

CHAPTER 1

An Introduction to Exchange-Traded Funds

EXCHANGE-TRADED FUNDS WERE INTRODUCED AS "SOMETHING TO TRADE"

Many of mankind's great innovations owe at least some of their success to serendipity. A popular legend suggests that serendipity helped mankind learn the usefulness of fire. As the story goes, one of our ancestors came upon the site of a fire that had been started by lightning. This early human discovered that the fire had burned an animal's carcass. The "cooked" meat tasted better than raw meat, and men soon learned that cooking enabled humans to obtain nutrition more efficiently, freeing up time and providing energy for other pursuits.¹ This kind of serendipity has been a common theme in many of mankind's endeavors.

One of the best examples of serendipity in the financial markets—from several angles—is the early development of exchange-traded funds (ETFs). In attributing some features of exchange-traded funds to serendipity, we certainly do not mean to minimize the role of the developers of the early exchange-traded funds. They deserve full credit for the wisdom they displayed in designing the early ETFs introduced in Canada and the United States. Although it is not fashionable to credit regulators with a positive role in financial product development, regulators were almost certainly responsible for some of the shareholder protection features of ETFs. Human efforts

¹The precise circumstances under which mankind first encountered and controlled fire—and then began cooking—are likely to remain obscure. Some archaeologists and anthropologists embrace the idea that the first controlled fire became available when men made sparks by hitting stones together while making tools. For a very concise (if conjectural) history see Dreifus (2009). For more details, see Wrangham (2009), especially pp. 190–194.

notwithstanding, however, some key features became part of the ETF by accident. The features of early ETFs were so important that they are now serving as the basis for some revolutionary financial engineering that promises to reshape the fund industry in the United States and around the world.

We will look at the early history of ETFs in some detail in Chapter 2, so these background comments will be brief. The first viable open-end exchange-traded portfolio basket was introduced in Canada and began trading in 1989 as the Toronto Stock Exchange Index Participations (TIPs). It took nearly four more years for the American Stock Exchange to obtain approval from the Securities and Exchange Commission (SEC) to launch the Standard & Poor's 500 SPDR fund product in the United States. In both cases, the principal purpose of the product launch was to provide something for exchange members to trade.

The labels "exchange-traded fund" and "ETF" are applied to a number of financial instruments. The fact that investors can trade most of the products called ETFs throughout the day at market-determined prices that are close to the intraday value of an underlying portfolio or index is one common feature of these securities. Many so-called "ETFs" are neither funds nor investment companies, as defined by the Investment Company Act of 1940. The ETF label has been attached to some open-end structured notes and to a number of grantor trusts, including HOLDRS and various currency- and commodity-based instruments. Vanguard offers exchange-traded share classes of a number of its mutual funds. Vanguard calls these shares ETFs, but these share classes do not have some important features that characterize the ETFs descended from the original SPDR. While the structure of the product does not matter in every case, shareholder protection, tax treatment, and credit risk can vary significantly among the products casually referred to as ETFs. Some observers have called the nonfund instruments exchange-traded products (ETPs) or exchange-traded vehicles (ETVs), but these names have not caught on. "Exchange-traded fund" or the acronym "ETF" is the almost universal generic label for *all* exchange-traded portfolios, open-end structured notes, and securitized commodity products.

While we will discuss all the financial instruments commonly called ETFs, the most significant and useful of these are and will continue to be the true funds. We begin by focusing on two important characteristics of the investment company ETF that are, in some respects, serendipitous. These characteristics have helped attract investors and they have been important in the early success of ETFs. These characteristics also provide a basis for growth in the true fund ETF model well beyond its impressive beginnings. Not everyone attaches as much significance as I do to these two features, but I am convinced that they hold the key to developing better funds. The two key features of these ETFs are *shareholder protection* and *tax efficiency*.




