategies offer compensate for the risks you will suffer when they lose money? The whole alpha/beta definition is falling apart.

Even then, taking advantage of time-varying multidimensional risks requires technical knowledge. Do you know how to write a credit default swap contract, how to make stock momentum strategy work without drowning in transactions costs, how to take advantage of temporarily high put option premiums in the euro-zone, or even how reliably to buy a “value” portfolio? Because such questions are not easy, portfolio problems like this might certainly benefit from professional and specialized management, and such management ought to be able to charge a fee.

Perhaps some of the puzzling features of investment practice might be understood as a rough and ready way of adapting to this more realistic portrait of risks and returns. If so, some active management and dynamic trading represents a form of socially beneficial insurance provision.

Hedge funds might make more sense in this investment world. They can move to and from asset classes as risk premiums change, and by using leverage and derivatives they can alter overall exposures quickly without incurring the transactions cost of buying and selling large portfolios.

Many of these alpha-generating strategies and new “factors” suggest needed institutional developments. As a concrete and recent example, consider the “betting against beta” anomaly reexamined by Frazzini and Pedersen (2011a, b). They document that low-beta stocks get higher average returns than they should, and high beta stocks get lower returns than they should. Their interpretation is that many investors want more risk than the market portfolio provides, yet leverage is costly to obtain. These investors buy high-beta stocks instead of leveraging, driving up the prices of high-beta stocks, and vice versa for low-beta stocks. In this setting, arbitrageurs cannot help. The problem is the price of risk, needing wider risk-sharing, not an arbitrage (riskless profit) opportunity. To bring prices back to what they should be, we need low-cost vehicles to bring leveraged low-beta investments to the part of the investing public that wants them—which, perhaps not coincidentally, Frazzini and Pedersen’s company provides.

We have seen this kind of institutional development before. Small stocks were one of the first prominent anomalies, generating (it appeared) higher average returns than their betas justified. But it was hard for individual investors to hold a diversified portfolio of small stocks. Arbitrageurs could only do so much, because small stocks move together, so a concentrated portfolio bears undiversifiable risk. Small stock mutual funds were started, which allowed a mass of investors to participate. Fees and expenses of those funds contributed to revenue and measured GDP, in a way that the activities of individual investors holding small stocks did not. But they allowed the risk of small stocks to be widely shared and the small stock premium to decline.

So far I have made no mention at all of informational inefficiency, exploiting mispricings, superior information, or winning the zero-sum alpha game. I have not violated the average investor theorem. Given the new facts of empirical finance, a large role for active management exists without any of that at all. Of course, I do not claim that current portfolio practice, and especially hiring many different high-fee hedge funds, is an optimal strategy. But it isn’t necessarily as “naive” or “agency conflicted” as it otherwise seems.

Marketing

In the quest to explain the persistence of active management and its fees, one other analogy seems worth pursuing: marketing. Marketing and advertising have long been a puzzle to economists, along with readers of *Consumer Reports* and coupon-clippers everywhere. Why buy the brand name when the generic is nearly identical, and costs a lot less?

The money-management industry is essentially a marketing industry. Its practitioners take generic ingredients, package, label, advertise, and market them. Yes, it's puzzling that people don't buy the generic at Vanguard. It's puzzling that they don't buy the pieces and assemble their own, with E*TRADE. It's puzzling that they pay so much for the slight differences in ingredients that the active managers deliver. And it is equally puzzling that they pay for Coke, Clorox, Bayer, or bottled water; that they shop at Macy's not Target, Whole Foods not Costco, and a hundred other brand names.

This is not the place to digress into the "rationality" of marketing and advertising. Simply dismissing centuries worth of branding and advertising as naiveté and folly seems, well, its own form of naiveté. Perhaps by thinking of active fund management as an instance of this larger pattern, we may make some progress to understanding how it actually works.

Information Trading and Price Discovery

Much trading and active management, however, is clearly aimed at bringing information to the market, not at better sharing of time-varying and multidimensional risk. The first welfare theorem does not clearly apply to information production, so we have little a priori reassurance that the quest for trading profits produces the "right" amount—or, perhaps more importantly, the right *kind*—of information.

