Preliminary and incomplete.
Please do not cite without permission.

Search Costs, Product Differentiation, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds*

Ali Hortaçsu
University of Chicago
hortacsu@uchicago.edu

and

Chad Syverson
University of Chicago
syverson@uchicago.edu

November 2002

* Zvi Eckstein provided thoughtful suggestions regarding an earlier draft. We have also benefited from discussions with Lars Hansen, Boyan Jovanovic, Bob Lucas, Derek Neal, and seminar participants at INSEAD, NYU, the University of Chicago, the Society for Economic Dynamics annual meeting, and the Economic Research Center. Arie Toporovsky, Chris Xu, and Yoo-Na Youm provided research assistance. We are grateful for the financial support of the John M. Olin Foundation. Correspondence can be addressed to either of the authors at Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637.

Abstract

Two salient features of the competitive structure of the U.S. mutual fund industry are the large number of funds and the sizeable dispersion in the fees funds charge investors, even within narrow asset classes. Portfolio financial performance differences alone do not seem able to fully explain these features. We focus on the retail S&P 500 index funds sector, where we find similar patterns of fund proliferation and price dispersion. This suggests that costly investor search and non-portfolio fund differentiation may play an important role in the mutual fund industry. To quantify the welfare impact of such factors in the market for mutual funds, we construct a model of industry equilibrium in which consumers conduct costly search over differentiated products, and estimate the structural parameters of our model using panel data on fund fees and market shares in the retail S&P 500 index funds sector. Our results indicate that fairly small search costs can explain the considerable price dispersion in the sector, and that consumers value funds' observable non-portfolio attributes—such as fund age and the number of other funds in the same fund family—in largely plausible ways. Finally, we investigate the possibility that there are too many funds in the sector from a social welfare standpoint. We quantify the welfare impact of a counterfactual sector structure where entry is restricted to a single fund. We find that restricting entry would yield nontrivial gains from reduced search costs and productivity gains from scale economies, but these may be counterbalanced by sizeable losses from monopoly market power and reduced product variety.

I. Introduction

An investor seeking to hold assets in a mutual fund is a consumer with many choices: in 2001, there were 8307 U.S. mutual funds in operation. If we count different share classes for a common portfolio as separate options available to an investor, the implied total number of funds to choose from tops 13,000. Note as a comparison that there were a total of 7600 hundred companies listed that year on the NYSE, AMEX, and Nasdaq combined. The choice set of the mutual fund investor has also been growing robustly over time: there were 834 mutual funds in operation in 1980. This nearly quadrupled to 3100 by 1990, and almost tripled again by 2001.[1]

An additional, less documented feature of the mutual fund marketplace is that enormous dispersion in the fees investors pay to hold assets in funds persists, despite the large number of industry firms. This dispersion is not simply a result of fee differences across fund sectors; price dispersion within (even narrowly defined) sectors is very large. Table 1 summarizes this within-sector dispersion. The table shows fund fee dispersion moments—the coefficient of variation, the interquartile price ratio, and the ratio of the 90th to the 10th percentile price—for each of 22 fund objective sectors in 2000.[2]

As is evident in the table, the 75th-percentile-price fund in a sector typically has investor costs two to three times that of the fund at the 25th percentile. The 90th-10th percentile price ratios range from roughly 3.4 to nearly 10. The extrema of the distribution (not shown) can exhibit vast dispersion; the minimum-price aggressive growth fund, for example, imposed annualized fees of only 14 basis points (i.e., 0.14% of the value of an investor's assets in the fund), whereas the highest-price fund charged a whopping 1670 basis points. While some of these sectors are fairly broad and may include funds with very different portfolios, dispersion remains even within what are plausibly quite specialized sectors. The table's second panel shows the same dispersion measures for four randomly selected specialized sectors (these are

[1] The expansion of the choice set has accompanied a steady increase in the fraction of the population taking advantage of the mutual fund option. Only 6% of households held mutual funds in 1980, with $134 billion in assets in 12 million individual accounts. By 2001, though, fully 52% of U.S. households held assets in mutual funds.

[2] Throughout the paper, we refer to funds' annualized fees as "prices" to reflect the fact that from the investor's perspective, these fees are the cost of buying the right to hold assets in a fund. The fund data used to calculate these figures is from the Center for Research in Security Prices (CRSP). Roughly 95 percent of the funds in the CRSP database are matched to one of these sectors, which are categorized according to the Investment Company Data, Inc. (now Standard and Poor's Micropal) system. Fund prices are annualized investor costs, calculated according to the literature standard of annual fees (both management fees and 12b-1 fees if applicable) plus one-seventh of any loads. The paper's data appendix discusses the data in greater detail.

based on Strategic Insight's classification system, which is considerably finer than the above taxonomy). There is only a modest drop in price dispersion compared to the broader sectors.

Of course, fund portfolios can vary considerably even within narrow asset classes. Fund price dispersion might then reflect within-sector differences in demand or cost structures across portfolios. On the demand side, certain portfolios will outperform their sector cohorts; higher prices may just reflect investors' willingness to pay for better performance. As for costs, fund managers choose different securities with which to comprise their portfolio, and some of these securities may be more expensive to analyze or trade than others. Fund prices may reflect this fact.[3] Portfolio differentiation, too, may explain in part the other salient fact discussed above: the large number of industry funds. Investors may differ in their ideal portfolios and their current asset compositions. Perhaps thousands of funds (and the several hundred new funds each year) are necessary to meet the demand for the many risk-return profiles sought by investors.[4]

However, a look at the retail (i.e., non-institutional) S&P 500 index fund sector strongly suggests that the composition and financial performance of funds' portfolios are not the only factors explaining fund proliferation and fee dispersion. All funds in this $164 billion (in 2000) sector explicitly seek to mimic the same performance profile, that of the S&P 500 index. Thus any discrepancies among these funds' financial characteristics should be minimal, and the

[3] To check whether price dispersion plausibly reflects performance differences across funds, we regressed, for all mutual funds in the CRSP database between 1992 and 2000 (roughly 73,000 fund-year observations), both the gross annual returns and the within-year standard deviation of monthly returns on fund prices. We also include interacted Strategic Insight objective category (of which there are 193) and year effects in the regressions (hence estimated coefficients reflect the correlation between returns and prices within sector-year cells). The price coefficient in the gross annual return regression is actually negative: -0.610 (s.e. 0.218); more expensive funds have lower than average returns. Furthermore, the correlation between fund prices and the standard deviation of monthly returns within the year is positive: 0.722 (s.e. 0.536)—also the opposite sign one would expect if performance and price were closely linked. A more careful investigation would obtain measures of expected returns and use longer performance histories to measure within-fund return variation; however, given the magnitude of the observed price dispersion, these results suggest that the price-performance link is not an overwhelming determinant of the observed patterns in the data. The findings are also in line with the results in Carhart (1997), which suggest the impact of expenses on performance is negative and at least one-for-one.

[4] The competitive effect of entry appears to weak throughout the industry. We regress sectoral price (fee) dispersion moments on the number of sector funds, allowing for sector-specific trends. The results (available from the authors upon request) show that within-sector price dispersion increases with entry; both the unweighted and weighted standard deviations increase significantly as the number of funds in the sector rises. Perhaps surprisingly, the effect of entry on the central tendency of the sector price distribution is weak. Running the same specification with both unweighted and weighted mean prices as dependent variables yields coefficient estimates that are insignificant (the unweighted mean coefficient is positive and the weighted mean estimate negative). So while entry into a sector serves to further spread the observed distribution of prices, there is little evidence of a competitive effect lowering the average market price.

observed competitive structures of the sector (including the important fund proliferation and price dispersion issues discussed above) are likely to be driven by non-portfolio effects.

It is readily apparent that, despite the sector's financial homogeneity, the features of the broader mutual fund industry are equally prominent. There were 85 retail S&P 500 index funds operating in 2000, even though each one offered plausibly equivalent expected risk-return profiles, a number that seems well beyond the saturation point arising from simple portfolio choice motives. Entry has been brisk too: the number of funds in the sector has more than quintupled since 1992.[5] As for price dispersion, the highest-price S&P 500 index fund in 2000 imposed annualized investor fees nearly 40 times as great as those of the lowest-cost fund: 368 vs. 9.5 basis points. This striking divergence is not restricted to the far ends of the distribution; the 25^{th} and 75^{th} percentile prices are 47 and 169 basis points, respectively—an over three-and-one-half-fold difference, and as high as any of the broader sectors outlined in Table 1. The 90^{th}-10^{th} percentile ratio is also a comparatively high 10.8. Even more interestingly, high-price funds are not all trivially small. The highest-fee fund held 1.1 percent of sector assets—enough to make it the tenth-largest fund in the sector and not much smaller than the 1.4-percent share of the lowest-price fund.[6]

That so many funds, with such diffuse prices, operate even in a sector of plausibly financially homogeneous funds is a puzzle we seek to address. We investigate whether two sources of price dispersion are at work in the retail S&P 500 index fund sector: consumers' taste for product attributes other than portfolio composition, and/or informational (or search) frictions that deter investors from finding the fund charging the lowest management fee. Below, we model a competitive equilibrium that simultaneously incorporates these two elements. We derive equilibrium conditions and evaluate the ability of search and non-portfolio differentiation to qualitatively and quantitatively explain patterns in the data. A favorable feature of our model

---

[5] Throughout this paper, we follow the CRSP convention (also common in the literature) of treating each fund share class in multi-class funds as a separate fund. Multi-class funds are those which have a common manager and portfolio, but have different pricing schemes and asset purchase and redemption rules. To give an idea of the relative prevalence of multiple-share-class portfolios in the sector, if we were to count multiple-class funds as single funds, there would be about 50 funds in 2000. We also include exchange-traded funds (ETFs) based on the S&P 500 index. ETFs are essentially index fund portfolios that can be traded in a stock market. There is only one ETF for most of our sample, Standard & Poor's Depositary Receipts (SPDRs). In 2000, a second ETF started trading, Barclay's iShares S&P 500 Index Fund.

[6] The sizeable price dispersion is not driven simply by loads; similar spreads are observed just among the annual fees (the sum of management and 12b-1 fees) charged by the funds. The comparable dispersion measures for annual fees are as follows: 75^{th}-25^{th} percentile price ratio = 3.2; 90^{th}-10^{th} percentile ratio = 9.4; and max-min ratio = 31.4.

is that it nests as special cases the more commonly applied search-over-homogenous-products and differentiated-products-with-no-search models. This allows us to gauge the ability of both search and product differentiation in isolation to explain industry outcomes, and to measure their relative magnitudes when operating simultaneously. We then apply the model's estimates to measure the social welfare implications of having a large number of funds with diverse prices in a competitive structure driven by search over differentiated products.

We find that fairly small search costs can explain the considerable price dispersion in the sector. Furthermore, the distribution of search costs across investors has been shifting over time. Search costs were falling in the lower three quartiles of the distribution through the late 1990s. Conversely, costs at the high end of the distribution were rising during the same period. We speculate that this is a result of a composition shift: the documented influx of novice (and high-information-cost) mutual fund investors into the industry during the period. We also find that consumers appear to value funds' observable non-portfolio attributes, such as fund age, the number of other funds in the same fund family, and tax exposure, in largely expected ways. In our welfare calculations, we find that restricting entry into the sector to a single fund would yield nontrivial gains from reduced search costs and productivity gains from returns to scale. However, these gains may well be counterbalanced by losses from monopoly market power and reduced product variety. We believe that our results are consistent with search and non-portfolio differentiation playing an important role in the retail S&P 500 index fund sector. They also suggest that these two features may help explain observations about the industry as a whole.

The mutual fund industry's sheer size and ubiquity—over half of U.S. households have mutual fund holdings, and various institutional fiduciaries hold over $3 trillion in assets, in 250 million mutual fund accounts—is responsible in part for the considerable research interest that mutual funds have attracted. The primary thrusts of previous research regarded the construction of mutual fund portfolios, evaluation of fund performance vis-à-vis the theoretical efficient-portfolio benchmark, and the agency relationship between mutual fund shareholders and the fund manager.[7] These research lines have produced very important normative and positive insights into the operation of the mutual fund industry. However, this research focuses on mutual fund company operation in isolation from other market firms; that is, the research does not delve

[7] This literature, too vast to cite comprehensively, follows from the classic contributions on portfolio theory and empirical testing as well as the principle-agent framework. See, for example, Jensen (1968), Malkiel (1995), Falkenstein (1996), Gruber (1996), and Chevalier and Ellison (1997, 2000).

greatly into the competitive structure of this industry. This paper seeks to partially fill this gap. Since competitive forces directly affect the fortunes of individual funds and fund families as well as consumer welfare, we believe that better understanding of demand and supply determinants in this important industry is a very fruitful area of research.[8]

An important additional contribution of this paper is methodological: the empirical modeling and estimation framework developed here can be applied to many other industries in which search frictions co-exist with product differentiation. Unlike previous empirical applications of models of search equilibrium (mostly in labor economics), we do not have data on individual consumers' (workers in the labor context) decisions, and have to draw inferences regarding search behavior from aggregate price and quantity data.[9] In this respect, our methodology is closest to the work of Hong and Shum (2001), who discuss identification and estimation of equilibrium search models using only market-level data on prices. Our model extends their results considerably by showing that aggregate quantity data, when combined with price data, allows one to estimate a much richer model in which products are vertically differentiated. As we will note below, there are also close ties between our search framework and discrete-choice demand models of McFadden (1981) and Berry, Levinsohn, Pakes (1995).

We proceed as follows. The following section provides more detail regarding the competitive patterns in the retail S&P 500 index fund sector. In Section III, we construct a model of consumer search over differentiated funds and the optimal pricing responses of mutual fund companies. We go on to estimate the obtained equilibrium equations in Section IV. This is done under several different sets of assumptions regarding the relative importance of search and product differentiation, and the results from each are compared. We then test the estimates for robustness to particular empirical modeling assumptions. In Section V, we estimate welfare implications of limiting entry into the sector. Section VI concludes.

---

[8] We should note that there has been a recent increase of interest among financial economists in the competitive structure of the mutual fund industry. For example, Sirri and Tufano (1998) note that mutual fund flows are relatively insensitive to management fees and excessively sensitive to past performance (as opposed to expected performance). Khorana and Servaes (1999, 2001) empirically explore what market characteristics are correlated with fund entry and the market share of fund families. Barber, Odean, and Zheng (2001) find that fund inflows seem to be more responsive to some price instruments than others, and that advertising appears to be an effective tool for increasing fund assets. On the theoretical side, Massa (2000) and Mamaysky and Spiegel (2001) explore the driving forces of fund creation.

[9] Sorensen (2000) is the only attempt that we are aware of in the industrial organization literature to estimate search costs using consumer-level product choice data. Examples of structural estimation of search models in labor markets include Flynn and Heckman (1982), Eckstein and Wolpin (1991), Ridder and van der Berg (1998), Bontemps, Robin and van der Berg (2001).