EXHIBIT 1.1 Pre-1968: Buying and Selling Mutual Fund Shares at Yesterday's Net Asset Value

SHAREHOLDER PROTECTION

Four exhibits will help illustrate the value of shareholder protection² and how it is provided by most true fund ETFs. Exhibit 1.1 shows how mutual funds were priced for sales and redemptions prior to 1968. This diagram shows the pattern of fund intraday portfolio values during market trading hours for three consecutive trading days. At the end of each day, a mutual fund calculates its net asset value (NAV) per share based on the end-of-day value of the portfolio. Prior to 1968, the price at which investors invested in the shares of a fund or redeemed their shares was the net asset value *as of the previous day's close*.³

In Exhibit 1.1, the fund publishes its net asset value at the end of Day 1. That value is indicated by the circle at the end of the squiggly price line showing the pattern of intraday values for Day 1. Prior to 1968, that net asset value was the basis for fund share transactions until the following day's market close—and the calculation of a new net asset value. The share price for an order received on Day 2 is indicated by the dotted line extending to the right of the circle through the end of Day 2. Clearly, buying shares of the fund at Day 1's net asset value as the market rose on Day 2 was a great opportunity for trading profit—and for abuse of the fund's established shareholders by opportunistic investors. Correspondingly, if someone wanted to redeem shares in the fund, they would know from the intraday behavior of market

²Some of the material in the balance of this chapter first appeared in Broms and Gastineau (2006; 2007).

³The material described in this and the next few paragraphs is widely known, but not frequently discussed. A recent comprehensive description of mutual fund pricing over the years is available in Swenson(2005), pp. 270–294.

indexes on Day 2 that they could probably redeem at a higher fund share price by waiting until after the determination of net asset value on Day 2. As it became clear that the market was going to close lower on Day 3, redeeming fund shares at the net asset value from Day 2 would have seemed like a better idea than waiting for calculation of Day 3's lower net asset value. It would also be clear during the trading session on Day 3 that the price of buying shares would be lower if the purchase were deferred until Day 4. Backward pricing led to abuses by dealers and by traders who could avoid the fund sales charges or "loads" that were more common in that period than they are today.

In 1968, the fund pricing rules changed. The SEC implemented its Rule 22(c)(1), which required fund share transactions to be priced at the net asset value *next determined* by the fund after the order was received. This meant that anyone entering an order after the close of business on Day 1 would purchase or sell fund shares at the net asset value determined at the close on Day 2. Correspondingly, someone entering an order to purchase or sell shares after the close on Day 2 would be accommodated at the net asset value determined at the close on Day 3. This process is illustrated in Exhibit 1.2.

While any mutual fund share *trader* might have preferred the pre-1968 system, most *investors* would agree that the basic idea behind Rule 22(c)(1) was a sound one. Allowing traders to decide today to buy or sell shares at yesterday's price is unfair to established investors in the fund's shares. However, there is still a transaction fairness problem for fund investors with Rule 22(c)(1) in place. That problem is illustrated in Exhibit 1.3.

By pricing all transactions in the mutual fund's shares *at the net asset value next determined*, as required by Rule 22(c)(1), the fund still provides *free liquidity* to investors entering and leaving the fund. As Exhibit 1.3 shows, anyone purchasing mutual fund shares for cash gets a share of the securities positions already held by the fund and priced at net asset value.





EXHIBIT 1.2 Since 1968: Buying and Selling Mutual Fund Shares at the Net Asset Value Next Determined




Note: Share purchases and redemptions are priced at the next net asset value calculated by the fund.

EXHIBIT 1.3 Cash Moves In and Out of a Mutual Fund: The Fund Trades Securities to Invest Incoming Cash or to Raise Cash for Redemptions

The new investor typically pays no transaction costs. All the shareholders of the fund share any transaction costs associated with investing the new investor's cash in portfolio securities. Similarly, when an investor departs the mutual fund, that investor receives cash equal to the net asset value of the shares when the NAV is next calculated. All the remaining shareholders in the fund bear the cost of selling portfolio securities to provide this liquidity. To the entering or leaving shareholder, liquidity is essentially free. To the ongoing shareholders of the fund, the liquidity given transacting shareholders is costly. Over time, the cost of providing this free liquidity to entering and leaving shareholders is a significant and a perennial drag on the fund's performance. The cost of this free liquidity is increased by the fact that the purchase or sale is usually deemed to occur when the investor's order is delivered to an agent of the fund. The order entry time is often hours or even days before the fund manager actually receives the order and can act on it to buy or sell securities in the fund.