It is possible that *not enough* social resources are devoted to trading, because information is a public good. As French (2008) wrote, despite deploring the private costs of alpha-chasing: "I offer no evidence on whether society is buying too little or too much of this good. Price discovery, however, is an externality—each active investor pays the full cost of his efforts but captures only a tiny slice of the benefit—so there is no reason to think active investors purchase the optimal amount of price discovery."

The common complaints "the financial crisis proves markets aren't efficient," or that tech and mortgages represented "bubbles," are at heart complaints that there was not *enough* active information-based trading. All a more "efficient" market could have done is to crash sooner, by better expressing the pessimist's views. Remember, "efficiency" means that prices incorporate all available information, not that markets are clairvoyant. The definition of "efficiency" is widely misunderstood. I once told a newspaper reporter that I thought markets are pretty "efficient," and he quoted me as saying markets are "self-regulating!"

If information is *not* incorporated into market prices and to such an extent that simple strategies with big alphas can be published in the *Journal of Finance*, there are not enough arbitrageurs. If asset prices fall in "fire sales," only to rebound later, there are not enough buyers following the fire trucks. If credit constraints are impeding the flow of capital, there is a social benefit to loosening those constraints.

The literature on short-selling is revealing on this point. Short sellers uncover far more financial fraud than the Securities and Exchange Commission. Conversely,

some of the biggest alphas and “inefficiencies” occur when there is a technical or regulatory impediment to short seller’s activities. Lamont (2012) finds 2.4 percent monthly alpha to a portfolio of short-selling-constrained stocks, a large informational inefficiency. This is a concrete example of inadequate (because constrained) information-based trading.

Information trading produces more informationally-efficient prices, which are socially useful. With better market signals, companies raise capital more easily for valuable projects, and are signaled not to invest in poor projects or at poor times. True, the simple q theory, which predicts that corporate investment should be a perfect function of stock price relative to book value, is formally rejected, but its glass is also half full: There are strong correlations between stock prices and investment, over time (through the tech boom and bust of the 1990s and through the financial crisis (see Cochrane 1991; 2011, Figure 10)) and across industries (Google versus, say, GM). When issuing stock generates a lot of money, companies do it, and build factories or websites. Those who view asset market booms and following busts as “irrational” or “bubbles” point to the consequent investment booms and busts as examples of the social costs of inefficient markets, thereby endorsing the social value of more efficient markets.

Even without investment, more efficient prices provide better risk sharing, as in an endowment economy. If the owner of an apple tree and that of a pear tree hedge their risks by trading stock in the other tree, their risk-sharing improves when stock prices are more efficient. (Hirshleifer’s, 1971, famous analysis stating that efficiency is only socially beneficial if production is involved did not treat such risk-sharing).

Information trading is central to “liquidity provision” and thus the success of markets for risk sharing. Markets such as Consumer Price Index, GDP futures or hurricane catastrophe options failed because there was not enough information trading. This is an important external benefit. Indeed, in the public forum, hedge funds and high-frequency traders primarily defend their activities by touting their “market making” and “liquidity provision” for small investors. (Of course, they are also pandering to their regulators’ tastes here.)

The Puzzle of Information Trading

Still, the cacophony of trading seems like a lot of effort for these goals. The classic theory of finance predicts that information is perfectly reflected in prices, with no trading volume needed. Suppose Apple is trading at \$500 per share, but you know that the iPhone 6 will make Apple worth \$1,000 per share. If you approach an uninformed investor with an offer to buy Apple at \$600 per share, the index investor should answer: “No, you must know something I don’t know. I only buy and sell the entire index, so I don’t lose to people like you.” If you offer \$700, the index investor answers: “I don’t think you heard me. I only buy and sell the entire index.” You keep trying, bidding the price up all the way to \$1,000 per share, at which point you give up. The price rises, reflecting your information, but no trade occurred. This is a colloquial version of Milgrom and Stokey’s (1982) famous no-trade theorem.