II. An Overview of the Retail S&P 500 Index Funds Sector

Index mutual funds hold portfolios tied to a particular market index; S&P 500 index funds are the most popular type of these funds. Index funds are passively managed; fund managers do not actively choose the stocks that compose the fund's portfolio. Instead, managers attempt to mimic the return patterns of an index, the S&P 500 in this case. One way for the manager to accomplish this is to hold index equities at the same proportions as their weights in the index. Other (high-tech) index-matching methods can also be used, especially for those funds that track very broad indices with a large number of component stocks.

The fact that all funds in the S&P 500 index sector mimic a common return pattern is useful. Restricting our attention to funds which are plausibly (ex-ante, at least) financially homogeneous allows us to isolate the roles of search and (non-portfolio) fund differentiation.[10] The evidence does indeed suggest considerable financial performance homogeneity in our sample. Table 2 contains summary statistics of our sample funds' annual returns as well as the yearly average and standard deviations of their monthly returns.[11] As can be seen, the dispersion in return patterns is quite small compared to their averages. The interquartile range of annual returns is no greater than 0.65 percentage points, and typically about 3% of the median value.[12] The dispersion of average monthly returns across funds is similarly small: the maximum interquartile range is 0.052 percentage points and the average interquartile range-median ratio is 0.03. The same ratio for the standard deviation of monthly returns are typically less than 0.01,

[10] We exclude institutional S&P 500 funds from our sample, despite the fact that they also mimic the same performance profile, because we believe they operate in a fundamentally different product market than non-institutional funds. Institutional funds characterized by low fees and very high initial minimum investment levels (typically $1 million) and are specifically marketed to institutional and very-high-wealth clients. Thus there are likely to be large differences in customer characteristics and behavior between institutional and retail funds. By excluding institutional funds, we hope to further control for unobservable differences across funds that might confound our analysis.

[11] The statistics in the table correspond only to funds operating every month in the given year of observation. This eliminates any composition bias from funds with return data spanning only a possibly non-representative portion of the year.

[12] While these patterns are largely mirrored in the standard deviations of the return moments, outliers do enlarge the standard deviation of annual returns in some years, to as much as 13% of the mean in 2000. An index fund's performance can deviate from the index which it is tracking (and from other funds tracking the same index) because of several factors. These include idiosyncratic portfolio sales required to meet the particular daily activity needs of a fund, how much of the fund's assets are held in cash, and the timing of trades.

and the coefficient of variation is under 5% in every year. This small variation suggests ex-ante return homogeneity among the funds.

As we document above, despite this homogeneity in financial performance, the sector shares the fund proliferation and price dispersion traits seen in the broader mutual fund industry. The sector has also seen robust entry, and despite the large number of new funds in the sector, both average prices and price dispersion have increased. This is particularly puzzling given what should be the high substitutability of the funds' portfolios. We model below a competitive equilibrium that allows for the facts observed in the data while still allowing sector funds to be financially homogeneous.

Figure 1 plots the vigorous growth in both the number of retail S&P 500 index funds and their total net assets under management from 1995 to 2000. This growth was coincident with the overall growth in the mutual fund industry documented above, but was even stronger. The number of sector funds more than tripled from 1995 to 2000 (from 24 to 85), and sector assets grew over the period at twice the rate as did total equity fund assets.[13]

Figure 2 and Table 3 show the evolution of the sector's price distribution over our sample period. A striking feature of the figure, which plots the cumulative price distribution functions of sector funds, is that despite the entry of nearly a dozen funds a year, the rightward shift in the price distribution was so strong from 1996-1999 that a given year's distribution is almost strictly dominated by the next year's. (This trend interestingly reversed in 2000.) Along with the increase in average price evident in the figure, Table 3 documents a corresponding increase in price dispersion: both the interquartile and the 90^{th}-10^{th} percentile ranges increase markedly throughout the observation period. Moreover, these greater price spreads result even though the lower quantiles increase, because of larger increases in the higher end of the distribution.

A related and interesting facet of the sector's development is the evolution of the relative market shares of the high and low price segments. One can see from Figure 3 that while low price funds do indeed capture a dominant market share, the asset share of funds in the lowest price decile has fallen consistently since 1996. In contrast, the market shares of the upper quartile and decile rose during the same period. The proportional rise in the amount of sector

[13] The growth was driven in part by the growing popularity of passive management strategies among investors. The S&P 500 index fund sector was particularly able to capitalize on this preference shift partly because the original index fund, the First Index Investment Trust, is today the Vanguard 500 Index Fund. Its early entry into a growing sector has been seemingly very important: the Vanguard 500 Index Fund still is a dominant player in the sector.

assets held in top price decile funds has been especially stark, rising from virtually zero to more than 2%. The reallocation of market share to higher-priced funds resulted in a 34% rise in the sector's asset-weighted average price from 1995 to 2000, from 27.1 to 36.4 basis points.

These facts document that the proliferation of sector funds has been accompanied by upward trends in both the dispersion, as is the general case in the industry, and central tendency of their price distribution. Average investor costs have continued to rise despite the significant amount of entry into the index fund marketplace. Indeed, if the asset-weighted mean price was held constant at 1995 levels, total annualized fees would be $152 million less.

Figure 4 sheds additional light on this observation, plotting numbers of entrants in low- and high-price groups. From 1995 through 1999, 25 funds entered the market with prices above 100 basis points. In contrast, only six funds entered the market that charged less than 40 basis points. A decomposition of the asset-weighted average sector price growth, shown in Table 4, quantifies this more specifically by taking into account the funds' market shares. The decomposition—used in another context by Foster, Haltiwanger, and Krizan (2002)—breaks annual changes in the (asset-based) market-share-weighted average price (i.e., $\overline{P}_t = \sum s_{i,t} p_{i,t}$) into five components: within, between with fixed prices, between covariance, entry, and exit.[14] As Table 4 indicates, entry's net price effect is positive for all years except 2000 (note, however, that the share-weighted average price still increased from 1999 to 2000 despite the shift in the unweighted price distribution). Furthermore, for the two years with the largest average share-weighted price increases (1998 and 1999), entry accounted for most of the change. Another component leading to price increases is the increasing market shares of above-average-priced funds. Interestingly, the within component is negative from 1996-1999, indicative of the fact that continuing funds' prices decrease on average.

[14] This decomposition is defined as follows:

$$\Delta\overline{P}_t = \sum_{i\in C} s_{i,t-1}\Delta p_{i,t} + \sum_{i\in C}\Delta s_{it}\left(p_{i,t-1} - \overline{P}_{t-1}\right) + \sum_{i\in C}\Delta s_{it}\Delta p_{i,t} + \sum_{i\in N} s_{i,t}\left(p_{i,t} - \overline{P}_{t-1}\right) - \sum_{i\in X} s_{i,t-1}\left(p_{i,t-1} - \overline{P}_{t-1}\right)$$

where i denotes a particular fund, $s_{i,t}$ and $p_{i,t}$ are the asset share and price of fund i in year t, respectively, C is the set of incumbent funds in t continuing operations from t-1, N is the set of funds that enter in t, and X is the set of funds that exit in t-1. The first term—the "within" component—measures that portion of the price change accounted for by the price changes of continuing funds, holding asset shares constant. When a continuing fund lowers its fees from one year to the next, this term becomes more negative. The second term captures market share reallocation between funds that are priced above or below the previous period's average prices; it is negative when relatively low (high) price funds gain (lose) market shares. The third is another between-component that captures the covariance between price and market share changes among continuing funds. The fourth and fifth terms capture the effect of entrants and exits, respectively; they are positive when entrants (exits) have higher prices than the sector average.

The model of sector equilibrium that we construct below can shed light on many of these facts. First and foremost, it shows how a large number of funds with dispersed prices can be sustained in equilibrium despite financial homogeneity. Furthermore, it reflects a possible explanation for the shift to more expensive funds during the great asset run-up of the 1998-99. It allows us to estimate the shape of the investor search cost distribution at any given point in time, as well as to track the distribution's evolution throughout the sample. The model also allows us to estimate which fund attributes are valued by investors in their purchases.

III.1. Model—General Setup

Demand for sector funds is characterized by a continuum of investors searching over funds with varying attributes. We assume that fund attributes are vertical characteristics, and that all investors share a common utility function. Thus conditional on investing in fund j, an investor receives indirect utility equal to $u_j(W_j;\ \theta)$, where W_j is a vector of fund attributes and θ is a set of parameters that characterize how the attributes affect utility.[15] While the model requires few specific assumptions on the form of $u_j(\cdot)$ a priori, we keep things straightforward by assuming that utility is a linear function of fund characteristics:

$$u_j = \underline{W}_j \beta - p_j + \xi_j, \quad (1)$$

where $\underline{W}_j$ are those elements of W_j other than price p_j and an unobservable component ξ_j. Note that the coefficient on the price term has been normalized to –1, so utilities are expressed in terms of the unit of price measurement. Here, given the nature of the good, that unit is basis points. Thus one can think of u_j as specifying the utility of the fund per dollar of assets the investor holds in it.

Given these preferences, the fund delivering the largest u_j would gain 100% market share if search were costless. However, because search is costly in our framework, market shares are distributed across funds. We assume these search costs are heterogeneous in the investor population and have distribution $G(c)$. Investors must incur their particular search cost whenever

[15] We have explored allowing a more general preference specification that incorporates horizontal taste differences across consumers; i.e., where the conditional indirect utility of fund j is now specific to a particular consumer i: $u_{ji}(W_j;\ \theta_i)$. However, we have not yet been able to demonstrate that market outcomes driven by horizontal differentiation are identifiably separable in our data from those caused by across-consumer variation in search costs. For now, we allow the idiosyncratic portion of consumer purchase behavior to be explained by search, and later compare the ability of this model to fit the data with that of a model without search but with horizontal taste shifters.

they wish to find out the indirect utility offered by a particular fund, with the exception of the first fund they search (this assures all investors desiring to hold assets in the sector end up doing so regardless of their search cost level). For tractability, we assume that investors search with replacement, and are allowed to "revisit" previously searched funds.[16] We also restrict investors to only purchase shares of one S&P 500 index fund.

Define investors' belief about the distribution of funds' indirect utilities $H(u)$. Then the optimal search rule for an investor with search cost c_i is to search for another fund as long as

$$c_i \leq \int_{u^*}^{\bar{u}} (u - u^*) dH(u), \quad (2)$$

where $\bar{u}$ is the upper bound of $H(u)$, and u^* is the indirect utility of the highest-utility fund searched to that point. This is a standard condition in sequential search models; search continues if the marginal cost of search is no greater than the expected marginal benefit. We simplify matters by assuming that investors observe the empirical cumulative distribution function of funds' utilities. That is, label the N funds by ascending indirect utility order, $u_1 < \ldots < u_N$. Then

$$H(u) = \frac{1}{N} \sum_{j=1}^{N} \mathrm{I}\left[u_j \leq u\right]. \quad (3)$$

Thus investors know the available array of indirect utilities; they just do not know which fund provides what utility level until engaging in costly search.

The optimal search rule above yields the following critical cut-off points in the search distribution:

$$c_j = \sum_{k=j}^{N} \rho_k \left(u_k - u_j\right), \quad (4)$$

where ρ_k is the probability that fund k is sampled on each search (these probabilities are common across investors and known), and c_j is the lowest possible search cost of any investor who purchases fund j in equilibrium. The intuition behind this expression is as follows. Optimal search continues until the investor's expected benefit from searching is lower than the search cost. The right-hand side of expression (4) is the expected benefit of additional search for an

[16] The search-with-replacement assumption simplifies matters in search models involving a finite number of (as opposed to continuum of) firms offering products, because we do not need to worry about how investors' beliefs about $H(u)$ evolve as certain funds are removed from consideration. This deviation from reality is small when there are a large number of funds. The revisit assumption implies of course that the investor's benefit from searching is relative to the best fund yet searched, rather than the particular fund in hand at any given time.

investor who has already found fund j. This is product of the probability ρ_k that another search yields a higher-utility fund (recall $u_k > u_j$ if $k > j$) and the corresponding utility gain $u_k - u_j$, summed over all funds superior to fund j. Draws of funds with utilities less than u_j will be ignored in this calculation, as investors are allowed to costlessly revisit funds already searched. Note that this expected benefit declines in the fund's index (in fact, $c_N = 0$). Thus as long as an investor's search cost is lower than c_j, he or she keeps searching until a fund offering greater utility than fund j is found. On the other hand, more search is not worthwhile for any investor having found fund j but with search costs greater than c_j.[17] Note that this implies the product index is declining in the ordinal ranking of critical search cost values; i.e., while $u_1<\ldots<u_N$, $c_N<\ldots<c_1$.

We can use this optimal search behavior to identify elements of the search cost distribution. Funds' market shares can be written in terms of the search cost c.d.f. by using the search-cost cutoffs above. Consider the lowest-utility fund, u_1. This fund has a high critical search value, c_1, because the expected search benefit to any investor having already found this fund is large. Therefore only those investors with very high search costs ($c > c_1$) purchase the fund; all others continue to search. At the same time, though, not *all* investors with $c > c_1$ purchase the fund; only those (unfortunate) ones who happen to draw Fund 1 first—which happens with probability ρ_1. Thus the market share of the lowest-utility fund is

$$q_1 = \rho_1\left(1 - G(c_1)\right) = \rho_1\left(1 - G\left(\sum_{k=1}^{N} \rho_k \left(u_k - u_1\right)\right)\right). \quad (5)$$

Now consider the market share of the second-lowest utility fund, Fund 2. Again a fraction of the highest-search-cost investors ($c > c_1$), unable to afford a second search, find Fund 2 first and purchase it. But a subset of investors with search costs $c_1 < c < c_2$ also purchase Fund 2; namely, those who find Fund 2 on their first search, or those search only to find Fund 1 and

[17] It is interesting to note here the links between this demand system implied by the present framework and that implied by a multiproduct logit. This model has purely vertically differentiated products, but still implies a nondegenerate market share distribution (even if all prices were equal), because the distribution of search costs across investors creates a type of horizontal differentiation. The standard logit model also introduces products that are (almost) purely vertically differentiated, but builds in horizontal differentiation (and its resulting market share dispersion) directly into the preference function with the i.i.d. random utility term.

keep searching until they draw Fund 2. The former case happens with probability ρ_2. The latter occurs with probability $\rho_2/(1-\rho_1)$.[18] Thus the total market share of Fund 2 is

$$q_2 = \rho_2\left(1-G(c_1)\right)+\frac{\rho_2}{1-\rho_1}\left[G(c_1)-G(c_2)\right]=\rho_2\left[1+\frac{\rho_1 G(c_1)}{1-\rho_1}-\frac{G(c_2)}{1-\rho_1}\right]. \quad (6)$$

Analogous calculations, detailed in the appendix, produce a generalized equation for the market shares for funds 3 through N:

$$q_j = \rho_j\left[\begin{array}{l} 1+\frac{\rho_1 G(c_1)}{1-\rho_1}+\frac{\rho_2 G(c_2)}{(1-\rho_1)(1-\rho_1-\rho_2)}+\sum_{k=3}^{j-1}\frac{\rho_k G(c_k)}{(1-\rho_1-\ldots-\rho_{k-1})(1-\rho_1-\ldots-\rho_k)} \\ -\frac{G(c_j)}{(1-\rho_1-\ldots-\rho_{j-1})} \end{array}\right] \quad (7)$$

These equations form a linear system that links observed market shares to $G(c_1),\ldots,G(c_{N-1})$—the fractions of the population with search costs less than the distribution's critical values. Moreover, we know $G(c_N) = 0$, because (4) implies $c_N = 0$ and search costs cannot be negative.[19]

The market supply side is comprised of F funds that choose prices to maximize current profits. The profits of fund j can be written as

$$\Pi_k = Sq_j(p,\underline{W})(p_j - mc_j), \quad (8)$$

where S is the total size of the market, p_j and mc_j are the price and (constant) marginal costs of fund j, and q_j is the market share of fund j given the entire price and characteristics distributions of sector funds.