Exhibit 1.4 shows that exchange-traded funds work differently from mutual funds. For most exchange-traded funds, creations and redemptions



Note: All securities transfers are priced at net asset value.
 CBA = Cash Balancing Amount

EXHIBIT 1.4 ETF Creation and Redemption Is In-Kind: Transaction Costs Are Paid by Entering and Leaving Investors

of ETF shares are typically made *in kind*. In a creation, a basket of portfolio securities is deposited with the fund in exchange for fund shares. In a redemption, fund shares are turned in to the fund in exchange for a basket of portfolio securities. We will describe the ETF creation and redemption process in more detail in Chapter 3; but the key feature of this process for the protection of the ETF's ongoing shareholders is that the creating or redeeming entity—in most cases, the portfolio trading desk of a major investment firm acting for a market maker in the ETF shares—is responsible for the costs of investing in the portfolio securities for deposit and the costs of disposing of portfolio securities received in the redemption of outstanding fund shares.⁴ Market makers expect to pass these transaction costs on to investors when the market maker trades fund shares with investors. The cost of entering and leaving a fund varies, depending on the level of fund share trading activity and the nature of the securities in the fund's portfolio. For example, the cost of trading in small-cap stocks can be much greater than the cost of trading in large-cap stocks.

ETFs are different from mutual funds in the way they accommodate shareholder entry and exit in at least two ways: (1) The trading costs associated with ETF shareholder entry and exit are ultimately borne by the entering and exiting investors, not by the fund. (2) An exchange-traded fund does not have to hold cash balances to provide for cash redemptions. An ETF can stay fully invested at all times.⁵ *As a result of these differences, the performance experienced by ongoing shareholders in an ETF should, over time, handily surpass the performance experienced by ongoing shareholders of a conventional mutual fund using the same index or active management investment process.* Ironically, even though the exchange-traded fund was designed to be traded throughout the trading day on an exchange, the ETF is a much better product than a conventional fund for the shareholder *who does not want to trade*. On the other hand, as any mutual fund market timer will tell you, a mutual fund is a better product to trade than an ETF because the shareholders of the mutual fund have traditionally paid the market timer's trading costs.

The conventional mutual fund structure that provides this free liquidity to investors who enter and leave the fund was behind the problems of late trading and market timing that provoked the mutual fund scandals of 2003 and 2004. The SEC has spent a great deal of time and effort trying to

⁴The market makers, or dealers acting on behalf of market makers, even pay a modest creation or redemption fee to cover the fund's administrative expenses associated with creation and redemption of the ETF shares.

⁵We will see in Chapter 6 that lazy ETF portfolio managers don't always keep the fund's cash fully invested.

deal with the problem of market timing trades in mutual funds without eliminating the free liquidity which ongoing shareholders in mutual funds give entering and leaving shareholders. Some fund companies have made a variety of operational “patches” as they attempt to restrict market timing trades. In connection with implementing Rule 22(c)(2), the SEC created a complex and costly reporting structure with almost mandatory redemption fees on mutual fund purchases that are closed out within a week. In the final analysis, the elimination of free liquidity—most easily through the exchange-traded fund in-kind creation and redemption process—is the only way to eliminate market timing without imposing unnecessary costs on all fund investors. Even if there is no such thing as a market timer in the future, long-term investors will fare better in funds that protect them from the costs of other investors entering and leaving the fund.⁶

TAX EFFICIENCY

One of the most frequently discussed advantages of the investment company exchange-traded funds is tax efficiency. The tax efficiency most commonly associated with ETFs is essentially capital gains tax-deferral until the investor chooses to sell the fund shares. Tax deferral in an investment company ETF is a natural consequence of Subchapter M of the Internal Revenue Code which permits fund share redemptions in-kind (delivering portfolio securities to departing fund shareholders) without tax impact inside the fund. Subchapter M mandates that a redemption in-kind initiated by a shareholder does not give rise to a taxable capital gain that would have to be distributed to shareholders of the fund.⁷

This kind of tax efficiency is obviously most important in ETFs that hold common stocks and other securities that can appreciate in value, but some bond funds have capital gains at times. The ability to avoid capital gains distributions even benefits tax-exempt investors because it prevents the build-up of unrealized gains inside an ETF. The build-up of unrealized gains in a mutual fund portfolio can lead to portfolio management decisions that adversely affect tax-exempt shareholders. When the choice facing a portfolio manager is (1) to realize gains on appreciated portfolio securities and distribute taxable capital gains to the fund’s shareholders

⁶For a comprehensive discussion of the 2003–2004 scandals and the market structure behind them, see Gastineau (2004).