The theory that prices reflect information with zero trading volume is of course dramatically at odds with the facts. The classic theory also ignores costs. If information traders cannot earn positive alpha, and if producing information and trading on it takes any time and resources, the information traders won't bother, and nobody is left to make prices reflect information. For this reason, as Grossman and Stiglitz (1980) wrote, informationally efficient markets are impossible.

The standard compromise model (Grossman and Stiglitz 1980; Kyle 1985; and a huge literature) posits “informed” traders who receive a signal about a firm's value, “liquidity” traders who for unspecified reasons must trade, and “market makers” who intermediate, charging a bid-ask spread to defend themselves against the informed traders.

Now, all current theories of trading rely on some sort of “irrationality” or other artificial assumptions. “Liquidity traders” are the classic example. Other models, like Scheinkman and Xiong (2003), posit slightly irrational dogmatic beliefs so each information trader can believe he or she is smarter than average. Many models, such as Acharya and Pedersen (2005), write down overlapping generations of agents without bequests who die every week or so, forcing them to trade.

But these assumptions are convenient shortcuts for getting trading into the model for other purposes, such as studying price discovery and liquidity. They are not there to describe microfoundations of socially destructive trading that needs remediation by policy. The “irrationality” that breaks the no-trade theorem, or the irrationality of the liquidity traders, is not typically deeply micro-founded in the psychology literature, as in true behavioral finance. People live more than a week, and leave bequests.

The fact staring us in the face is that “price discovery,” the process by which information becomes embedded in market prices, uses a *lot* of trading volume, and a lot of time, effort, and resources. And we are only beginning to understand it.

The empirical literature offers tantalizing glimpses of this process. A very small taste of this vast literature: The period after a news announcement often features high price volatility and trading volume, in which markets seem to be fleshing out what the news announcement actually means for the value of the security. For example, Lucca and Moench (2012, Figure 6) show a spike in stock-index trading volume and price volatility in the hours just *after* the Federal Reserve announcements of its interest rate decisions. The information is perfectly public. But the process of the market digesting its meaning, aggregating the opinions of its traders, and deciding what value the stock index should be with the new information, seems to need actual shares to trade hands.⁴ Perhaps the common model of information—essentially, we all agree on the deck of cards, we just don't know which one was picked—is wrong.

Securities such as “on the run” or benchmark bonds, where “price discovery” takes place, have higher prices than otherwise identical securities. Traders are

⁴ Banerjee and Kremer (2010) and Kim and Verrecchia (1991) offer models in which such disagreement about public information leads to trading volume.

willing to suffer lower average returns in order to participate in the information-trading game, in much the same way as money holders suffer lower returns for the transactions services money provides (see Cochrane 2003 and references therein). Similarly, “liquidity” seems to be extremely valuable to investors and has been so for a long time, even though none of us feel the need to trade every 10 minutes.

Markets in financial securities are set up, and exist, almost entirely to be markets for *information trading*, and high-frequency “liquidity provision,” that we find hard to fathom. They are not really markets for the *securities* themselves. We could easily handle individuals’ lifetime saving and dissaving needs, and firms’ need to issue and retire equity, with orders-of-magnitude less volume, in much sleepier bank-like institutions. Yes, we could each avoid being the negative-alpha part of price discovery by only buying index funds. It’s a bit of a puzzle that we don’t. It’s also a good thing we don’t, or there would be no traders making prices efficient.

But as with active management, perhaps we should work just a little harder before dismissing the hundreds of years of trading activity, and the entire existence of the New York Stock Exchange, Chicago Mercantile Exchange, and other markets, as monuments to human folly, or before advocating regulations such as transactions taxes—the perennial favorite answer in search of a question—to reduce trading volume whose size, function, and operation we do not understand. Are we sure that they should not be transactions subsidies?

And before we deplore, it’s worth remembering just how crazy passive indexing sounds to any market participant. “What,” they might respond, “would you walk in to a wine store and say ‘I can’t tell good from bad, and the arbitrageurs are out in force. I sure won’t pay you 1 percent for recommendations. Just give me one of everything’?”