This maximization implies the standard first-order condition for p_j:

$$q_j(p,\underline{W})+(p_j-c_j)\frac{\partial q_j(p,\underline{W})}{\partial p_j}=0. \quad (9)$$

I.e., the producer balances the profit-increasing effect of a marginal price increase (working through the market share at that price level) with the indirect negative effect that a price rise has on the fund's market share.

The elasticities faced by the fund are of course determined in part by the derivatives of the share equations (7). We show in the appendix that these derivatives are:

---

[18] The total probability that a search sequence yields only Fund 1 draws until a Fund 2 draw is $\rho_1\rho_2 + \rho_1\rho_1\rho_2 + \rho_1\rho_1\rho_1\rho_2 + \ldots = \rho_1\rho_2/(1-\rho_1)$. When summed with the probability that the first draw is Fund 2 (ρ_2), this yields $\rho_2/(1-\rho_1)$ as the probability that an investor with $c_2 < c < c_1$ buys Fund 2.

[19] Notice that since $G(c_N) = 0$, the market share equations only include N-1 values of $G(c)$.

$$\frac{dq_j}{dp_j} = -\frac{\rho_1 \rho_j^2 g(c_1)}{1-\rho_1} - \frac{\rho_2 \rho_j^2 g(c_2)}{(1-\rho_1)(1-\rho_1-\rho_2)} - \sum_{k=3}^{j-1} \frac{\rho_k \rho_j^2 g(c_k)}{(1-\rho_1-\ldots-\rho_{k-1})(1-\rho_1-\ldots-\rho_k)}$$
$$-\frac{\rho_j \left(\sum_{k=j+1}^{N} \rho_k\right) g(c_j)}{(1-\rho_1-\ldots-\rho_{j-1})} \tag{11}$$

Note that the density of the search cost distribution $g(c)$, evaluated at the cutoff values for those funds offering lower utility than j (i.e., $k < j$), affects the demand elasticities faced by the funds. To see why, consider investors' reactions to an increase in Fund j's price. The price increase decreases u_j. This will have two distinct effects on the critical search cost cutoff values. For $k < j$, c_k decreases (see (4)); if you hold a fund of lower quality than j, additional search becomes less appealing when Fund j has its offered utility reduced. Thus some investors with search costs less than c_{j-1} who would have formerly continued searching and serendipitously found fund j will no longer continue searching. The sales of j lost through this channel is directly related to the density of the investor population at these higher search costs, which is embodied in the first j-1 terms of (11). The second, more direct quantity effect of a price increase is created by the increase in c_j when u_j falls. That is, the benefit of continued search becomes greater for investors who would have purchased j at the original price. Some investors on the margin will now choose to continue searching and end up purchasing higher utility funds than j. The number of such marginal investors is embodied in $g(c_j)$. The final term in (11) captures this loss.

This link between fund prices and the p.d.f. of the search cost distribution, as well as the connections between market shares and the distribution's c.d.f. shown above, play an important role in empirically identifying the model. We discuss this below.

III.2. Model—Identification

The market share equations (7) show how we can map from observed market shares to the c.d.f. of the search cost distribution evaluated at the critical values. If we know (or assume) the sampling probabilities ρ_j, then all $G(c_j)$ are nonparametrically identified from market share data. Solving the linear system (7) to recover $G(c_1),\ldots,G(c_{N-1})$ and using the fact that $G(c_N) = 0$ gives all critical values of the c.d.f. If the sampling probabilities are unknown, and must be estimated, the probabilities as well as the search cost distribution can be parameterized as $\rho(\omega_1)$

and $G(c;\omega_2)$, respectively. Given ω_1 and ω_2 of small enough dimension, observed market shares can be used to estimate these parameters.

While market share data can be mapped into the c.d.f. of the search cost distribution, the actual distribution cannot in general be traced out using only share information. This is because market shares do not generically identify the *level* of the critical search cost values $c_1,\ldots,c_N$, only their relative positions in the distribution. However, shares *do* identify search cost levels in the special but often-analyzed case of homogeneous (in all attributes but price) products with unit demands; i.e., when $u_j = u' - p_j$, where u' is the common indirect utility delivered by the funds. In this case, (4) implies

$$c_j = \sum_{k=j}^{N} \rho_k \left(u' - p_k - \left(u' - p_j\right)\right) = \sum_{k=j}^{N} \rho_k \left(p_j - p_k\right). \quad (12)$$

Now, given sampling probabilities (either known or parametrically estimated), $c_1,\ldots,c_{N-1}$ can be calculated directly from observed fund prices using (12).

In the more general case where products differ in other attributes additional to prices, we require additional information to identify the cutoff search cost values. We find this information in fund companies' optimal pricing decisions. The logic of our approach is straightforward. We need to recover the p.d.f. of the search cost distribution (evaluated at the cutoff points). These values enter the derivatives of the market share equations with respect to price, (11). If we assume Bertrand-Nash competition, the first order conditions for prices (9) imply:

$$\frac{\partial q_j(p)}{\partial p_j} = -\frac{q_j(p)}{\left(p_j - c_j\right)}. \quad (13)$$

Note that we observe prices and market shares in the data. Therefore, given knowledge of marginal costs c_j, we can compute dq_j/dp_j. From (11), these derivatives form a linear system of $N - 1$ equations that can be used to recover the values of the search cost density function $g(c)$ at the critical values $c_1,\ldots,c_{N-1}$. If marginal costs are not known, they can be parameterized along with the search cost distribution and estimated from the price and market share data.

Once both the values of the search cost c.d.f. and p.d.f. (evaluated at the cutoff search costs) have been identified, we can recover the *level* of these cutoff search costs c_j in the general case of heterogeneous products. To do so, we note that by definition the difference between the c.d.f. evaluated at two points is the integral of the p.d.f. over that span of search costs. This difference can be approximated using the trapezoid method; i.e.,

$$G(c_{j-1}) - G(c_j) = 0.5\left[g(c_{j-1}) + g(c_j)\right](c_{j-1} - c_j) \quad (14)$$

We invert this equation to express the differences between critical search cost values in terms of the c.d.f. and p.d.f. evaluated at those points:

$$c_{j-1} - c_j = \frac{2\left[G(c_{j-1}) - G(c_j)\right]}{g(c_{j-1}) + g(c_j)} \quad (15)$$

Given the critical values of $G(c)$ and $g(c)$ obtained from the data above, we can recover the c_j, and from these trace out the search cost distribution.[20] In non-parametric specifications, a normalization is required: the demand elasticity equations do not identify $g(c_N)$, so a value must be chosen for the density at zero-search costs (recall that $c_N = 0$).[21]

Finally, we can also use the critical values of the search cost distribution to estimate the indirect utility function (1). The implied indirect utilities of the funds u_j are derived from the cutoff search costs via the linear system (4) above.[22] We regress these sum of these values and the respective fund's price (because of the imposed unit price coefficient) on the observable characteristics of the fund to recover β, the weights of the characteristics in the indirect utility function. We must be careful, however, as the unobservable components ξ are likely to be correlated with price, which would result in biased coefficients in ordinary least squares regressions. We use instrumental variables in an attempt to avoid this problem.

III.3. Model—Estimation

Our approach to estimating the model is to build up from the simplest version of the model, adding complexities (sometimes at the cost of parametric assumptions) as we go along, and compare the performance of the various versions in explaining the data. This both builds intuition for the reader regarding the effect of the various assumptions involved in the modeling

---

[20] Of course, since we only identify the search cost distribution at the cutoff values c_j, we don't identify the c.d.f through its entire domain. Any monotonically increasing function between the identified cutoff points could be consistent with the true distribution; our trapezoid approximation essentially assumes this is linear. The approximated c.d.f. converges to the true function as the number of funds increases.

[21] The intuition for this is that demand elasticities are determined by the actions of searchers on the margin between two funds. Given that search is responsible for spreading output shares across funds, and that changes in indirect utilities move shares on the margin only between adjacent funds, there are only N-1 margins for N funds. Thus the markup/elasticity equation system only identifies the first N-1 cutoff values of the search cost density function.

[22] Note that in our current setup, (4) implies that $u_1 = 0$, so fund utility levels are expressed relative to the least desirable fund. This normalization results from our assumption that all investors purchase a fund; if we added an outside good that could be purchased without incurring a search cost, we could alternatively normalize the utility of this good to zero.

process as well as, we hope, argues for the importance of considering departures from standard modeling frameworks.

The Basic Model. We begin by assuming that funds are homogeneous: the only characteristic that matters to S&P 500 index fund investors is price. As noted above, homogeneity implies that $u_j = u' - p_j$, where u' is common to all funds, and given sampling probabilities, the cutoff search cost values can be computed directly from observed prices.

Consider the case where funds have equal sampling probabilities; i.e., $\rho_j = 1/N \; \forall \; j$. In this case the market share equations (7) simplify to:

$$q_j = \frac{1}{N} + \sum_{k=1}^{j-1} \frac{1}{(N-k+1)(N-k)} G(c_k) - \frac{1}{(N-j+1)} G(c_j). \qquad (16)$$

As noted in the above discussion, this system theoretically identifies nonparametrically the c.d.f. of the search cost distribution. The implied $G(c_j)$ values when combined with the computed c_j from the version of (4) with $\rho_j = 1/N \; \forall \; j$ and $u_j = u' - p_j$ would allow us to trace out the search cost distribution.

However, this simple version of the model is rejected straightaway by the data. To see why, note that when sampling probabilities are equal and funds homogeneous, there will be a negative and monotonic relationship between price and market share. A glance at Figure 5, which plots the fund price vs. market share (both in logs) for the 2000 funds, shows that this is not the case in the data.[23] While there is a clear negative correlation between price and market share, the relationship is far from monotonic. (Recall as an example of this departure from monotonicity that the highest-fee fund was the 10th-largest of the over 85 funds in 2000.)

Given that this simplest version of the model is unsupported in the data, we consider two possible modifications that would make the predictions of the model more consistent with the observed market structure in the retail S&P 500 funds sector. One preserves the assumption of fund homogeneity but allows nonuniform sampling probabilities. This would break the basic model's implication of monotonicity between fund price and market share required by letting certain higher-priced funds to be sampled with a higher relative probability than some of their lower-priced competitors. The other modification, already alluded to above, is to presume that

---

[23] Other years show similar patterns; we chose 2000 because it had the largest number of funds.

funds are differentiated in characteristics other than price. In this case, a higher-priced fund may have a larger market share than its cheaper competitor because it provides other attributes valued by investors. We consider each of these more complex versions in turn.

Unequal Sampling Probabilities. Allowing sampling probabilities ρ_j to vary across funds begs the question, What determines the sampling probabilities? There are a number of plausible determinants. Broadly speaking, if sampling probabilities do differ, they are likely correlated with a fund's relative exposure in the marketplace. One possible measure of exposure is a fund's advertising expenditure. Unfortunately, we do not observe this at the fund level.[24] Another candidate would be a variable related to the size of the fund, to capture the dynamics of the search process and possible "social learning" effects. However, this poses the conceptual problem of essentially using market share to explain market share.

We propose instead having the sampling probability be a function of fund age. Specifically, we specify

$$\rho_j = \frac{Age_j^{\alpha}}{\sum_{k=1}^{N} Age_k^{\alpha}}, \qquad (17)$$

where we estimate the parameter α from the data. If $\alpha > 0$, the probability that an investor samples a fund increases in its age. We interpret this specification as a reduced-form embodiment of a market-exposure process driven by reputation and longevity effects. The unequal sampling probability implies that among a set of funds that have the same price, the older fund will have the larger market share. This is largely supported in the data: there are 11 pairs and 8 triples of funds with equal prices and unequal ages in our sample, and the older (oldest) fund has the largest market share in 8 of the pairs and 6 of the triples.

To estimate the structural parameters of the model (the search cost distribution and sampling probabilities), we use the market share equations (equation (7)), and the first-order conditions for pricing (equation (9)) and use nonlinear least squares. We parameterize the search

[24] Some studies (e.g., Sirri and Tufano 1998) have used 12b-1 fees and load charges—which we do observe—to proxy for advertising and marketing expenditures. This is not a particularly strong proxy for our purposes, as many funds advertise despite not charging either 12b-1 fees or loads, and funds that do charge such fees likely spend a considerably different amount on advertising from the actual levels of the charges. Of particular trouble for our methodology are the former group of funds: the 12b-1/load advertising proxy would imply a large number of funds that did not advertise, and it is not clear what the sampling probability would be in such cases.

costs as being lognormal with E[ln(c)] = μ and Var[ln(c)] = σ^2.[25] We also let the mean marginal cost be a free parameter that is estimated.[26] We add further flexibility to the model by allowing the mean search cost to vary across years.

The results from this estimation are presented in Table 5. Observe first that the parsimoniously parameterized model explains both price and market share extremely well, leaving only 1-2% of variation unexplained. Notice, too, that the estimated search costs are quite small. For 1995, the estimated median search cost is 5.4 basis points, implying a cost of $5.40 to search one more fund in which to invest $10,000. The interquartile range of search costs spans 1.17 to 26 basis points. We also find a downward trend in the search costs (although not all year dummies are statistically significant). According to our estimates, the median search cost in year 2000 is 3 basis points ($3 per $10000 investment), with an interquartile range from 0.64 (64 cents) to 14 basis points ($14). Our estimates imply that the large amount of price dispersion between these funds (up to 350 basis points!) can be explained by modest search costs.