⁷For more details on ETF tax treatment, including opportunities for capital gains tax deferral, see the extended discussion of this topic in Chapter 4.

or (2) to hold overvalued securities and avoid realizing capital gains, the portfolio manager faces a conflict between the interests of tax-exempt and taxable investors. This conflict of interest between taxable and tax-exempt investors—inevitable in a conventional mutual fund—disappears in an ETF. Even modest fluctuations in an ETF's shares outstanding from offsetting creations and redemptions give the fund portfolio manager opportunities to deliver the fund's lowest cost holdings of a security in redemptions and gradually increase the fund's average cost basis in each position.

With exchange-traded funds, the decision to change the portfolio can be based solely on investment considerations, not on the tax basis of portfolio securities. The conflict between taxable and tax-exempt shareholders disappears because the achievement of tax efficiency in ETFs is largely a matter of careful designation of tax lots so that the lowest cost lots of a security are distributed in-kind in redemptions and high cost lots are sold to realize losses for the fund when a sale is necessary or appropriate.

Exchange-traded funds grow by exchanging new fund shares for portfolio securities that are deposited with the fund. Redemptions are also largely in-kind. Investors sell their fund shares on the exchange. Dealers buy the fund shares and turn them in to the fund in exchange for portfolio securities. This process lets ETF managers take full advantage of the redemption in-kind provision of the Internal Revenue Code by delivering their lowest cost tax lots without realizing gains that must be distributed to the fund's shareholders. The rules for ETF redemption permit the fund manager to remove a high-cost tax lot from the redemption basket and sell it for cash to realize losses inside the fund.

The early developers of exchange-traded funds were aware of this tax treatment and its ability to defer capital gains taxes for fund investors, but the tax efficiency it gives ETFs was by no means a significant objective in the early development of exchange-traded funds. It is largely serendipitous that most well-managed investment company exchange-traded funds don't distribute taxable capital gains to their shareholders.⁸ Creation and redemption in-kind not only transfers the cost of entering and leaving the fund to the shareholders who enter and leave, it can also help defer capital gains taxes until a shareholder chooses to sell the fund shares.⁹

⁸Later chapters offer a few examples of ETFs that can be much less tax-efficient.

⁹Interestingly, tax deferral helps encourage "shareholder loyalty" to an ETF. An investor in a mutual fund will usually receive taxable gains distributions that, if reinvested, increase his basis in some of the fund shares as their value increases over time. When he sells the higher basis mutual fund shares, the higher basis reduces the capital gains tax on the sale. An investor in an investment company ETF with securities as its principal portfolio holdings should almost never receive a capital

The in-kind ETF creation/redemption process is an efficient, even elegant, solution to several of the obvious problems that continue to plague the mutual fund industry. A growing number of fund industry observers believe that the exchange-traded fund structure will eventually replace conventional mutual funds. To make that happen, however, the serendipity of early ETF development needs to be harnessed through creative financial engineering to overcome weaknesses in the index ETF structure and extend the best ETF features to a wider range of portfolios. We need to innovate selectively and constructively. To extend the advantages of ETFs to a broader range of investment instruments we need to look at some of the other features of ETFs in more detail. Two important topics that are often misunderstood about ETFs are *transparency* and *intraday trading*. Both can be very useful, but it is possible to have too much of what is usually a good thing.

COST TRANSPARENCY IS DESIRABLE, BUT TRADING TRANSPARENCY IS COSTLY

ETFs are appropriately praised for their cost transparency and inappropriately praised for their portfolio transparency. Every investor should be able to identify all of a fund's costs. However, if any fund is going to serve the interests of its shareholders, the portfolio manager needs to implement portfolio changes without revealing the fund's trading plans until *after* trading to implement a portfolio composition change has been completed.

Under current regulations, all investment companies registered with the SEC must reveal their portfolio contents quarterly with a 60-day lag. This means that the holdings of a mutual fund at the end of the March quarter must be published by the end of May. The purpose of this rule is to provide investors with *complete* information on the contents of the fund portfolio at regular intervals. The 60-day lag protects the portfolio manager's ability to trade without revealing the trades until they have been completed. If an investor really wants to know the precise content of a fund portfolio and believes that this knowledge is important for some purpose, there are a lot of

gains distribution. Consequently, the basis of all the ETF shares will stay at the investor's original cost. The tax due on sale of the ETF shares will tend to be greater than the tax due on sale or redemption of an otherwise comparable mutual fund position. An investor with both mutual funds and ETFs can usually defer more taxes by selling mutual fund shares first when he needs money for living expenses. The ETF shareholder will be a more loyal shareholder simply because he wants to continue to defer his tax liability.

benchmark index ETFs and mutual funds that provide daily portfolio transparency. However, the evidence is overwhelming that trading transparency enables scalpers to front-run fund transactions—whether the transactions are in index funds or in actively managed funds. This kind of transparency is contrary to the best interests of the fund’s investors.

