Mutual Fund Families, Market Structure, and Fee Competition

Richard Grice* and Ahmed Guecioueur†

February 5, 2026

Abstract

We show that the positioning of fund *families* influences price competition between individual mutual *funds*. We document that investors consider a limited set of families, and model the equilibrium consequences: a competition network results, and funds have the incentive to internalize this market structure when setting fees. Calibrating our model with prospectus download data, we confirm that fees charged by seemingly homogeneous S&P 500 index funds are predicted by their families' positions in the consideration-based competition network. Counterfactual market structure analyses reveal that fees are at one-third the level of monopoly pricing but three times higher than competitive benchmarks.

Keywords: Competition Network, Index Funds, Fee Dispersion, EDGAR, Consideration
JEL Codes: G20, D40, G11, D12, L10

*NUS Business School, National University of Singapore. Email: rpgrice@nus.edu.sg
†University of Washington, Foster School of Business. Email: ahmedgc@uw.edu
First version: May 15, 2023. We are grateful to Simona Abis, Byeong-Je An, Frederico Belo, Philip Bond, Maxime Bonelli, Winston Wei Dou, Doron Levit, Anton Lines, Pierre Mabille, Massimo Massa, Joël Peress, Clemens Sialm, Maria Ana Vitorino and Ran Xing, as well as participants at the 2023 Financial Intermediation Research Society (FIRS) Conference, the 8th NSF Network Science and Economics Conference, the 14th Annual Hedge Fund Research Conference, the 2023 AFA PhD Poster Session, the 35th Australasian Finance and Banking Conference, the 11th Wharton-INSEAD Doctoral Consortium and at INSEAD for helpful comments. Guecioueur gratefully acknowledges the past support of the Henry Grunfeld Foundation.

1 Introduction

Fund families are a key feature of the asset management industry, with almost 800 complexes in operation (Investment Company Institute, 2025). Funds may organize themselves into families in part to share costs and exploit returns to scale. Yet families are also willing to *incur* costs to maintain distinct identities,[1] including advertising, and do so even in cases when a direct flow effect is absent (e.g., Gallaher, Kaniel, and Starks, 2015). Why?

In this paper, we argue that a *family*'s identity grants its offered *funds* a distinct source of market power when setting their fees. Family identities matter because they shape which funds investors consider, thereby influencing the competitive landscape that individual funds face. We analyze a mechanism in which investors do not search the entire universe of mutual funds when selecting among alternatives in a given category or style. Rather, an investor considers funds in that category that are offered by a small set of families that she knows, trusts, or otherwise finds salient. These limited and heterogeneous patterns of family consideration generate a *consideration-based competition network* among funds. When setting fees, funds internalize this market structure, so a fund's equilibrium fee reflects both its own performance and the performance of other funds offered by families that tend to be considered alongside it.

We formalize this mechanism in a new model of oligopolistic fee competition with networked demand. Investors are organized into segments according to the set of families they tend to consider simultaneously. Conditional on a segment's consideration set, a representative investor allocates wealth across the considered funds and a risk-free asset by solving a portfolio choice problem. Limited consideration implies that funds are perceived as imperfect substitutes, and each segment's demand depends on the expected risk-adjusted ex-fee returns of the funds it considers. Aggregating across segments yields a demand system in which pairwise demand linkages define a competition network with adjacency matrix $\boldsymbol{E}$, whose elements capture the strength of direct competition between funds.

On the supply side, each fund sets its fee taking its position in this network as given. A fund's pricing decision therefore internalizes a host of influences that lead investors to consider some families together but not others.[2] Funds choose fees to maximize fee revenue (since marginal costs are normalized). Because a higher fee charged by one fund reallocates demand toward its co-considered competitors, funds' best responses are strategic complements along

[1] For example, Franklin Templeton maintains 18 distinct brand names under which it offers funds to investors. Source: Franklin Templeton 2024 annual report.

[2] We later empirically calibrate the model using the *outcomes* of investor consideration processes, which allows us to remain agnostic about the exact drivers of investor consideration. As we discuss in Section 4.3, treating these consideration patterns as given is conservative for our empirical tests.

the edges of the competition network. We show that this fee-setting game admits a unique Nash equilibrium in which the vector of equilibrium fees $\boldsymbol{f}^*$ for the entire fund category is a closed-form function of the competition network, $\boldsymbol{f}^*(\boldsymbol{E})$. These equilibrium fees resemble a centrality measure: funds that optimally set higher fees are thus funds that are more central in the competition network, according to our model-implied measure.