Finally, observe that *Age* is a very important determinant of sampling probability. The estimated exponent is 2.3, indicating that fund age affects searchers in a non-linear fashion. However, *Age* could well be a proxy for many factors that might determine differential investment flows across funds. We experimented with other variables that could affect flows, such as the number of funds in a fund family, or the presence of a load (which means that the fund is being marketed by brokerages), but found the model did not fit the data nearly as well.

Heterogeneous Funds. Freeing up sampling probabilities allows the model to fit the data quite successfully. However, specifying the probabilities is somewhat unsatisfying because of its arbitrariness. If we instead take the tack of modifying the model by allowing heterogeneous funds, we do not need to rely on functional form assumptions to estimate the search cost distribution.[27]

---

[25] We found that a lognormal specification fit the data much better than normal, exponential, uniform, and gamma distributions. We show below that the lognormal specification fits nonparametric estimates of the search cost distribution well.

[26] The "errors" in our estimating equations can thus be interpreted as measurement error of marginal cost, and measurement error in computing market shares.

[27] Given the discussion of the assumed indirect utility function (1), this statement may seem contradictory. Note, however, that the search cost distribution is entirely determined by the values of the u_js, which can for those purposes be an arbitrary functional form of fund attributes. The linear structure placed on the utility function

Additionally, there is considerable casual evidence that investors do care about funds' non-portfolio characteristics, and as such a model which incorporates this possibility may be more realistic. According to a survey administered by Capon, Fitzsimmons and Prince (1996) to a sample of 3,386 mutual fund holders, investors rank factors like fund manager reputation, scope of funds in the family, responsiveness to inquiries as more important selection criteria than management fees. Investors also value the provision of additional services like check writing privileges and brokerage services.

To incorporate non-price differentiation into the model, we assume again that sampling probabilities are equal and funds share a common marginal cost. As shown above, the search cost distribution as well as fund utilities u_j are nonparametrically identified in this case. We use the observed fund market shares and prices to trace out the implied search cost distribution for each sample year using the procedure in section III.2. We assume that the marginal costs of the funds are identical at 8 basis points (i.e. it costs \$8 to serve a customer with \$10,000 investment). We select this level of marginal cost since the lowest price charged by the funds in our sample is 9 basis points, and the nonparametric procedure outlined in section III.2 can run into numerical problems when the price charged by a fund is lower than its marginal cost. Experimenting with marginal costs as high as 15.6 basis points (the estimate from Table 5) and as low as 0 basis points yielded similar findings.

Table 6 contains the results and reveals some interesting features of the implied search cost distributions. In the early part of the observation period, the upper percentiles of the search cost distribution are trending down (or in the case of the 99^{th} percentile, staying roughly constant). This trend continues throughout for the 75^{th} percentile. However, it is reversed after 1997 for the 90^{th}, 95^{th}, and 99^{th} percentiles, which all rise substantially through 1999, until falling again in 2000. Thus it appears that while search costs have been falling for a substantial fraction of investors, they have actually been rising at the upper reaches of the distribution.[28]

---

through (1) is helpful in the second stage of the estimation below, which allows us to put utility weights on the various attributes.

[28] Note that we do not report any results for quantiles lower than 50% in some years. The current framework implies the highest market share fund (Vanguard) offers the highest utility, and that the bulk of this market share is comprised of low-search-cost investors. Since we can only identify discrete cut-off points on the search cost distribution, and since Vanguard's market share was above 50% in all the years in our sample, our data yields little information about the lower tail of the search cost distribution. That is, any inference regarding the lower tail of the search cost distribution depends on functional form assumptions. As a result, we choose not to report any

These patterns are more visible if one "interpolates" our estimated quantiles of the search cost distribution by fitting a parametric distribution. We plot in Figure 6 the nonparametrically identified search cost quantiles for 1996 and 2000 alongside log-normal distributions fitted to these quantiles using a least-squares criterion (we only plot two years to ease clutter). This exercise implies that search costs have declined considerably for investors below the 85th percentile: the interpolated median of the 1996 search cost distribution is 1.5 basis points, whereas the 2000 median is 0.2 basis points. On the other hand, search costs appear to have risen for investors above this level, creating divergence in search costs across the population of investors.

This divergence is concurrent with three events in the industry. First, the period was a time of massive asset inflows into the sector (and into equity mutual funds in general). The second was the influx of novice mutual fund investors entering the market. Figures from the Investment Company Institute indicate that participation in the mutual fund markets rose one-third between 1996 and 2000, going from 37.2% of households holding at least one fund to 49%. Further, these new investors were measurably different from incumbent investors. As Table 7 shows, they typically had lower incomes: the percentage of mutual fund holders with incomes below $35,000 doubled between 1996 and 2000. One could reasonably make the case that these new participants had a more difficult time searching for the best fund: according to another ICI survey, 53% of mutual fund holding households in 2000 with income less than $35,000 had internet access, compared to 81% of households with incomes over $75,000. This may in part be driving our finding of an increase in search costs at the upper tail of the distribution. At the same time, the fact that overall internet use among mutual fund investors grew (from 62% of mutual fund investors in 1998 to 68% of investors in 2000—the only two years for which we have data) may account for part of the decreases seen in the bulk of the distribution.

The third concurrent event was, as discussed earlier, considerable entry of high-priced S&P 500 index funds, and the concomitant growth in the high-price segment of the sector. We interpret the above findings regarding the increase at the upper end of the search cost distribution and the changes in investors' measurable characteristics over the period as suggestive evidence that the observed shift in market shares toward higher-price funds may have been driven by the

---

corresponding values here. However, we have much more information regarding the higher quantiles of the search cost distribution that is not dependent on any assumed function form for the distribution.

entry of green mutual fund investors, with significantly higher search costs than prior investors, into the market. That is, the confluence of costly search and a composition shift in the investing population might explain one of the features of the S&P 500 index fund sector discussed above (and to some extent, the mutual fund industry as a whole). This implication of our model is only suggestive, of course—we would need data on the purchase behavior and financial market experience of individual investors to test this notion definitively—but we do find the story tantalizing.

We now use the fund utility values u_j (calculated from the critical search cost values using (4)) to infer the contribution of funds' observable characteristics to investor utility. We assume a linear functional form (1) for the indirect utility function, both to keep things computationally simple and to stay consistent with the bulk of the discrete-choice literature. Since we impose that the coefficient on price in the u function is –1, we run the following regression to recover an estimate of β:

$$u_{jt} + p_{jt} = X_{jt}\beta + \eta_{jt} \qquad (18)$$

where u_{jt} is the indirect utility of fund j in year t calculated from (4), p_{jt} is the observed price of the fund, X_{jt} are observed fund characteristics, and η_{jt} is an error term that includes the unobservable attribute ξ_{jt}.

The utility function regression results are presented in Table 8. We include ten attributes in X_{jt}. There are dummies indicating whether a fund charges a load and what type (front or rear/deferred). Loads are a pricing element (which we have already amortized into our price measure), but they are also indicative of funds that are sold through brokers rather than purchased directly from the mutual fund company. Customers may value broker services. We also include a dummy if the fund is an exchange-traded fund (i.e., SPDRs or Barclay's iShares) to control for the special features of ETFs such as enhanced liquidity and intra-day valuation. We measure the number of additional share classes attached to the fund's portfolio (for a single-share-class fund this value is zero). We also include the number of other funds managed by the same management company in order to capture any value from being associated with a large fund family, such as the ability of investors to move money among family funds with little or no transaction fees, reduced paperwork, etc. Fund age is included in the regressions as well;

customers might prefer an older, more established fund.[29] (Both the number of managed funds and age enter in logs to embody diminishing marginal effects.) Investors may also prefer a fund manager with a longer tenure, perhaps as a signal of skill or experience. Thus we include the current fund manager's tenure as well, in years (tenure varies across funds and can do so independently of the fund's age). And while all of the funds in our sample seek to match the return profile of the S&P 500 index, they do exhibit some small differences in their financial characteristics. These can result from skilled trading activities by a fund's management; it is possible through careful market timing to lower tax liability, improve relative rates of return, and reduce volatility, despite having a severely constrained portfolio. We thus include measures of tax exposure (the taxable distributions yield rate), the yearly average of the ratio of monthly fund returns to those of the S&P 500 index, and the standard deviation of monthly returns. To the extent that fund buyers prefer any persistent positive variations in financial performance, these controls should capture much of this effect.

The utility function results are presented in Table 8. Qualitatively speaking, the results seem largely sensible. The coefficient on the exchange-traded fund dummy is positive and significant, as are the number of funds managed by the same company and fund age. Furthermore, higher tax exposure affects utility negatively and significantly. The number of other share classes sharing a common portfolio as well as the fund-S&P 500 index return ratio have positive coefficients, although these impacts cannot be statistically distinguished from zero. These characteristics all enter the estimated utility function in the expected direction. The load dummy coefficients are more ambiguous. The rear/deferred load dummy enters positively and significantly, but the front load dummy is negative and insignificant. If investors preferred the advice of a broker, we might expect both coefficients to be positive, although we can't rule out that this is the case due to the imprecision of the front-load coefficient. The most puzzling result regards the coefficient on the standard deviation of monthly returns. It is positive and significantly so. Ceteris paribus, investors should prefer a fund with less return volatility, not greater.

---

[29] Because ξ_j—which is included in η_j—is likely to be correlated with fund age (funds with higher levels of desired unobservable attributes are more likely to survive), we estimate (18) using instrumental variables. We use the instruments suggested by Berry, Levinsohn, and Pakes (1995): own-product attributes (to instrument for themselves in the demand equation) as well as summary measures (average levels) of the attributes of two sets of the other products in the sector: those produced by the same company and those produced by competitors.

IV. Welfare Implications

We take the ability of the our model to explain several features of the data and to provide what we feel are sensible estimates of search costs and utility function weights, as evocative evidence that search and product differentiation play important determining roles in the retail S&P 500 index fund market. In this section, we ask whether given our estimates the presence of search causes a social welfare loss from too many funds. Since most financial portfolio needs could be met by a single S&P 500 index fund, it is plausible that the 85 funds existing in the free-entry equilibrium are too many from a social welfare standpoint.

Economists have long recognized that entry can be socially inefficient. Mankiw and Whinston (1986) formalize this notion in a model where entry causes a loss of scale economies due to the over-spreading of output across production units. Stiglitz (1987) and Stahl (1989), in theoretical exercises of particular relevance to our framework, highlight socially harmful effects on buyer search behavior as the number of producers increases. They show that when search costs have a non-degenerate distribution across consumers, an additional entrant can increase the total search costs incurred by customers. Furthermore, entry in a search equilibrium can actually increase the market power of incumbents. The basic means by which this effect operates is that an increased number of products reduces the probability of finding a lower price with each (costly) search. This can effectively lower the demand elasticities facing producers.[30]

It is important to note that it is entirely possible for a free-entry (zero-profit) equilibrium to create a net social loss. Prospective entrants, when calculating whether entry into a market is privately worthwhile, do not fully internalize certain external effects that their entry may have (such as the loss of scale economies or induced search behavior of customers). The point at which the social return of an additional entrant falls to zero may come substantially before the point at which the expected profit of the marginal entrant is zero.

To highlight the welfare implications of our estimates, we consider a polar case as a counterfactual: restricting entry into the sector to only one fund—the Vanguard 500 Index Fund. As mentioned above, it was the first entrant into the sector and is still the dominant player.

[30] Of course, there are counterbalancing positive influences of entry on social welfare. Entry can reduce distortions of market power by increasing incumbents' demand elasticities, but as mentioned above, this depends on the particulars of the search process. The second possible benefit of entry is through product-variety, as analyzed in Dixit and Stiglitz (1977) and Salop (1979), among other places. If goods are heterogeneous, and there are across-consumer differences in how attributes are valued, then additional products can be welfare enhancing. We will attempt to quantify this benefit in our welfare calculations using our fund data.

While this counterfactual is perhaps not a particularly realistic case, the stark nature of the case does serve to underscore the nature and size of the possible welfare impacts that the many sector funds have had.

Search Losses. The welfare loss from search comes from two sources. One loss occurs because some investors end up purchasing a fund that does not offer the highest indirect utility. Another loss is through investors' direct expenditure of resources on search, which is the sum over all investors of the product of an investors per-fund search cost and the number of funds the investor searches before purchasing. We compute each of these numbers below.

The loss from ending up at a fund with lower utility ranking than *N*—i.e., the gain in utility that would occur of the fund with the highest utility value (Fund *N*) had 100% market share—is easily calculated as the market-share-weighted utility difference:

$$Loss_1 = \sum_{j=1}^{N} q_j \left(u_N - u_j\right) = u_N - \sum_{j=1}^{N} q_j u_j \,, \qquad (19)$$

where the utility levels used in the calculation are those measured in the nonparametric specification above. These calculated losses are values are in the first numerical column of Table 11. The estimated losses fluctuate substantially in levels from year to year. However, as measured relative to u_N, the utility of the Vanguard fund, the loss is dropping throughout the period, from a high of 8.8% in 1995 to 2.6% in 2000. This reflects in part an increase in the relative market shares of the SPDRs.[31] These costs are not huge, not surprising since the highest-utility fund holds a majority position, but at the same time are not trivial. As will be discussed below, given the total assets in the sector, the estimated losses in dollar terms run into the several-hundred-millions.

The search loss due to direct search cost expenditures relative to the monopoly counterfactual is simply the total expenditures in the observed free-entry equilibrium, since there will be no search when entry is restricted to a single fund. To measure this loss, we need to use parametric estimates of the search cost distribution rather than the nonparametric distributions recovered above. This is because, as we describe below, we must take draws from our estimated

[31] Of course, as discussed previously, the sector also saw during this time a drop in the asset share of the lowest-priced funds and an increase in the share of high-price funds. Since the low-cost SPDRs gained significant market share throughout the period, the overall shift to higher-priced funds was driven by an asset shift in the middle parts of the price distribution from lower- to higher- priced funds, as well as market share losses by Vanguard not made up for by SPDRs.

search cost distribution in the computations. Our nonparametric identification scheme matches the critical values of the search cost distribution, but the portion of the distribution between is simply interpolated. Parameterizing the distribution smoothes this out in a reasonable way.

We re-estimate the heterogeneous fund model, but now impose a lognormal functional form on the search cost distribution. Figure 6 compares the estimated parametric distributions from this model with their nonparametric counterparts for 1996 and 2000. Notice that the correspondence, both qualitatively and quantitatively, between the two versions is reasonable. This result makes us confident in going forward with computing the total welfare costs of search using the parametric search cost functions.

To calculate the expected loss for search, recall from above that an investor with search costs higher than c_1 will stop at the first fund. If investors sample funds evenly, $(N–1)/N$ of investors with search costs between c_1 and c_2 also stop at the first fund they sample. However, the fraction $1/N$ among such investors who sampled the lowest utility fund on their first try will search again, incurring another search cost.