Trading transparency can be very costly to a fund’s investors. Whether a fund is attempting to replicate an index or to follow an active portfolio selection or allocation process, portfolio composition changes cannot be made efficiently if “the market” knows in advance what changes a fund will make in its portfolio. A number of recent studies have highlighted the costs of index composition changes. Benchmark indexes like the S&P 500 and the Russell 2000 do not make efficient portfolio templates for long-term investors. Investors in index funds based on these and other popular, transparent indexes are disadvantaged by the fact that anyone who cares will know what changes the fund must make before the fund’s portfolio manager can make them.¹⁰ When transparency means that someone can earn an arbitrage-type profit at the expense of a fund’s shareholders by front-running a fund’s trades, transparency is not desirable.

The cost to ongoing shareholders of preannounced portfolio composition changes in index ETFs should be eliminated whenever possible. The best way to improve index fund performance is to use “silent” indexes. Silent indexes keep portfolio composition changes confidential until after the fund has traded. This requires new procedures for the management of indexes and for the management of index funds. A similar procedure that protects trading confidentiality can be used for actively managed exchange-traded funds. Nearly everyone seems to agree that actively managed funds require confidential treatment of portfolio composition changes until after the fund has traded. Only recently have investors and regulators begun to understand the costs that index transparency imposes on index fund investors. Making portfolio changes confidentially and efficiently requires some changes in the ETF creation/redemption process and in ETF trading. We will discuss these issues in greater depth in later chapters.

INTRADAY ETF TRADING

Intraday trading in ETFs is useful to many investors and traders and there is no question that hyperactive trading in some ETFs has helped call attention

¹⁰This problem is discussed at length in Chen, Noronha, and Singal (2006), Gastineau (2002a), Gastineau (2002d), Gastineau (2005), and Gastineau (2008). The major issues are also discussed in Chapter 5.

to these funds. The fact that some of the most actively traded ETFs make regular appearances on the daily most active stock trading list has stimulated interest in the funds. Active trading has played a major role in the adoption of ETFs for some short- and intermediate-term risk management applications. However, large traders often have an intraday trading advantage over individual investors, particularly individual investors who use less actively traded ETFs. There is asymmetry in the amount and kind of market information available to large traders on one hand and small investors on the other hand. Any trading information asymmetry is unimportant when an ETF trades, say, 10,000,000 shares per day. The bid/asked spread on such actively traded funds is usually just a penny or two. Information asymmetry is *very* important when an ETF is not actively traded and the bid/asked spread is wide. We will discuss NAV-based ETF trading as an antidote to this asymmetry in trading information in Chapter 8.

Many individual investors have a stake in being able to make small, periodic purchases or sales in their fund share accounts. The prototypical investor of this type is the 401(k) investor who invests a small amount in her defined contribution retirement plan every payroll period. The mutual fund industry has developed elaborate procedures that permit small orders for a large number of investors to be aggregated and for cash to enter or leave the fund to accommodate small investors at net asset value. There are ways to adapt ETF procedures so that these small periodic purchasers, while paying a little more than they have paid in the past to cover the transaction costs of their entry and exit, will still be accommodated at low cost. The snowballing push for greater transparency in the costs of defined contribution accounts like 401(k) plans will make fund cost and performance comparisons easier—to the advantage of ETFs and the investors who use them.

COMPARING ETF AND MUTUAL FUND ECONOMICS

Exhibit 1.5 provides an economic comparison of ETFs and mutual funds with the advantages of the ETF cost structure measured in terms of improved investment performance for fund shareholders. In the first column, the possible ETF advantage is listed. The information in parentheses in that column is an estimate of the *range of improved annual investment performance a long-term shareholder who uses an ETF rather than a mutual fund will enjoy*. As these numbers indicate, the advantage of an ETF over a comparable mutual fund can vary over a wide range. Of course, in some cases there is no material difference between ETFs and mutual funds, but sometimes both types of funds would benefit if the fund industry changed current practices.