In our framework, fee-setting decisions are made by *funds*, but these fees are shaped by *family* identities. One of our model's contributions is to highlight a direct analytical link between investors' consideration patterns (of families) and the market structure of funds within a given category, captured by the competition network $\boldsymbol{E}$:

$$\begin{array}{c}\text{Family Consideration}\\ \text{Patterns of Investors}\end{array} \quad \rightarrow \quad \begin{array}{c}\text{Market Structure Encoded}\\ \text{in Competition Network } \boldsymbol{E}\end{array} \quad \rightarrow \quad \begin{array}{c}\text{Fund Fees } \boldsymbol{f}^*(\boldsymbol{E})\\ \text{Set in Equilibrium}\end{array}$$

The model takes investor consideration patterns as an input and produces predictions of the equilibrium fees $\boldsymbol{f}^*$ levied by the entire cross-section of funds. Our model allows us to test for the presence of our proposed mechanism by calibrating its input to the observed consideration patterns of investors, thereby computing the family competition network $\widehat{\boldsymbol{E}}$ and finally a vector $\widehat{\boldsymbol{f}}$ of predicted fees for the cross-section of funds. Ultimately, the test of our model is to compare its predicted fees $\widehat{\boldsymbol{f}}$ to observed fees $\boldsymbol{f}^{obs}$.

We test whether family identities lend pricing power to their offered funds by taking our model to the data. The key empirical challenge is to measure the identities of which families investors consider together; these consideration patterns determine the market structure $\boldsymbol{E}$ in our model. We empirically measure family consideration using SEC EDGAR search outcome data: an investor's acquisition of mutual fund prospectus information reveals the families from which she considers potential choices. Our revealed consideration approach based on observable EDGAR search data avoids the need to make assumptions about which family attributes matter to investors, or to measure all such attributes. Similarly, because the model calibrates directly to this search outcomes dataset, there are no parameters for preferences, search costs, or behavioral frictions to estimate. We thus remain open-minded about the (potentially heterogeneous) drivers of individual investors' decisions, instead focusing our test on funds' fee-setting responses to any patterns of consideration.

Our descriptive analysis of the EDGAR data suggests family identities shape investors' consideration patterns. We document that mutual fund investors tend to consider a small and varied subset of families when deciding upon their initial investments: on average, an investor's consideration set contains only 2-3 families. Whatever the origin of these patterns of family consideration by investors, funds have an incentive to internalize them when setting

their fees.[3] Furthermore, the sets of considered families overlap so that all families are considered by at least some investors, leading to a high degree of heterogeneity in the composition of families across consideration sets.

As part of our mechanism test, we produce model-implied fees $\widehat{\boldsymbol{f}}$ by calibrating our model to its inputs; namely, observed fund returns, and the market structure $\boldsymbol{E}$ of the families that offer them, computed based on observed consideration sets from our SEC EDGAR data. The final step in our test is to compare these predictions to observable fees $\boldsymbol{f}^{obs}$.

To most sharply test our model's predicted fees, we compare them to the fees charged by S&P 500 index funds. We thus apply our framework to a single category of seemingly homogeneous funds that nevertheless continue to exhibit substantial price dispersion (as first documented by Hortaçsu and Syverson (2004) and Elton, Gruber, and Busse (2004)). To the extent that index funds are perceived as a single, undifferentiated category of products, they serve as a useful laboratory to assess the impact of positioning at the family level, which we hypothesize influences fund pricing power in general. We focus only on retail share classes when collecting observable fees $\boldsymbol{f}^{obs}$ for our comparison (since institutional fees may be negotiated bilaterally).

Comparing $\widehat{\boldsymbol{f}}$ to the observed fees $\boldsymbol{f}^{obs}$ levied by S&P 500 index mutual funds, we find a strong positive relationship. The cross-sectional association is statistically indistinguishable from one-to-one, even when controlling for fund size (a proxy for marginal costs) and prospectus downloads (a proxy for attention and reach). Furthermore, model-implied fees retain their strong predictive power even after controlling for a host of observable family attributes such as total assets, number of funds offered, and passive share; indeed, our model-implied measure almost entirely subsumes the ability of observable characteristics to predict fund fees.

The baseline calibration test therefore validates our paper's main hypothesis. Since the model – calibrated solely from family consideration patterns – accurately predicts fund fees, we conclude that family identities meaningfully influence funds' pricing decisions. Furthermore, family positions within the broader competitive landscape predict fees, rather than family attributes *per se*.

To isolate the role of investor consideration from other ingredients of the model, we complement our baseline test with a series of permutation analyses. These additional tests hold fixed simple features of the EDGAR data – such as how often each family is considered and the

[3]This observed outcome could be driven by investors engaging in a costly and preference-directed search process, or simply by following behavioral heuristics to allocate their limited attention (Barber and Odean, 2008). The features of a family that may attract certain investors are likely to be heterogeneous and potentially unobservable (e.g., family brand perceptions); our use of observed consideration sets allows us to avoid these measurement challenges. We investigate the composition of investors' consideration sets in Section 7.

distribution of investors' consideration set sizes – while randomly perturbing which families are co-considered with which others. Across three such designs the explanatory power of models calibrated on these permuted inputs fall markedly relative to the baseline, and the explanatory power generated by the observed pattern of consideration lies in the extreme right tail of the permutation distributions. These results show that the detailed structure of investors' consideration sets, not just overall reach or attention, is necessary to explain the cross-section of fees, and that the EDGAR-based consideration data are informative rather than being noise.