Continuing on with the above logic, the expected number of searches conducted by an investor with search cost $c_{k+1}<c< c_k$ will be:

$$E[n(c)\,|\,c_{k+1} < c < c_k] = \sum_{j=k+1}^{N} jp_{k,j} \prod_{l=0}^{j-1} (1 - p_{k,l}) + k \prod_{l=0}^{k} (1 - p_{k,l}) \qquad (20)$$

where

$$p_{k,l} = \frac{N-k}{N-l}. \qquad (21)$$

Integrating the incurred search costs over the population, we find that consumer surplus lost due to search in the observed equilibrium is given by:

$$Loss_2 = \sum_{\mathrm{k=1}}^{\mathrm{N}} \int_{\mathrm{x_k}}^{x_{k+1}} E[n(c)]cg(c)dc \qquad (22)$$

We evaluate the integral in (22) numerically, taking 100,000 draws from the search cost distribution. The two rightmost columns of Table 9 present the expenditures on search costs. The average number of funds searched by investors grew in conjunction with the number of funds in the sector, increasing over 2.5 times (from 9.2 to 24.9) from 1995 to 2000. However, the total average expenditure on search did not rise at the same time; indeed, they dropped by over 60%. This is due to the fact that average search costs were falling throughout the period.

As compared to the search-imposed loss from buying inferior funds, the direct search cost expenditures are somewhat smaller, typically 40-60% as large.

Additional Welfare Implications. We have ignored three other welfare implications of entry in the above analysis. First, there is no product variety welfare loss above from restricting entry to only Vanguard, despite the fact that products are differentiated. This is because this fund already offers the highest indirect utility among all funds. However, this result is an artifact of assuming purely vertical differentiation. If in reality funds are also horizontally differentiated or if there are cross-sectional differences in investors' valuations of fund attributes, both distinct possibilities, then there may well be a product-variety welfare loss from restricting entry.

Second, as formalized in Mankiw and Whinston (1986), entry can cause a loss of scale economies in production resulting from an over-spreading of output across production units. If entrants take away market share from incumbents, they decrease the output that the incumbent would have produced in the absence of entry, lowering average productivity levels if there are returns to scale in production. The resulting productivity loss is a social cost. Thus restricting entry may hold an additional benefit of productivity gains.

Third, of course, is the effect that entry has on competition and the reduction of market power. Granting a monopoly (an unregulated one at least) would open the likely possibility of greater distortions from market power. The welfare loss will of course grow in the elasticity of demand.

We first attempt to quantify product variety losses of monopoly by estimating a logit demand system and backing out the implied welfare loss when entry is restricted. Assuming that the demand system follows the standard assumptions of the logit model (see, e.g., Anderson, de Palma, and Thisse 1992), this method has two benefits. First, it is computationally simple to implement. Moreover, the logit model has well-known biases that result in it tending to overestimate the benefit of variety.[32] Thus we can interpret the welfare estimates of the logit demand system as an upper bound on the loss from restricted entry that we can compare to the gain from reduced search computed above.

[32] The most prominent of these is the implication of a pure variety benefit. That is, the logit demand system implies a welfare gain even if entering products are exactly the same as incumbents in all respects but ε_{ij}, because additional products raise the expected value of $\max(\varepsilon_{ij})$. See, for example, Petrin (2002) for a more complete discussion of the welfare measurement biases of standard logit models.

In the standard logit framework, consumers (indexed by i) choose among funds to maximize their indirect utility function $u_{ij}(X_j;\theta)$, which is a function of product attributes X_j and a parameter vector θ. This indirect utility function is a linear combination of two components:

$$u_{ij}\left(X_j;\theta\right)=\delta_j\left(X_j;\theta\right)+\varepsilon_{ij}, \quad (23)$$

where, as in our model above, δ_j is a product-specific element common to all consumers that captures the "mean" utility delivered by the product. The additional term ε_{ij} is a random (to the econometrician) taste shifter assumed to be i.i.d. Weibull across both products and consumers. The market share of fund j in this case is given by

$$s_j=\frac{\exp\left(\delta_j\right)}{1+\sum_{k=1}^{J}\exp\left(\delta_k\right)}. \quad (24)$$

Notice that in this case that the natural logarithm of the ratio of fund j's market share to that of the outside good equals δ_j. (Recall that the mean utility of the outside good has been normalized to zero, so $\exp(\delta_0) = 1$.) Thus the logit specification suggests estimating the indirect utility function by regressing the log of funds' market share ratios relative to the outside good on their price and observable characteristics. As above, δ_j includes the unobserved product attribute ξ_j, so we must use instrumental variables estimation.

Small and Rosen (1981) demonstrate that the compensating variation measure of consumer welfare resulting from a change in the choice set in the logit model is

$$\frac{\Delta E}{N}=-\frac{1}{\lambda}\left[\ln\sum_j\exp\left(\underline{\Gamma}\right)\right]_{\underline{\Gamma}^o}^{\underline{\Gamma}^f}, \quad (25)$$

where ΔE is the total compensating variation for the entire market, N is the number of consumers, λ is the marginal utility of income (the coefficient on price in the logit), $\underline{\Gamma}$ is the vector of mean utilities δ_j, and $\underline{\Gamma}^o$ and $\underline{\Gamma}^f$ are respectively the pre- and post-change values of $\underline{\Gamma}$.[33]

Table 10 reports the results from the logit estimation. The upper panel reports the estimated utility weights. Note that they are qualitatively similar to those we obtain above from

[33] The expression is a discrete-choice analogy to the standard compensating variation calculation in continuous-choice demand systems. In that case, the welfare change from an attribute variation is a combination of the variation's utility effect and the quantity-consumed difference induced by the change. Here, when attributes change in a discrete-choice framework, it is not the amount of the goods that are purchased but rather their relative probabilities of purchase.

our model. The attribute coefficients have the same signs in the two models in all cases but one, which happens to be the most puzzling of the prior specification (monthly return standard deviation). In the present case, the coefficient for this attribute has the expected negative sign. The lower panel of the table shows the computed the implied welfare loss from reduced product variety when entry is restricted to Vanguard. The implied losses are of the same order of magnitude as the benefits that would be gained from reduced search. Furthermore, the product variety benefits are rising with the number of funds throughout the sample, going from 18.0 basis points in 1995 to 47.1 in 2000, although much of this increase may be driven by the pure-variety benefit implication of the logit discussed above, since there are many more funds at the end of the sample.

We now estimate the potential social loss due to the reduction in scale economies caused by the over-spreading output across the many funds. For a subsample of our data (approximately 80% of the fund-year observations), we have data on funds' total costs. These cost numbers were obtained from Lipper Analytical Services, who use funds' annual expense statements to compute costs as a share of the average net assets managed by the fund over the course of the year. We use this data to estimate a starkly simple cost function that yields an estimate of the fixed yearly operating cost of the fund. We estimate a cost function that is linear in the total net assets of the fund; i.e.,

$$C_t = F + c \cdot TNA_t. \quad (26)$$

That is, the total cost of operating a fund during a given year (i.e., "production" costs) are equal to a fixed cost F plus the product of the fund's total net assets under management and a constant marginal cost c. We assume both the fixed and marginal costs are constant over the time period of our sample.

While extremely simple, this specification offers an easily characterized welfare loss of entry through the loss of scale economies. The social cost of an additional fund is simply the fixed cost, because any assets can be shifted among funds at a the same marginal cost. If on the other hand returns to scale in the industry were characterized by declining marginal costs, for example, this loss is much harder to measure, because it depends on how many assets are taken from which funds and those funds' initial assets levels. While clearly over-simplistic, we believe our simple functional form for costs is a reasonable first-order approximation.

The cost function estimates imply an annual fixed cost of $1.22 million (s.e. = $0.16 million) and a marginal cost of $0.00182 per dollar of assets—18.2 basis points (s.e. = 0.1 basis points).[34] The coefficients are estimated fairly precisely, particularly for marginal cost, and the goodness of fit is high (R^2 = 0.98). Both coefficient estimates strike us as reasonable; the estimated marginal cost is less than most of the prices of funds in our sample, although surely in reality there is some variation in this value across funds.[35]

The total welfare cost of the loss of scale economies in our counterfactual case is simply the sum of the fixed costs of the extra funds. In 2000, for example, since there were 82 funds, 81 funds paid total fixed costs that were $98.8 million above those incurred by a monopolist producer.[36] In reality, it is likely that a single index fund with such a large amount of assets could run up against increasing marginal costs. This would of course result in a lower implied average productivity gain from reducing the number of funds. Thus our estimates should be considered an upper bound.

Table 11 combines all of the estimated welfare effects. It shows, in both basis points and dollar values, the total welfare changes that would be induced by the imposition of a Vanguard monopoly. The estimates imply that there would be an estimated welfare benefit of restricting entry in every year except 2000. The gains from reduced search costs and fewer resources spent on fixed fund operating costs would be larger than the welfare loss from reduced product variety, if there were horizontal taste differences across investors. The gains range from $161.2 million in 1995 to $557.2 million in 1999. In 2000, however, the product variety loss is great enough to make a monopoly market structure less socially desirable that the observed equilibrium; a net welfare loss of $122.3 million is implied. To put these values in perspective of the prices

---

[34] To do a back-of-the envelope check on this marginal cost estimate, consider that the average mutual fund account (across all fund types, not just in our sector) in 1997 was just over $26,000. Our estimate then implies a marginal annual cost per account of around $47, which is just higher than the estimated $30-40 cost of providing an account with basic administration (accounting, printing statements, etc.) reported in *The Economist* (2002). Adding additional trading and marketing marginal costs to this estimate would bring its value closer still to our implied marginal cost.

[35] We checked the robustness of the cost function results by including fund fixed effects to remove the influence of any across-fund differences in fixed costs may have on the marginal cost estimate. Doing so resulted in a marginal cost estimate of 17.9 basis points (s.e. = 0.2 b.p.). We also estimated a specification that allowed fixed costs to be different for funds that are part of a multiple share class portfolio, and found no significant difference.

[36] We are ignoring here the small number of funds, typically quite small in terms of total assets, that we know existed in the sector but do not have data for. (There were 85 funds observed in operation during some time in 2000, for example.) If the fixed costs of these other funds are included, these numbers would be slightly higher.

investors face, we have also calculated the increase in Vanguard's price that, once it was a monopolist, would leave the average investor in the sector indifferent between the observed market structure and the imposed monopoly.[37] Considering that Vanguard's price is just under 20 basis points through most of our sample, the implied price changes that would make investors indifferent are substantial.

These calculations come with a host of caveats. We have entirely ignored any losses that may arise due to the deadweight loss of monopoly (not to mention any transfer in surplus from the consumer to the producer that may matter to those concerned about distributional effects). These could well be quite large. While these market power losses could plausibly be avoided by regulating the monopolist, regulation induces its own well known welfare costs (rent-seeking, corruption, etc.) that could be themselves substantial. Also, the product variety welfare and fixed cost savings estimates were made using assumed functional forms that, while facilitating straightforward interpretation, at the same time result in estimates that should be considered bounds rather than point estimates. Therefore it is likely that the true welfare loss from less variety in the monopoly counterfactual is less than estimated, and the actual productivity increase that would obtain from restricting entry would be smaller than the numbers above. These in combination would tend to cancel each other out, so what would happens on net depends on the relative size of the departure of the actual welfare impacts from their calculated bounds.

Given these cautionary notes, we are reluctant to prescribe policy based on our welfare findings. However, we do look at them as being within the bounds of reasonableness, and as such may offer guidance when thinking about the impact that search and product differentiation have on welfare in the retail S&P 500 index fund sector in particular and the mutual fund industry as a whole.

V. Conclusions

We have presented evidence that key features of the U.S. mutual fund industry are driven by factors beyond financial portfolio heterogeneity alone. We focus on an asset segment, retail S&P 500 index funds, where all fund are characterized by nearly homogeneous return patterns.

---

[37] To keep things simple, we are assuming inelastic demand here, so no assets would be pulled out of the sector in response to a price change.

Despite the homogeneity, we find that this sector exhibits the fund proliferation and price dispersion patterns seen in the broad industry.

We consider a combination of search frictions and non-financial fund differentiation as explanations for these observations. We model a sector equilibrium where investors with heterogeneous search costs shop over differentiated funds, and these funds compete with each other in prices mindful of investors' search behavior. We estimate the model using data on the funds in the retail S&P 500 index fund sector and find that the search costs necessary to sustain the observed price dispersion are of plausible magnitude. Indeed, the estimated search costs exhibit much less dispersion than the price variation they support. This is in part driven by the fact that our framework allows funds to be vertically differentiated. Our model furthermore lets us back out the contribution of various observable fund attributes to the indirect utility of investors who purchase the funds. The implied utility weights are qualitatively sensible.

Our estimated search cost distributions also shed light on developments in the sector over the course of our sample. We observe considerable entry of high-price funds into the sector that was accompanied by a shift in assets toward more expensive funds. This is despite the fact that technological improvements over the same period plausibly led to decreases in the average cost of search. The estimated search cost distributions offer an explanation for these seemingly divergent features. We find that while average search costs were declining, costs for those at the upper deciles of the distribution actually tended to increase through our sample years. This spreading of the distribution was concurrent with the well documented increase in households' first-time participation in the mutual fund market. We take this as suggestive evidence that novice investors with high search/learning costs were causing this shift of assets into higher-price funds and supporting the high-fee entrants.

We also consider the welfare implications that search costs and differentiated products pose when there are a large number of financially identical funds. We find that the total costs sunk into the search process, both directly and indirectly through investors purchasing lower-utility funds, are sizeable. These costs could be saved by restricting entry into the sector to a monopolist fund. Of course, this could well cause welfare losses of its own due to the reduction in product variety and the deadweight loss from increased market power. On net, our rough calculations indicate that imposing monopoly might be socially beneficial. Given the myriad simplifying assumptions that are required in such calculations, however, we are reluctant to make any policy recommendations too strongly.

While we chose a particular mutual fund asset class to control for financial performance heterogeneity while highlighting our search-cost / differentiated-product explanation, we think our results also suggest at least partial explanations for the fund proliferation and large fee dispersion seen in the mutual fund industry as a whole. Much more work needs to be done, however, to fully characterize these impacts.

A. Market Share Equations

The market share of a particular fund depends on the probabilities that investors with search costs less than or equal to its corresponding cutoff search cost value from (2) finds the fund before finding another fund with a critical value higher than their search costs.

More formally, consider an investor with search cost $\hat{c}$. Then the probability that our investor ends up purchasing a particular Fund k is equal to either (a) zero if $c_k > \hat{c}$, since the investor will always keep searching upon drawing Fund k, or (b) if $c_k < \hat{c}$, the probability that he draws k before drawing any other fund with a critical value less than $\hat{c}$. The latter probability is equal to ρ_k times the sum of the probabilities of draw sequences where all funds have $c_k < \hat{c}$. We can express these probabilities, and therefore the corresponding market share equations (7), compactly using combinatorics. We derive these equations below.