High-Frequency Trading and Market-Making

It’s especially hard to see why high-frequency trading is needed. Price discovery every millisecond doesn’t seem necessary to guide corporate investment or individual risk sharing and hedging.

High-frequency trading reminds us in the extreme that the amount of trading based on a well-understood or “fundamental” piece of information about a company’s cash flow is minuscule. Models in which an informed trader possesses a “signal” about the value of a liquidating dividend just don’t describe the vast majority of trading. High-frequency traders do not trade on earnings reports 20 milliseconds ahead of the market.

Instead, high-frequency traders—and even most “low-frequency” day and week traders—look at patterns of prices, volumes, and past trading activity, not “information” or opinion about firm fundamentals.

They may describe their strategy as “statistical arbitrage,” removing the small predictability of high-frequency price movements (and grossly misusing the term “arbitrage.”) Sometimes they defend their social function as “market makers” or “liquidity providers.” If so, market making is a far more dynamic process than simply posting bid-ask spreads, as the standard theory envisions! If you ask their critics, they

are artfully front-running demand from less-sophisticated investors, subtracting “liquidity,” worsening “price impact,” choking bandwidth with quickly-canceled orders and removing the economic rewards to genuine information trading. Their activity may also answer the interesting question of how information spreads from one informed trade to the whole market. Somebody has to notice the price pattern and pile in.

However we come to understand these issues, the social costs and benefits of high-frequency trading are clearly not at all related to the minor (as a fraction of GDP) resources devoted to them—the cost of possibly useless fiber-optic cable, co-located servers, and the time of smart programmers who could be developing better iPhone games. The social question for high-frequency trading—like all of finance, really—is whether it screws up markets or makes them more efficient and “liquid.”

There isn’t yet much evidence or theory on this point, but isolated events suggest doubts about liquidity-provision and efficiency. For example, in the May 6, 2010 “Flash Crash,” the Standard and Poor’s 500 fell 6 percent in a few minutes after a large sell order arrived, and promptly recovered in less than an hour, only after a five minute trading halt. Kirilenko, Kyle, Samadi, and Tuzun (2011) who study this event (see their Figure 1) document that high-frequency traders absorbed demand for about four seconds before turning around and selling along with everyone else. On July 19, 2012, Coke, McDonalds, IBM, and Apple saw price sawtooths: sharp rises exactly on each hour, reversed by the next hour. Vigna and Lauricella (2012) offer some amazing graphs.⁵ These movements were widely attributed to an algorithm placing big orders exactly on the hour—and other algorithms not picking up on the inefficient signal abundantly obvious to the human eye. These palpable inefficiencies suggest a market with very little “liquidity provision,” not the opposite.

The structure of markets, with design and regulation stemming from the days of human trading, could be at fault. Prices must jump in discrete intervals—once 1/8 dollar, now 1 cent. Limit orders must be filled in strict time priority: if order A arrives before order B, order A must be filled completely and B gets nothing. Yet time is continuous. A’s order need only arrive a millisecond before B’s, and A wins the pot. (Traders report that the ability to quickly cancel limit orders that are in the back of the line is another advantage of very high speed.) You can see an arms race for speed emerge. It’s worth spending a lot on computers to speed up trades by a few milliseconds.

If my hunch is correct, it suggests an obvious solution: Suppose that an exchange operated on a discrete clock, as a computer does in order to let signals settle down before processing them. The exchange could run a once-per-second, or even once-per-minute, matching process, with all orders received during the period treated equally. If there are more buy than sell at the crossing price, orders are

⁵ The website <http://www.nanex.net/FlashCrash/OngoingResearch.html> is devoted to weird behavior in high-frequency markets.

filled proportionally. Such an exchange would eliminate extremely high-frequency trading, because there would be no gain or loss from acting faster than a minute.

Would this system be an improvement, to efficiency and liquidity? Would exchanges choose such systems if they were allowed to do so? The Taiwan Stock Exchange already matches limit orders once every 90 seconds (Barber, Lee, Liu, and Odean 2008). Is its performance atrociously worse? These are all good questions! High-frequency trading is a ripe area of research.