EXHIBIT 1.5 Using ETFs to Deliver Better Investor Performance

ETF Advantages	Possible ETF Problems	Solutions
Shareholder protection (<0.1% to >5.0%)	Uncertain transaction costs. Fairness of execution.	Net asset value (NAV)-based trading process improves on basic ETF shareholder protection pp. 197-225
Lower operating costs/cost transparency (0 to 0.35%)	Under the alternative minimum tax (AMT), embedded costs cover fees.	New fund delivery structures pp. 272-273
Capital gains tax efficiency (0 to 2.5%)	None	None needed
Taxable/tax-exempt conflict (0 to 1.0%)	None	None needed
Indexing (equal)	Inefficient indexes hurt all investors: The more popular the index, the greater the performance drag from index transparency.	Silent indexes as portfolio templates pp. 110-126
Active management (equal)	Confidentiality in portfolio changes is essential to all investors.	Same portfolio composition disclosure as mutual funds pp. 187-195

Column two in Exhibit 1.5 lists some possible problems with today's ETF or mutual fund structure and column three offers solutions for implementation, in some instances in a new generation of ETFs. In a few cases (such as the need for more efficient indexes), the silent index solution is equally applicable to conventional mutual funds that follow an indexing strategy. Common sense leads to the conclusion that it is not in any fund investor's interest to pay significant index change transaction costs that the fund incurs because its index is transparent.

Each of these advantages, problems, and solutions merits a more detailed discussion than is possible in this introduction. The page numbers in the solutions column indicate where you will find detailed discussions in the chapters that follow. However, a few of the features that are part of most ETFs merit at least a brief discussion at this point.

The first ETF advantage reflects the value of shareholder protection from the cost of investors entering and leaving a fund as discussed in connection

with Exhibits 1.3 and 1.4 earlier in this chapter. The return comparison in parentheses removes the allocation of all entry and exit costs from ongoing mutual fund shareholders and assigns them to entering and leaving shareholders who use ETFs. *In an ETF transaction, a shareholder pays only the cost of his own entry to and exit from the fund. The mutual fund shareholder pays a pro rata share of the entry and exit costs of all fund buyers and sellers for as long as he owns the fund shares.* The transaction costs don't disappear, but the long-term investor in an ETF pays them only when he personally transacts, not every time any investor enters or leaves the fund.

There have been few appropriately designed studies of the shareholder performance cost of the *flow* of cash into and out of mutual funds. In a study published in 1999, Roger Edelen, then a professor at Wharton, measured the cost of flow for a sample of 166 equity and hybrid mutual funds using data from 1985 through 1990.¹¹ He calculated the cost of flow in terms of its adverse effect on fund shareholder performance at 1.43 percent (143 basis points) per year in the average fund in his sample. The shareholder turnover in the sampled funds was low enough that it is clear that market timing and late trading was not a significant factor in the cost of flow to these funds' shareholders. Shareholder turnover in most large mutual funds is lower today than it was in Edelen's sample. Some transaction costs associated with accommodating flow are also probably lower today.

In a more recent paper, Edelen, Evans, and Kadlec (2007) examined the cost of flow in a larger sample of more recent mutual fund data. They found the average annual cost of flow to be .75 percent (75 basis points), partly because the average mutual fund shareholder stays in the fund longer than he did 10 years earlier. If the cost of flow for the average mutual fund investor (not the average mutual fund) is .75 percent per year for the \$5.6 trillion in equity and hybrid mutual funds in the United States at the end of 2009, this represents a performance loss to investors of \$42 billion per year. If the average cost of flow is as low as 0.50 percent per year, the cost to investors is still \$28 billion per year. This lost performance dwarfs the costs that have been attributed to mutual-fund market-timing transactions under any reasonable assumptions.

Note the wide range we use for the value of shareholder protection from the cost of flow (less than 0.1 percent to more than 5.0 percent per year) in Exhibit 1.5. The less than one-tenth of 1 percent number is representative of some very big, large-cap mutual funds with very low shareholder turnover.

¹¹See Edelen (1999). For a more detailed discussion of this paper, see Gastineau (2004).