First, define $\hat{c}_l$ as the largest fund cutoff search cost value that is less than $\hat{c}$:

$$\hat{c}_l \equiv \max\{c_j \mid c_j < \hat{c}\}, \qquad \text{(A.1)}$$

An investor with search cost $\hat{c}$ will stop searching once a fund with $c_j \le \hat{c}_l$ is drawn, because at this point the benefit of additional search is less than its cost. Consider drawing such a fund as a "success" event. Then the probability of T failures (i.e., where all T draws are of funds with $c_j > \hat{c}_l$) is equal to

$$\sum_{\bar{a}\mid a_1+\cdots+a_l=T}\left[\frac{T!}{a_1!\ldots a_l!}\rho_1^{a_1}\ldots\rho_l^{a_l}\right], \qquad \text{(A.2)}$$

where the l subscript denotes the fund with $c_j = \hat{c}_l$, ρ_j is the probability of sampling Fund j, a_j is the number of times that the fund is drawn in the sequence of T draws, and the sum is taken over all combinations of the a vector $[a_1,\ldots,a_l]$ that sum to T.

It is known from combinatorics theory that the above summation simplifies to $(\rho_1 + \ldots + \rho_l)^T$. Therefore the probability that Fund k is chosen after T failures—that k is the "success" draw—is $\rho_k(\rho_1 + \ldots + \rho_l)^T$. This expression must be summed over values of T (from $T = 0$, an immediate success where k is the first fund drawn, to $T = \infty$, the limit of possibility) to obtain the total probability that an investor with search cost $\hat{c}$ purchases Fund k (alternatively, the probability that k is chosen given $c_k \le \hat{c}_l$). That is,

$$\Pr(\text{Fund } k \text{ chosen} \mid c_k \le \hat{c}_l) = \rho_k \sum_{T=0}^{\infty}(\rho_1 + \ldots + \rho_l)^T = \frac{\rho_k}{1-\rho_1-\ldots-\rho_l} \qquad \text{(A.3)}$$

This is the probability (b) in the second paragraph above.

Of course, this probability depends on the value of $\hat{c}$, since this determines the particular value of ρ_l. Investors with differing search costs thus have various probabilities of finding a "success" in Fund k. For example, the highest-search-cost investors (i.e., those with $\hat{c} > c_1$) have $\hat{c}_l = c_1$, so they only take one fund draw and purchase whichever fund find. Thus the contribution to Fund k's market share from investors with search costs above c_1 is therefore $\rho_k[1 - G(c_1)]$. Analogously, for all $k \ge 2$, the market share comprised of investors in the next-lower segment of the search cost distribution (between c_2 and c_1) are those investors with that only draw Fund 1 until Fund k is chosen. This market share is then $[G(c_1) - G(c_2)]\rho_k/(1 - \rho_1)$. These are the values discussed in the text for $k = 2$.

This is easily generalized. Investors with search costs between c_3 and c_2 purchase k if $k \geq 3$ and if k is the first fund drawn besides Fund 1 or 2. Thus their contribution to the market share of Fund k is $[G(c_2) - G(c_3)]\rho_k/(1 - \rho_1 - \rho_2)$. Generically, investors with search costs between cutoff values c_j and c_{j-1} account for a market share for Fund k equal to the following, if $k \geq j$:

$$\frac{\rho_k}{1-\rho_1-\ldots-\rho_{j-1}}\left[G(c_{j-1})-G(c_j)\right]. \qquad \text{(A.4)}$$

The total market share for k is then the sum of these values for all segments of the search cost distribution above c_k:

$$q_k = \rho_k\left[1-G(c_1)\right]+\frac{\rho_k}{1-\rho_1}\left[G(c_1)-G(c_2)\right]+\ldots+\frac{\rho_k}{1-\rho_1-\ldots-\rho_{k-1}}\left[G(c_{k-1})-G(c_k)\right] \qquad \text{(A.5)}$$

Grouping common $G(c)$ terms, factoring out a ρ_k, and evaluating at $k = j$ yields expression (7) in the text.

B. Derivatives of Demand Curves

The generalized market share equations are given in (7):

$$q_j = \rho_j\left[\begin{array}{l} 1+\dfrac{\rho_1 G(c_1)}{1-\rho_1}+\dfrac{\rho_2 G(c_2)}{(1-\rho_1)(1-\rho_1-\rho_2)}+\displaystyle\sum_{k=3}^{j-1}\frac{\rho_k G(c_k)}{(1-\rho_1-\ldots-\rho_{k-1})(1-\rho_1-\ldots-\rho_k)} \\ -\dfrac{G(c_j)}{(1-\rho_1-\ldots-\rho_{j-1})} \end{array}\right] \qquad \text{(A.6)}$$

we want to take price derivatives of these equations. Notice is where prices enter into the equations: the indirect utilities provided by the funds include their prices, and these indirect utilities are in turn embodied in the cutoff search cost values $c_1,\ldots,c_j$ above.

So with this in mind we can take the price derivative of the above:

$$\begin{aligned} \frac{dq_j}{dp_j} = {} & \frac{\rho_1\rho_j g(c_1)}{1-\rho_1}\frac{dc_1}{dp_j}+\frac{\rho_2\rho_j g(c_2)}{(1-\rho_1)(1-\rho_1-\rho_2)}\frac{dc_2}{dp_j}+\sum_{k=3}^{j-1}\frac{\rho_k\rho_j g(c_k)}{(1-\rho_1-\ldots-\rho_{k-1})(1-\rho_1-\ldots-\rho_k)}\frac{dc_k}{dp_j} \\ & -\frac{\rho_j g(c_j)}{(1-\rho_1-\ldots-\rho_{j-1})}\frac{dc_j}{dp_j} \end{aligned} \qquad \text{(A.6)}$$

Now recall from (4) that

$$c_j = \sum_{k=j+1}^{N}\rho_k\left(u_k-u_j\right). \qquad \text{(A.7)}$$

Combined with the fact that the derivative of the u with respect to price is –1, this implies

$$\frac{dc_r}{dp_j} = \begin{cases} -\rho_j, & \text{if } r < j \\ \displaystyle\sum_{k=j+1}^{N}\rho_k, & \text{if } r = j \\ 0, & \text{if } r > j \end{cases}. \qquad \text{(A.8)}$$

The above takes advantage of our normalization assumption, $du_j/dp_j = -1$. Substituting this into (A.6) gives:

$$\frac{dq_j}{dp_j} = -\frac{\rho_1 \rho_j^2 g(c_1)}{1-\rho_1} - \frac{\rho_2 \rho_j^2 g(c_2)}{(1-\rho_1)(1-\rho_1-\rho_2)} - \sum_{k=3}^{j-1} \frac{\rho_k \rho_j^2 g(c_k)}{(1-\rho_1-\ldots-\rho_{k-1})(1-\rho_1-\ldots-\rho_k)} - \frac{\rho_j g(c_j)\sum_{k=j+1}^{N}\rho_k}{(1-\rho_1-\ldots-\rho_{j-1})} \tag{A.9}$$

This is equation (11) in the text.

In the special case of equal sampling probabilities ($\rho_j = 1/N \; \forall \; j$), this simplifies to:

$$\frac{dq_j}{dp_j} = -\sum_{k=1}^{j-1} \frac{g(c_k)}{N(N-k+1)(N-k)} - \frac{(N-j)g(c_j)}{N(N-j+1)}. \tag{A.10}$$

Data Appendix

The bulk of our performance and characteristic data on mutual funds comes from the CRSP mutual fund database for the years 1992-2000. This data includes a considerable amount of information about the funds (as mentioned previously, different share classes for the same asset pool are considered separate funds). Most of the data is compiled annually, but monthly information on returns and assets under management are also available. We observe the year the fund was established, the identity of the fund's manager and managing company, the starting date of the manager's tenure, and whether the fund in a given year. Annual performance and portfolio characteristics included in the data are income and capital gains distributions. Pricing information includes annual management fees; 12b-1 fees; and front, rear, and deferred load levels.

We use this data to compute a number of supplementary variables for our analysis. These include fund age, the total number of funds managed by a fund's management company (including those outside the S&P 500 index category), the average and standard deviation of monthly returns, and measures of gross return benchmarked to the return of the S&P 500 index.

We supplement the CRSP data with mutual fund cost data from Lipper Analytical Services for 1995-2000. Lipper examines funds' year-end reports to gather cost data in a number of categories. These cost numbers are aggregated and combined with asset data to yield a total annual costs. We have cost data for approximately 80% of our sample fund observations over 1995-2000.

Two measures are central to our empirical work: price and market share. Unfortunately, for mutual funds, neither concept is as straightforward as it often is for other products. Several central issues arise; we discuss those relevant to each measure in turn.

There are a number of ways in which mutual fund prices vary from traditional concepts of price. One of these is that, invariably, rather than being priced as a simple dollar level, mutual fund prices are fractional charges related to the size of asset flows into or stocks held in the fund. Furthermore, there are several dimensions along which mutual funds can be priced. The standard margin employed by all funds is the annual management fee. This is a percentage of the assets that a fund owner holds in a fund that is withdrawn from the owner's account and used to reimburse the fund management company. Some funds also charge a 12b-1 fee, which is also an annual charge earmarked for fund distribution expenses that is a flat percentage of assets held in the fund. There are also one-time loads imposed by some funds. The size of loads are expressed as a percentage of fund flows into or out of a fund. There are two types of loads typically employed by the industry. Front-end loads are charged at the time of a purchase of fund assets. Back-end (or deferred) loads are charged at a time of withdrawal.

Funds differ in both the pricing margins and their levels. Furthermore, because of the stock/flow distinction between annual fees and loads, as well as the timing discrepancy between front- and back-end loads, it is not a trivial matter to identify a single price for each fund. We use the approach, common in the literature, of measuring fund price by adding annual fees (both management and 12b-1 fees) to one-seventh of the sum of all load levels. The one-seventh fraction is obtained from the stylized fact that a typical mutual fund account is held for about seven years. Loads are incurred only when there are flows into or out of a fund, not on any asset stocks held.

Hence the price is meant to incorporate the shareholder's annualized cost of the load. We consider this price measure to be a reasonable compromise between a number of competing factors.

Market size is also less than empirically clear-cut in this industry. While market size measurement issues are not uncommon, they typically revolve around defining the boundaries of the market. In our case, however, this is secondary. What is more difficult is defining the unit of purchase. On one hand, assets are a plausible dimension along which to measure market shares. By doing so we are implicitly assuming that mutual fund customers evaluate on an annual basis whether to continue to hold their assets in a particular fund or to move them elsewhere. (This is further complicated by the fact that loads apply to flows but management fees apply to stocks.) On the other hand, we define market share in standard markets as the share of total purchases accounted for by a producer; i.e., the flow of funds, not the stock. This measure is itself complicated by the fact that our data do not allow full measurement of gross fund purchases. Because we only measure total assets under management, we only observe net flows into a fund. Net flows hide the size of the gross flows underlying them, and gross inflows are the preferable measure to base flow-centered market shares upon. (For example, in our 2000 data, it was not uncommon to see negative net changes in fund assets over the course of the year. It is not exactly clear just how one could define a flow-based market share measure from these negative values.) We are partly able to circumvent this problem because we observe monthly asset and return data. Using these we can back out monthly net flows and then sum the positive net monthly flows to get an aggregated measure of gross flows.

References

Anderson, Simon P., André de Palma, Jacques-François Thisse. *Discrete Choice Theory of Product Differentiation*. Cambridge, MA: MIT Press, 1992.

Berry, Steven, James Levinsohn, and Ariel Pakes. "Automobile Prices in Market Equilibrium." *Econometrica*, *63*(4), 1995, pp. 841-90.

Berry, Steven and Joel Waldfogel. "Free Entry and Social Inefficiency in Radio Broadcasting." *RAND Journal of Economics*, Forthcoming.

Bresnahan, Timothy F. and Peter C. Reiss. "Entry and Competition in Concentrated Markets." *Journal of Political Economy*, *99*(5), 1991, pp. 977-1009.

Burdett, Kenneth and Kenneth L. Judd. "Equilibrium Price Dispersion." *Econometrica*, *51*(4), 1983, pp. 955-69.

Carhart, Mark M. "On Persistence in Mutual Fund Performance." *Journal of Finance*, *52*(1), 1997, pp. 57-82.

Carlson, John A. and R. Preston McAfee. "Discrete Equilibrium Price Dispersion." *Journal of Political Economy*, *91*(3), 1983, pp. 480-93.

Diamond, Peter. "A Model of Price Adjustment." *Journal of Economic Theory*, 3 (June 1971): 156-68.

Dixit, Avinash K., and Joseph E. Stiglitz. "Monopolistic Competition and Optimum Product Diversity." *American Economic Review*, *67*(3), 1977, pp. 297-308.

"Everyman, Farewell." *The Economist*, Jan. 10th, 2002.

Foster, Lucia, John Haltiwanger, and C. J. Krizan. "The Link Between Aggregate and Micro Productivity Growth: Evidence from Retail Trade." NBER Working Paper 9120, 2002.

Goolsbee, Austan and Amil Petrin. "The Consumer Gains from Direct Broadcast Satellites and the Competition with Cable Television." NBER Working Paper 8317, 2001.

Hong, Han and Matthew Shum. "Estimating Equilibrium Search Models." Working paper. 2001.

Hunter, Matthew. "Putnam Fund Consolidation May Signal Industry Trend." *The American Banker*, March 20, 2002, p. 5.

Investment Company Institute. *Mutual Fund Fact Book, 2002*. Washington, DC. 2002.

Mankiw, N. Gregory and Michael D. Whinston. "Free Entry and Social Inefficiency." *Rand Journal of Economics*, *17*(1), 1986, pp. 48-58.

Mortensen, Dale T. "Equilibrium Wage Distributions: A Synthesis." *Panel Data and Labor Market Studies*. Amsterdam; Oxford and Tokyo: North-Holland, 1990, pp. 279-96.

Nevo, Aviv. "A Practitioner's Guide to Estimation of Random-Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy*, *9*(4), 2000, pp. 513-48.

Petrin, Amil. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy*, *110*(4), 2002, pp. 705-29.

Small, Kenneth A. and Harvey S. Rosen. "Applied Welfare Economics with Discrete Choice Models." *Econometrica*, *49*(1), 1981, pp. 105-30.

Sorensen, Alan T. "Price Dispersion and Heterogeneous Consumer Search for Retail Prescription Drugs." UCSD Working Paper, 2001.

Stahl, Dale O. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, *79*(4), 1989, pp. 700-12.

Stahl, Dale O. "Oligopolistic Pricing with Heterogeneous Consumer Search." *International Journal of Industrial Organization*, *14*(2), 1996, pp. 243-68.