Housing, Consumer Credit, and the Size of Regulated Finance

The growth of housing finance and consumer credit raises a different set of issues. It's useful to divide the mortgage business into three parts: mortgage origination, mortgage refinancing, and mortgage-backed securities.

The increase in fees for residential loan origination is easily digested as the response to an increase in demand. The increase in housing demand may indeed not have been "socially optimal" (!). There are plenty of government policies and perhaps a few market dislocations to blame. But it doesn't make much sense to criticize growth in the financial industry for responding to this increase in demand, whatever its source, or for passing along the subsidized credit—which was and remains the government's explicit intention to increase—with the customary fee.

The large fees collected for refinancing mortgages are a bit more puzzling. US mortgages are strangely complicated, predominantly featuring fixed rates, no penalty for prepaying when interest rates fall, limited recourse, and a complex refinancing option. Other countries have gravitated to much simpler contracts. The now-familiar structure of US mortgages emerged after only the Great Depression, when new federal agencies started issuing them. Before the Great Depression, US mortgages lasted only five to ten years and required only the payment of interest. The principal was due at the end of the loan, and was typically refinanced (Green and Wachter 2005, p. 95). Today, the structure of mortgage contracts is pretty much dictated by what the government agencies that dominate the market will buy and guarantee.

These observations suggest that such complex contracts are not a market necessity. However, a glance at my cellphone contract and frequent flyer miles rules suggests to me that price discrimination by needless complexity might be part of the story as well.

Still, collecting fees when interest rates decline or consumers refinance is not conceptually part of GDP. They are state-dependent transfers dictated by the terms of an option contract. And we are unlikely to see a lot of refinancing as interest rates eventually rise.

There was a lot of financial innovation in mortgage-backed securities, some of which notoriously exploded. But here again, whether we spend a bit of GDP filling out forms or paying fees is clearly the least of the social benefit and cost questions. The "shadow banking" system was prone to a textbook systemic run, which happened. This fragility, not the size or fraction of GDP, is the important issue.

A good part of this innovation, such as creating off-balance-sheet, special-purpose vehicles and tailoring securities in order to game credit ratings, was clearly designed to engineer around ill-conceived regulations. That part counts as a regulatory failure needing reform, rather than a market failure needing additional regulation.

Yet much of this financial innovation has the potential to be of large social benefit. Suppose that mortgages were bundled into securities, intermediated by mutual funds whose values float, just like those of equity mutual funds, and held around the world in retirement accounts, pension funds, and our endowments' portfolios, without government guarantees at every step. This would be a terrific financial structure. Though mortgage-backed securities are a bit opaque, they are nowhere near as opaque as the entire balance sheet of, say, Citigroup. Furthermore, such a structure would be immune to runs, bankruptcies, and bailouts, thus requiring minimal regulation. And the fees required to fill out the mortgage-backed security paperwork would surely be less than the bank and regulatory paperwork, regulation, and compliance costs of the current system.

Concluding Remarks

The size and revenues of the finance industry increased because fee income for refinancing, issuing, and securitizing mortgages rose along with the rise in housing transactions and house prices, and because asset-management fee income rose along with a shift to professional management from "roll-your-own" portfolios and a rise in asset values. Compensation to employees with skills in short supply increased. Fee schedules themselves declined a bit. These facts suggest "demand shifted out," not "something big changed in the structure of this industry."

Demand that shifts out can shift back again. Demand for financial services evaporated with the decline in housing and asset values in the 2008 recession and subsequent period of sclerotic growth. Much of the "shadow banking system" has disappeared. For example, asset-backed commercial paper outstanding rose from \$600 billion in 2001 to \$1.2 trillion in 2007—and now stands at \$300 billion. Financial credit market debt outstanding in the flow of funds rose from \$8.6 trillion in 2000 to \$17.1 trillion in 2008—and now stands at \$13.8 trillion. Employment in financial activities rose from 7.7 million in 2000 to 8.4 million in 2007—and is now back to 7.7 million (according to the Bureau of Labor Statistics). Study of "why is finance so big," using data that stops in 2007, may soon take its place alongside studies of "why are Internet stocks so high" in 1999 or studies of "why is there a Great Moderation" in 2006.