The more-than 5 percent annual cost figure applies to some small-cap funds with high shareholder turnover. Clearly, the cost of accommodating market timers and late traders in some funds implicated in the 2003–2004 “scandals” was well in excess of 5 percent per year.¹² There is some cost disadvantage to a mutual fund’s ongoing shareholders relative to an ETF when there is *any* flow. Most ETFs eliminate these costs completely for a fund’s ongoing shareholders with in-kind creation and redemption of their fund shares.¹³

The only “problem” that limits the ability of ETFs to deliver this degree of shareholder protection is that the true transaction costs associated with buying and selling shares of an ETF can be difficult for an investor to determine in advance of an intraday trade. The information available to investors on the intraday value of an ETF and the total cost of intraday trading is not as useful or as readily available as it should be. Calculations of intraday fund portfolio values are made and disseminated, but many investors do not have easy access to the every 15-second net asset value proxy calculations for existing ETFs. In any event, as we will see in Chapter 8, these intraday value proxies are of little value to traders. Furthermore, intraday ETF trading execution costs are difficult or impossible to measure accurately. In many instances intraday trading is much more costly than a new way to trade ETFs that is described in Chapter 8. The solution to the cost and implementation problems that intraday trading creates for some investors is an NAV-based trading process that increases the transparency of ETF transaction costs and, consequently, improves ETF structural shareholder protection without compromising the ETF “gold standard” whereby

¹²The analyses made in connection with financial settlements paid by parties associated with the 2003–2004 market-timing scandals reveal that market timing was practiced by many fund share traders who did not have formal or informal arrangements with fund managers or distributors. In at least one case, “nonarrangement timing” accounted for more shareholder costs than arrangement timing. Furthermore, these analyses document some of the trading and dilution costs Edelen (1999) and Edelen, Evans, and Kadlec (2007) found in connection with ordinary fund share purchases and redemption transactions. See Anand (2006) and U.S. Securities and Exchange Commission File No. 3-11814 (2007), especially Sections III–V for more detail on the costs of market timing to mutual fund investors. One of the most comprehensive discussions of the impact *any* purchase or sale of mutual fund shares has on the broadly defined transaction costs, opportunity costs, and dilution experienced by ongoing fund shareholders can be found in Greene and Hodges (2002).

¹³A small number of mutual funds levy asset-based fees on entering and/or redeeming shareholders in an attempt to protect ongoing shareholders from the approximate cost of flow. This approach is better than ignoring the problem, but such fees cannot capture changes in costs over time or relative to trade size.

investors entering and leaving the fund pay the costs of their own (and only their own) entry and exit.

The second advantage of exchange-traded funds listed in Exhibit 1.5 is that ETFs frequently offer lower operating costs and greater cost transparency than conventional mutual funds. Some of the reduction in operating costs and increase in cost transparency is associated with the elimination of costs associated with shareholder accounting at the fund level. Some of these shareholder accounting costs still have to be borne by someone. The financial intermediary that provides fund share transaction and custody services to the investor may ultimately charge the cost of these services to investors. In addition to the expenses embedded in the fund's expense ratio, sales and advisory charges are often paid outside the fund by ETF investors who use the services of an advisor to select ETF shares.

Unbundling costs can create a problem for taxable investors—particularly for investors who are subject to the alternative minimum tax (AMT). The embedded costs of mutual funds, because they are often taken out before the fund's income distributions are made, are deducted from the income that taxable investors receive. A separately billed advisory fee is not fully deductible to the average U.S. taxpayer and may not be deductible at all to an investor who falls under the alternative minimum tax regime. There can be significant tax savings for many taxable investors if advisory and administrative costs and sales charges are embedded in the financial instrument, rather than billed as separate fees. The solution to this tax problem is a new fund delivery structure that provides tax efficiency by re-embedding some of the costs that have been taken out of exchange-traded funds. Embedding costs in a fund can serve an investor's interests as long as the re-embedding does not reduce cost transparency or increase after-tax costs.

Returning to the next items listed in Exhibit 1.5, capital gains tax deferral and taxable/tax-exempt investor conflict of interest elimination are unmitigated gains for all ETF investors. There are no obstacles to realizing these advantages, so no solution is necessary. These important gains flow to ETF investors automatically.

With respect to the last two issues listed in Exhibit 1.5, performance penalties associated with transparency in indexing and the need for confidentiality of an active manager's trading activity, the solutions for the two fund structures are essentially identical: Eliminate portfolio composition trading transparency. Index funds should be based on efficient indexes. There are some very efficient published indexes available today. An outstanding example of an efficient broad-market index is the Wilshire 5000.¹⁴

¹⁴See Gastineau (2006b).