Stiglitz, Joseph E. "Competition and the Number of Firms in a Market: Are Duopolies More Competitive than Atomistic Markets?" *Journal of Political Economy*, *95*(5), 1987, pp. 1041-61.

Table 1: Price Dispersion within Fund Sectors

Sector	N	Mean Price	Coefficient of Variation	75th to 25th %ile Ratio	90th to 10th %ile Ratio
Aggressive Growth	1237	227.3	0.529	2.56	4.05
Balanced Growth	452	202.9	0.526	2.99	4.89
High-Quality Bonds	824	143.3	0.682	2.94	6.75
High-Yield Bonds	327	213.5	0.473	2.82	4.34
Global Bonds	332	222.6	0.482	2.61	4.63
Global Equities	428	272.7	0.387	1.92	3.38
Growth and Income	950	194.7	0.797	3.35	7.19
Ginnie Mae	173	179.4	0.544	3.24	5.37
Gov't Securities	435	163.8	0.642	3.33	6.38
International Equities	1209	263.0	0.461	2.35	3.88
Income	205	208.5	0.508	2.83	4.60
Long-Term Growth	1756	219.7	0.495	2.63	4.21
Tax-Free Money Mkt	441	72.4	0.502	1.94	3.40
Gov't Securities Money Mkt	427	70.4	0.740	2.25	6.00
High-Quality Muni Bond	502	173.5	0.657	3.44	5.89
Single-State Muni Bond	1309	192.5	0.471	2.09	5.02
Taxable Money Mkt	528	97.8	0.860	2.31	9.94
High-Yield Money Mkt	62	209.4	0.468	2.07	4.68
Precious Metals	35	293.0	0.408	1.96	3.61
Sector Funds	501	238.6	0.438	2.16	3.77
Total Return	301	224.4	0.474	2.29	4.60
Utilities	87	230.6	0.422	2.11	4.35
Selected Specific Sectors					
Financial Sector	65	254.4	0.366	1.78	3.27
Health Sector	62	252.7	0.398	2.07	2.95
Tax-Free Bonds (California)	133	187.5	0.480	2.09	5.11
Technology Sector	171	238.2	0.426	2.18	3.45

Note: Prices computed from CRSP data (annual fees + one-seventh of total loads) expressed in basis points. Data for 2000.

Table 2: Financial Homogeneity of Non-Institutional S&P 500 Index Funds

Variable	Statistic	1995	1996	1997	1998	1999	2000
No. of Funds		22	23	33	51	64	75
Annual Gross Return (%)	Median	36.74	22.29	32.63	28.08	20.20	-9.54
	IQR	0.58	0.65	0.60	0.63	0.63	0.55
	Mean	36.72	21.99	32.51	28.13	20.09	-9.31
	SD	0.51	1.37	0.54	0.97	0.71	1.17
Avg. Gross Monthly Return (%)	Median	2.652	1.858	2.633	2.450	1.791	-0.705
	IQR	0.079	0.110	0.148	0.212	0.224	0.037
	Mean	2.860	1.877	2.687	2.548	1.859	-0.672
	SD	0.152	0.151	0.167	0.236	0.181	0.127
Std. Dev. Monthly Returns (%)	Median	1.491	3.128	4.584	6.195	3.774	4.939
	IQR	0.016	0.023	0.029	0.028	0.024	0.039
	Mean	1.491	3.132	4.572	6.199	3.808	4.931
	SD	0.025	0.039	0.052	0.052	0.116	0.221

Note: Sample limited to funds reporting returns in every month of given year.

Table 3: Evolution of Non-Institutional S&P 500 Index Funds Price Distribution

Statistic	1995	1996	1997	1998	1999	2000
Number of Funds	24	33	45	57	68	82
Minimum	19.83	18.0	18.0	18.0	18.0	9.45
10th Percentile	21.5	20.0	23.2	25.7	27.4	28.0
25th Percentile	43.8	45.0	40.0	46.0	47.1	47.1
Median	82.5	60.0	84.7	93.0	102.0	92.0
75th Percentile	127.9	124.9	149.3	160.9	193.4	168.6
90th Percentile	168.4	174.7	255.9	262.1	301.7	292.4
Maximum	306.4	306.4	331.4	331.4	335.4	368.4
Interquartile Range	84.1	79.9	109.3	114.9	146.2	121.4
90th-10th Percentile Range	146.9	154.7	232.7	236.4	274.2	264.4
Mean	93.5	91.5	109.3	117.3	133.9	127.5
Asset-Weighted Mean	27.2	26.9	26.7	30.8	32.9	36.5

Note: All prices in basis points

Table 4. Decomposition of Asset-Weighted Average Price Growth

Year-to-Year Change	1995-1996	1996-1997	1997-1998	1998-1999	1999-2000
Growth in Asset-Weighted Mean Price	-0.236	-0.271	4.113	5.004	0.678
Within Component	-0.129	-1.589	-0.100	-0.144	0.309
Between—Fixed Prices	-0.366	1.263	0.647	0.640	0.724
Between—Covariance	-0.026	-0.556	0.052	0.029	-0.212
Entry	0.284	0.611	3.514	4.678	-0.145
Exit	0	0	0	0.198	-0.002
Net Entry Effect (Entry – Exit)	0.284	0.611	3.514	4.479	-0.143

Note: All measures are in basis points.

Table 5: Search Model with Unequal Sampling Probabilities

Parameters	Estimates
Log(Mean Search Cost)	-7.497*
	(1.187)
Variance of Logged Search Cost Dist.	2.304*
	(0.235)
Mean Marginal Cost, basis points	15.53*
	(5.41)
α	2.307*
(*Age* exponent in sampling probability)	(0.295)
1996	-0.276
	(0.255)
1997	-0.007
	(0.276)
1998	-0.475*
	(0.236)
1999	0.189
	(0.278)
2000	-0.609
	(0.649)
R^2, (prices)	0.948
R^2, (quantities)	0.989
Mean Squared Error	9.03e-5
Median Search Cost (1995), basis points	5.5
IQR of Search Cost (1995), basis points	1.2-26

Note: Standard errors, computed using the delta method, are in parentheses.

Table 6: Nonparametric estimates of search cost quantiles

Year	Percentile of Search Cost Distribution (basis points): 50th	75th	90th	95th	99th
1995	N/A	8.9	23	53	112
1996	N/A	5.0	19	25	114
1997	N/A	3.0	15	38	161
1998	N/A	2.7	18	50	170
1999	0.8	2.3	42	76	236
2000	0.2	1.6	23	52	190

Table 7: Characteristics of Households Holding Mutual Funds
(Source: ICI Fundamentals Research Briefs, 1996-2000)

A. Income

Year	% of Mutual Fund Holders with Income Less Than: $25,000	$35,000	$50,000	$75,000	$100,000	> $100,000
1996	N/A	18	41	58	73	N/A
1997	12	31	38	57	72	N/A
1998	13	28	47	62	72	77
1999	15	30	49	62	78	78
2000	17	37	49	66	77	79

B. Age

Year	< 25	25 to 34	35 to 44	45 to 54	55 to 64	> 65
1996	N/A	28†	32	19	11	10
1997	4	18	28	20	15	15
1998	3	18	26	22	15	16
1999	3	19	27	23	13	15
2000	2	18	28	25	14	14

† Separate data for under age 25 not available in 1996. Reported number reflects all households with head younger than 35.

Table 8. Utility Function Estimates and Fund Attribute Summary Statistics

Attribute	Utility Weight, basis points (s.e.)	Mean	Std. Dev.
Constant	133.2	N.A.	N.A.
	(85.9)		
Front Load Dummy	-56.2	0.278	0.449
	(30.8)		
Rear/Deferred Load Dummy	87.2*	0.272	0.446
	(32.3)		
Exchange-Traded Fund	219.0*	0.023	0.149
	(106.4)		
Number Other Share Classes	17.6	1.621	1.337
	(10.9)		
ln[No. Funds in Same Mgmt. Company]	35.1*	4.252	1.217
	(12.9)		
ln[Fund Age]	148.8*	1.393	0.728
	(57.4)		
Manager Tenure (yrs.)	0.60	2.922	2.776
	(11.9)		
Income+Capital Gains Yield (%)	-9.33*	3.248	3.363
	(4.34)		
Avg. Monthly % Diff. btw. Fund and S&P 500 Returns	85.8	-0.081	0.108
	(126.3)		
Std. Dev. Monthly Returns (%)	58.2*	4.455	1.293
	(10.2)		
N	309	309	309
R^2	0.295		
Mean of Dependent Variable	755.3		

Notes: Fund age instrumented for using characteristics of other funds. See text for details.

Table 9. Loss From Search

Year	Average Loss from Buying Inferior Funds, basis points	Average Number of Funds Searched	Total Average Expended Search Costs, basis points
1995	49.9	9.2	28.5
1996	43.3	12.7	23.3
1997	37.3	16.5	16.7
1998	33.5	19.4	14.2
1999	47.2	22.2	18.8
2000	22.4	24.9	10.9

Table 10. Logit Demand System Estimates and Welfare Implications

A. Utility Function Estimates

Attribute	Estimated Utility Weight	Standard Error
Constant	-9.860*	0.799
Price	-0.013*	0.004
Front Load Dummy	-0.710	0.371
Rear/Deferred Load Dummy	1.178*	0.523
Exchange-Traded Fund	4.452*	0.853
Number Other Share Classes	0.087	0.120
ln(No. Funds in Same Mgmt. Company)	0.403*	0.105
ln(Fund Age)	1.374*	0.521
Manager Tenure	0.072	0.101
Income+Capital Gains Yield	-0.013	0.037
Average % Monthly Return Difference	1.067	1.048
Std. Dev. Monthly Returns	-0.132	0.081
N	309	
R^2	0.476	

B. Implied Product Variety Welfare Loss of Vanguard Monopoly

Year	Loss (basis points)
1995	17.9
1996	19.7
1997	27.1
1998	35.1
1999	38.8
2000	46.8

Table 11. Summary of Welfare Changes In Vanguard Monopoly Counterfactual

	(1)	(2)	(3)	(4)	(5) = \$1.22m x [(1)-1]	(6) = (2) x (3)	(7) = (2) x (4)	(8) = (5)+(6)+(7)	(9) = (8) ÷ (2)
Year	Funds	Assets (\$billion)	Savings from Search (b.p.)	Product Variety Cost (b.p.)	Fixed Costs Savings (\$million)	Savings from Search (\$million)	Product Variety Cost (\$million)	Net Welfare Change (\$million)	Indifference Monopolist Price Change (b.p.)
1995	24	22.0	78.4	-17.9	28.1	172.5	-39.4	161.2	73.3
1996	33	39.4	66.6	-19.7	39.0	262.4	-77.6	223.8	56.8
1997	44	70.6	54.0	-27.1	52.5	381.2	-191.3	242.4	34.3
1998	57	118.0	47.7	-35.1	68.3	562.9	-414.2	217.0	18.4
1999	68	174.8	66.0	-38.8	81.7	1153.7	-678.2	557.2	31.9
2000	82	163.8	33.3	-46.8	98.8	545.5	-766.6	-122.3	-7.5

Figure 1: Number and Assets of Retail S&P 500 Index Funds

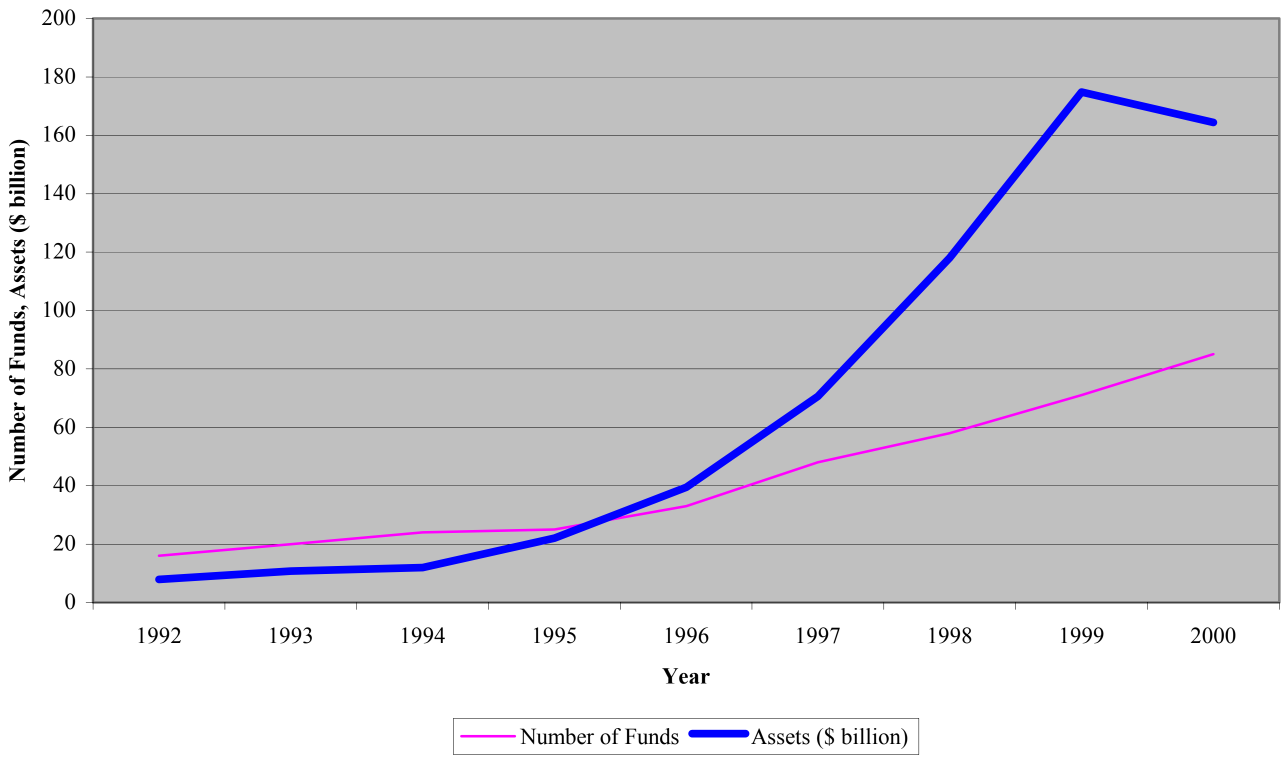

Figure 2.