An older literature on the size of the financial system, forgotten in the current debate, studies the socially inefficient resources devoted to cash management in the face of positive interest rates, measuring social costs as the area under the money demand curve. Lucas (2000) concluded that finance was about 1 percent of GDP too big by this measure. The fragility of those cash-management schemes can now

be added to the list of social costs. Zero interest rates have eliminated these costs for now, and if the Fed continues to pay market interest on reserves, those costs can remain largely eliminated in the future.

The size question for the finance industry going forward, under the Dodd–Frank regulatory structure, is likely to be how many resources are devoted to regulation, regulatory compliance, lobbying to influence those regulations, and the distortions they induce. The social cost question remains how to create a financial system that is not prone to runs, crashes, and bailouts, even if that costs a few percentage points of GDP. Unless sovereign debt bites us first.

Many puzzles remain in the structure of the finance industry. The persistence of high-fee active management chosen by sophisticated institutional investors remains a puzzle. To some extent, as I have outlined, this pattern may reflect insurance provision, that is, the dynamic and multidimensional character of asset-market risk and risk premiums. To some extent, this puzzle also goes hand in hand with the puzzle of why price discovery seems to require so much active trading, and whether and how information trading provides valuable “liquidity.” It is possible that there are far too *few* resources devoted to price discovery and market stabilization. In the financial crisis, we surely needed more pools of cash prepared to pounce on fire sales, and more opportunities for negative long-term views to express themselves.

Surveying the current economic literature on these issues, it is certain that we do not very well understand the price-discovery and trading mechanism, nor the economic forces that allowed high-fee active management to survive so long.

Unless we adopt the arrogant view that what we don’t understand must be bad, it is clearly far too early to make pronouncements such as “There is likely too much high-cost, active asset management,” or “Society would be better off if the cost of this management could be reduced.” Such statements are not supported by theory or evidence. Nor is their not-so-subtle implication that resources devoted to greater regulation—by politicians and regulators no less naive than current investors, no less behaviorally-biased, armed with no better understanding than academic economists, and with much larger agency problems and institutional constraints—will improve matters. This proposition amounts to Samuel Johnson’s dictum on second marriages, “the triumph of hope over experience.”

References

- Acharya, Viral V., and Lasse H. Pedersen. 2005. “Asset Pricing with Liquidity Risk.” *Journal of Financial Economics* 77(2): 375–410.
- Banerjee, Snehal, and Ilan Kremer. 2010. “Disagreement and Learning: Dynamic Patterns of Trade.” *Journal of Finance* 65(4): 1269–1302.
- Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2008. “Just How Much Do Individual Investors Lose by Trading?” *Review of Financial Studies* 22(2): 609–632.
- Berk, Jonathan B. 2005. “Five Myths of Active Portfolio Management.” *Journal of Portfolio Management* 31(3): 27–31.
- Berk, Jonathan B., and Richard C. Green. 2004. “Mutual Fund Flows and Performance in Rational Markets.” *Journal of Political Economy* 112(6): 1269–95.