Even though its changes are transparent, the Wilshire 5000 is more efficient than most indexes because it is designed to have very few composition changes. Some inherently inefficient indexes are used for such small asset pools that scalpers who know what the ETF has to do to match the published composition changes in its index are not likely to try to front-run the fund's portfolio trades.¹⁵

Even if trading transparency is not always costly, there is no reason why the index templates for most index funds should not be silent indexes. All investors should have the opportunity to buy index funds based on silent indexes to protect themselves from the cost of index composition front-running trades.

In most discussions of actively managed ETFs, there has been appropriate concern expressed for the cost of achieving enough portfolio transparency to facilitate trading in ETFs without subjecting the fund's trades to the front-running risk that all of today's index funds experience to some degree. The SEC's Concept Release on actively managed ETFs stressed the importance of finding a solution to this problem.¹⁶ While solutions to this problem are addressed at greater length in Chapters 7 and 8, a full consideration of this topic is outside the scope of the present volume because many aspects of full-featured active management of ETFs do not yet have formal SEC approval. In brief, however, a non-transparent, actively managed ETF can trade efficiently if the manager offers no more information on the fund's portfolio composition than the manager of a conventional mutual fund must publish today. With net asset value based trading, there is no need for disclosure of portfolio changes until after the changes are completed. The level of disclosure required of mutual funds is sufficient to support efficient ETF trades at or near each day's net asset value. Funds that do not require the full measure of confidentiality available under today's rules for fund asset disclosure can reduce transaction costs for their entering and leaving shareholders and for market makers by providing more frequent portfolio disclosure.¹⁷ But more frequent disclosure is not essential. An investment process that requires the maximum permitted portfolio confidentiality can work well for an actively managed ETF.

¹⁵If one of these funds grows in response to unusually good performance, the manager may face the same kind of front-running costs that S&P 500 and Russell 2000 index funds experience regularly.

¹⁶United States Securities and Exchange Commission, (2001), www.sec.gov/rules/concept/ic-25258.htm.

¹⁷Many funds already publish their portfolios more frequently and/or with a shorter lag than the SEC requirement.

CONCLUSION

Fund issuers can build on the compelling advantages of today's exchange-traded funds to offer better and more varied portfolios. Nontransparent portfolios in actively managed and silent index funds can offer shareholders the protection from the cost of entry and exit by other fund shareholders and the tax efficiency that are inherent in the initial generation of SPDR-style exchange-traded funds. Active management and silent indexing require a slightly modified ETF structure and an improved trading mechanism for investors who buy and sell ETF shares. The new ETFs will also offer alternative fund share trading and delivery structures and systems to make ETFs useful to all investors. Transparent index funds will be challenged by silent index funds that provide improved performance as a result of lower fund portfolio transaction costs. Actively managed ETFs will feature flexibility in portfolio disclosure to permit the fund manager to determine the degree of transparency that is appropriate for a specific fund, within established regulatory constraints. These and other changes are coming and investors will benefit substantially from the new opportunities these changes provide.

In expressing confidence in the desirability—and the inevitability—of the improved exchange-traded funds described throughout this book, I am well aware of the obstacles facing innovators in the financial services industry. Professor John Y. Campbell, in his 2006 presidential address to the American Finance Association, addressed this issue:

I suggest that retail financial innovation is slowed by the cost of advertising and educating households, together with the weakness of patent protection for financial products. . . . I speculate that the existence of naive households permits an equilibrium . . . in which confusing financial products generate a cross-subsidy from naive to sophisticated households, and in which no market participant has an incentive to eliminate this cross-subsidy. . . . It may be difficult for new investment products to gain acceptance if sophisticated households, who are the natural early adopters, must give up the benefit of a cross-subsidy when they move from an existing product to a new product.

Campbell raised some important concerns, but there is every reason to believe that the most important ETF innovations will succeed. Mutual funds subsidize the *fund share trading costs* of short-term investors (market timers and all other mutual fund share traders), small investors (young investors and others with few assets), and investors who invest small amounts

periodically (largely owners of 401[k] and similar defined contribution retirement accounts). These trading-cost subsidies come at the expense of some of the most “sophisticated households” that hold mutual fund shares as long-term investments. The regulatory interest in thwarting mutual fund timers and traders is well known. Small investors and 401(k) contributors tend to be long-term investors. They will pay a transaction cost to buy and sell ETF shares, but new delivery mechanisms should minimize this cost and clarify the total ETF cost and performance advantage. The incentives for all long-term investors and regulators to embrace the ETF fund structure are compelling.

<http://www.pbookshop.com>