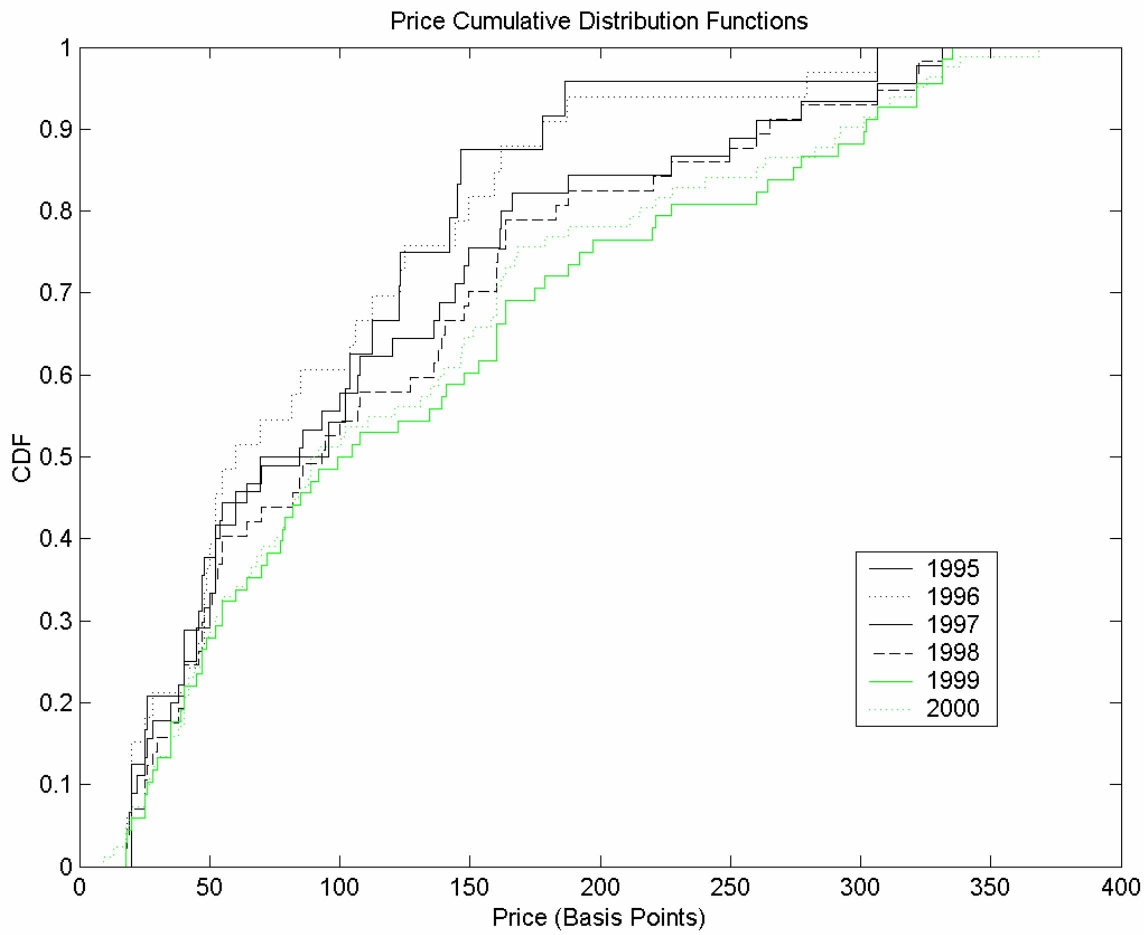

Price Cumulative Distribution Functions
CDF
Price (Basis Points)
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
0
50
100
150
200
250
300
350
400
1995
1996
1997
1998
1999
2000

Figure 3.

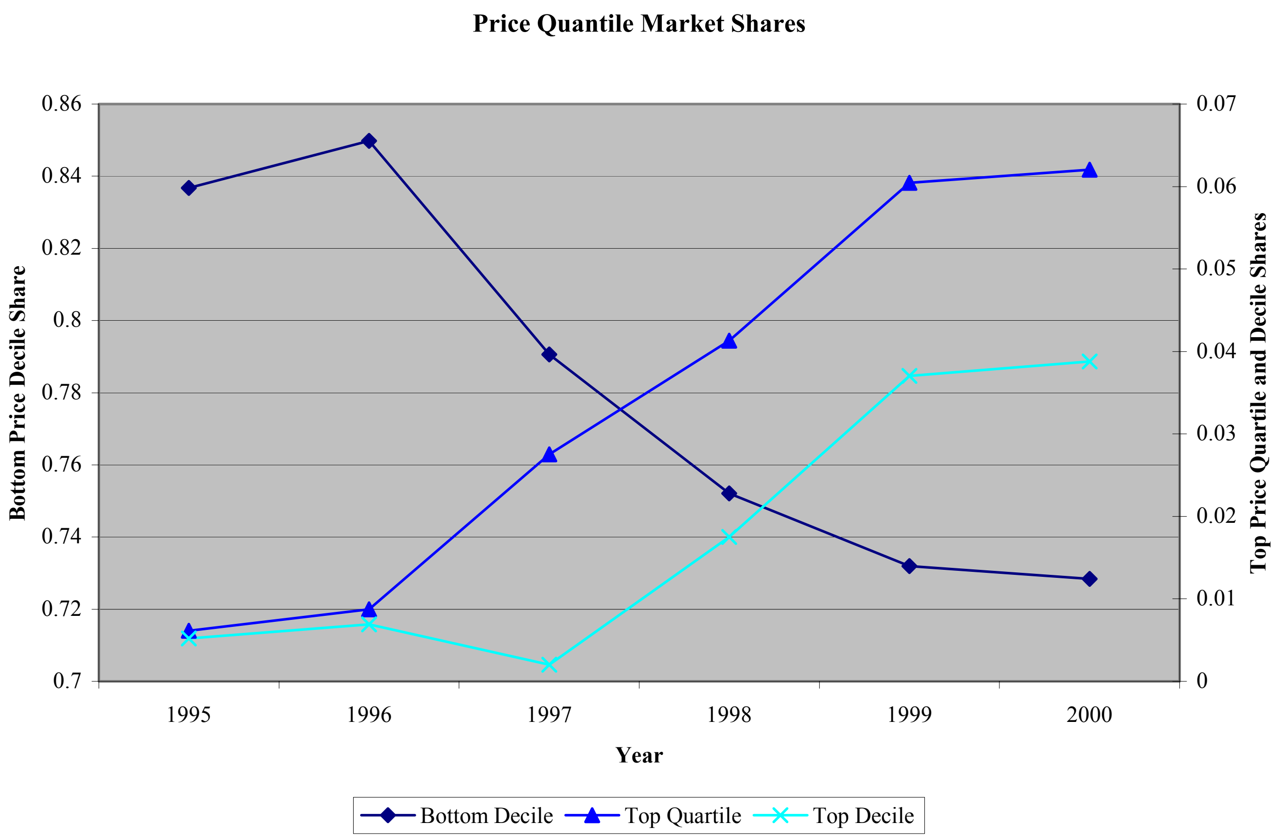
Price Quantile Market Shares
Bottom Price Decile Share
Top Price Quartile and Decile Shares
0.86
0.84
0.82
0.8
0.78
0.76
0.74
0.72
0.7
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
1995
1996
1997
1998
1999
2000
Year
Bottom Decile
Top Quartile
Top Decile

Figure 4.

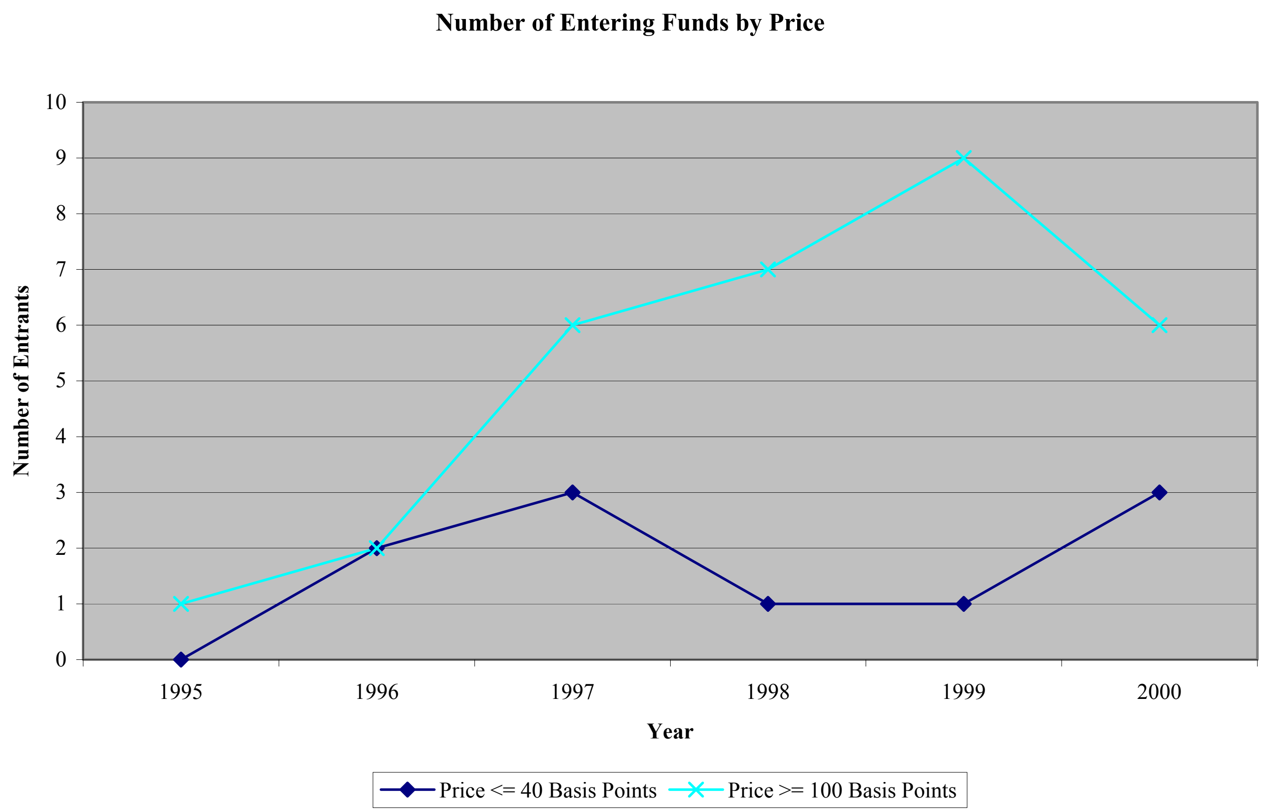

Number of Entering Funds by Price
Number of Entrants
10
9
8
7
6
5
4
3
2
1
0
1995
1996
1997
1998
1999
2000
Year
Price <= 40 Basis Points
Price >= 100 Basis Points

Figure 5

Log Price vs. Log Market Share, 2000
6
5.5
5
4.5
4
3.5
3
2.5
2
ln(Price in Basis Points)
-14
-12
-10
-8
-6
-4
-2
0
ln(Market Share in Year)

Figure 6

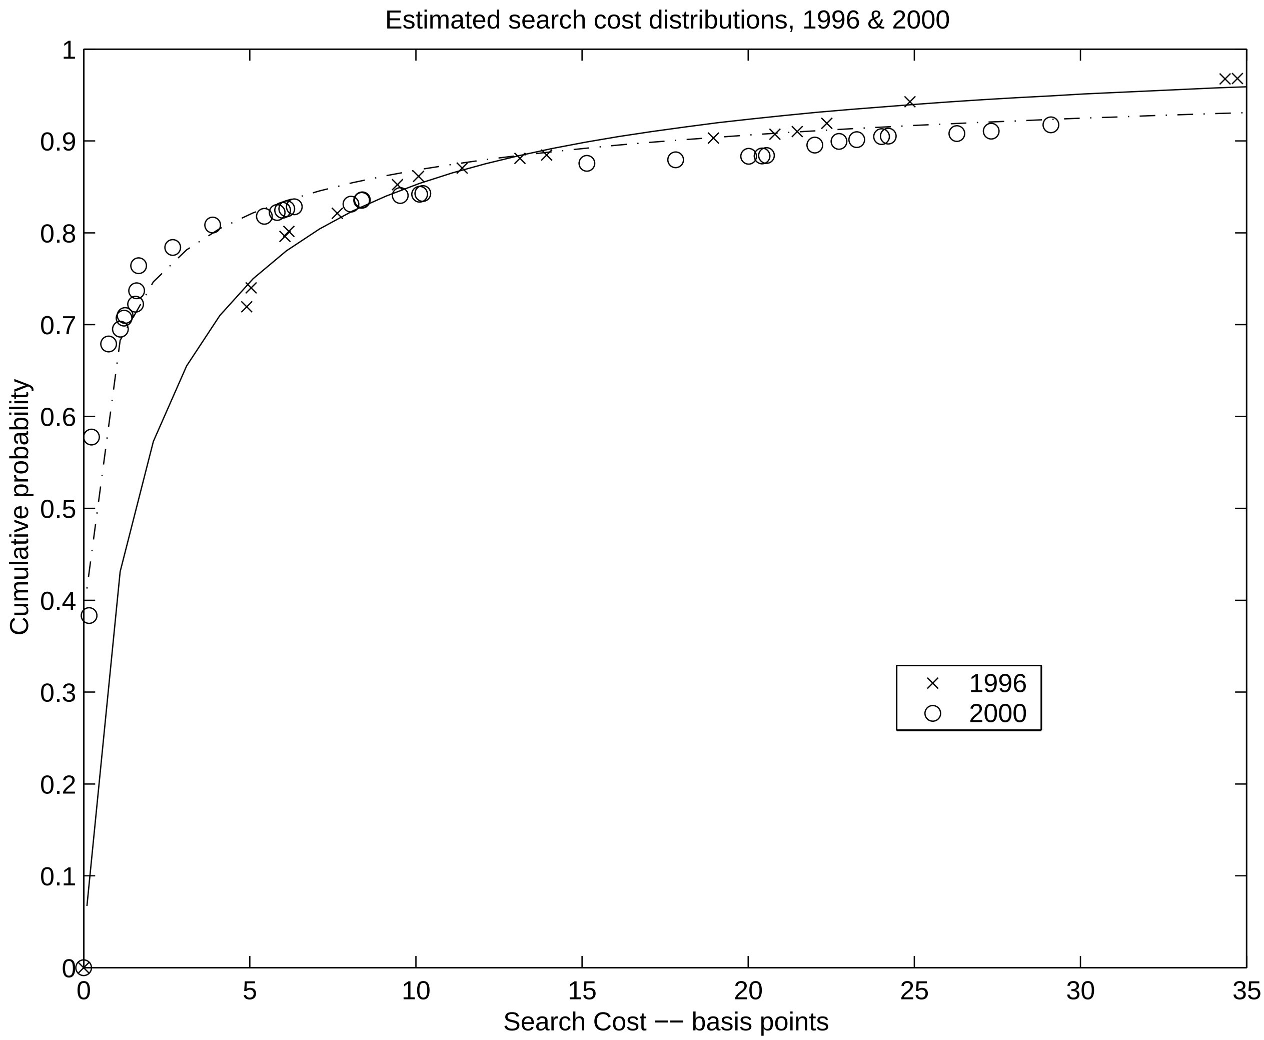

Notes: Solid line is parametric distribution for 1996, dashed line is parametric distribution for 2000.