- Berk, Jonathan B., and Jules H. Van Binsbergen.** 2012. "Measuring Managerial Skill in the Mutual Fund Industry." NBER Working Paper 18184.
- Carhart, Mark M.** 1997. "On Persistence in Mutual Fund Performance." *Journal of Finance* 52(1): 57–82.
- Chevalier, Judith, and Glenn Ellison.** 1997. "Risk Taking by Mutual Funds as a Response to Incentives." *Journal of Political Economy* 105(6): 1167–1200.
- Cochrane, John H.** 1991. "Production-Based Asset Pricing and the Link between Stock Returns and Economic Fluctuations." *Journal of Finance* 46(1): 207–234.
- Cochrane, John H.** 2003. "Stock as Money: Convenience Yield and the Tech-Stock Bubble." Chap. 12 in *Asset Price Bubbles* edited by William C. Hunter, George G. Kaufman, and Michael Pomerleano. Cambridge: MIT Press.
- Cochrane, John H.** 2011. "Discount Rates." *Journal of Finance* 66(4): 1047–1108.
- Cochrane, John H.** Forthcoming. "A Mean-Variance Benchmark for Intertemporal Portfolio Theory." *Journal of Finance*.
- Fama, Eugene F., and Kenneth R. French.** 1996. "Multifactor Explanations of Asset Pricing Anomalies." *Journal of Finance* 51(1): 55–84.
- Fama, Eugene F., and Kenneth R. French.** 2006. "Dissecting Anomalies." *Journal of Finance* 63(4): 1653–78.
- Fama, Eugene F., and Kenneth R. French.** 2010. "Luck versus Skill in the Cross-Section of Mutual Fund Returns." *Journal of Finance* 65(4): 1915–47.
- Frazzini, Andrea, and Lasse Heje Pedersen.** 2011a. "Betting against Beta." Paper available at <http://www.econ.yale.edu/~af227>.
- Frazzini, Andrea, and Lasse Heje Pedersen.** 2011b. "Embedded Leverage." Paper available at <http://www.econ.yale.edu/~af227>.
- French, Kenneth R.** 2008. "Presidential Address: The Cost of Active Investing." *Journal of Finance* 63(4): 1537–73.
- Goetzmann, William N., and Sharon Oster.** 2012. "Competition among University Endowments." NBER Working Paper 18173.
- Green, Richard K., and Susan M. Wachter.** 2005. "The American Mortgage in Historical and International Context." *Journal of Economic Perspectives* 19(4): 93–114.
- Greenwood, Robin, and David Scharfstein.** 2012. "The Growth of Modern Finance." July, 2012. http://www.people.hbs.edu/dscharfstein/Growth_of_Modern_Finance.pdf.
- Grossman, Sanford G., and Joseph E. Stiglitz.** 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70(3): 393–408.
- Hirshleifer, Jack.** 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61(4): 561–74.
- Kim, Oliver, and Robert E. Verrecchia.** 1991. "Trading Volume and Price Reactions to Public Announcements." *Journal of Accounting Research* 29(2): 302–21.
- Kirilenko, Andrei A., Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun.** 2011. "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1686004.
- Kyle, Albert S.** 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53(6): 1315–36.
- Lamont, Owen A.** 2012. "Go Down Fighting: Short Sellers vs. Firms." *Review of Asset Pricing Studies* 2(1): 1–30.
- Lamont, Owen A., and Richard H. Thaler.** 2003. "Can the Market Add and Subtract? Mispricing in Tech Stock Carve-Outs." *Journal of Political Economy* 111(2): 227–68.
- Lucas, Robert E., Jr.** 2000. "Inflation and Welfare." *Econometrica* 68(2): 247–74.
- Lucca, David O., and Emanuel Moench.** 2012. "The Pre-FOMC Announcement Drift." http://www.newyorkfed.org/research/staff_reports/sr512.html.
- Merton, Robert C.** 1971. "Optimum Consumption and Portfolio Rules in a Continuous Time Model." *Journal of Economic Theory* 3(4): 373–413.
- Milgrom, Paul, and Nancy L. Stokey.** 1982. "Information, Trade, and Common Knowledge." *Journal of Economic Theory* 26(1): 17–27.
- Pensions & Investments.** 2011. "Chart: Top Paid CIOs of Tax-Exempt Institutions." November 7. <http://www.pionline.com/article/20111107/CHART04/111109905>.
- Scheinkman, Jose, and Wei Xiong.** 2003. "Overconfidence and Speculative Bubbles." *Journal of Political Economy* 111(6): 1183–1219.
- Vigna, Paul, and Tom Lauricella.** 2012. "Sawtooth Trading Hits Coke, IBM, McDonald's, and Apple Shares." *Wall Street Journal*, July 19. <http://blogs.wsj.com/marketbeat/2012/07/19/sawtooth-trading-hits-coke-ibm-mcdonalds-and-apple-shares/>.