Our proposed mechanism is based on competition between funds (and their families) for the allocations of investors who can freely choose among multiple alternatives. To demonstrate this, we exclude consideration sets consisting of a single family – which include investors who are held captive within a family – then rerun our calibration to find the results almost unchanged. This analysis demonstrates that competitive pressures operate through the marginal investors who consider multiple families.

While our model is most relevant for the pricing of directly-sold funds, it retains significant predictive power for families that sell funds through indirect distribution channels, such as brokers. The correspondence between model-implied and observed fees is closest to one-for-one for those funds that do not charge any 12b-1 distribution fees, yet remains significant for funds that do. This suggests that even brokers and other intermediaries are likely to be influenced by family positioning.

We next use our model to quantify the degree of competition in the mutual fund industry by computing counterfactual fees under alternative market structures. We compare observed fees to three benchmarks, which we derive analytically: a monopoly scenario where each fund faces no competition, and two highly competitive scenarios where investors either uniformly consider all possible fund pairs or universally consider all available funds. Our counterfactual calibrations show that fund fees are approximately one-third of monopoly fees but roughly three times higher than fees under the competitive benchmarks. Thus, while competition among fund families constrains pricing, the competition remains imperfect.

Our primary focus is on the fee-setting behavior of mutual funds. We nevertheless complement our main findings on fund behavior with two additional analyses on investor behavior. First, we demonstrate that fund families are more frequently co-considered by investors when they are more similar in characteristic space; in particular, the passive share and the number of offered funds. This finding is consistent with observable family attributes influencing pricing power indirectly through their effect on investor consideration patterns, with families competing for attention primarily against nearby competitors in attribute space.

Second, we geo-locate investors to the county level, then use these geographic locations to infer the sociodemographic characteristics of investors. We are thus able to examine whether investor characteristics can shed light on the consideration set formation process. We find that various proxies for investor sophistication (such as education & income) predict a larger investor consideration set: investors who are more likely to be sophisticated tend to consider a larger set of families when acquiring fund prospectuses.

Related Literature. We contribute to a growing literature on the importance of fund families, the identities of which are a driver of investor fund flows (Sirri and Tufano, 1998; Sialm and Tham, 2016). Families actively compete to attract investor demand with their offerings (Massa, 2003; Nanda, Wang, and Zheng, 2004; Kostovetsky and Warner, 2020) or by affiliating with retirement plans (Pool, Sialm, and Stefanescu, 2016). We show that family identities affect funds' pricing power. This mechanism rationalizes the presence of a family-level component in individual fund fees (Iannotta and Navone, 2012).

Our paper is related to models that attribute fee dispersion to heterogeneity in investor search. Hortaçsu and Syverson (2004) show that variation in investors' search costs – equivalently, in the size of their consideration sets – can rationalize substantial dispersion in mutual fund fees, and Roussanov, Ruan, and Wei (2021) emphasize that marketing activities can influence the probability that a fund is searched. We build on these insights by using observed patterns of family *co*-consideration to construct an explicit competition network and by showing that a family's position in this network explains not only the extent of fee dispersion but also its cross-sectional structure: which specific families are able to sustain higher or lower fees.

Our paper is also related to a wider literature on competition and differentiation in the mutual fund industry, which has mostly focused on differentiation based on fund styles and attributes. As well as our focus on families (rather than funds) we introduce a new source of search outcome data for measuring market structure that can capture unobservable and heterogeneous drivers of consideration, thus complementing measures of fund similarity based on observable fund attributes such as portfolio overlap (Wahal and Wang, 2011), fund return styles (Hoberg, Kumar, and Prabhala, 2018) and prospectus textual similarity (Kostovetsky and Warner, 2020; Abis and Lines, 2024; Bonelli, Buyalskaya, and Yao, 2023).

Our findings also shed new light on the determinants of index fund fees, complementing work by Hortaçsu and Syverson (2004), Brown et al. (2023), and Loseto and Mainardi (2023). Simultaneously, we use index funds as a laboratory to cleanly identify broader empirical facts (e.g., Elton, Gruber, and Busse, 2004; DeHaan, Song, Xie, and Zhu, 2021; Gueciouer, 2023).

We contribute a novel network-based model of fund fee competition that maps heterogeneous investors' consideration patterns into a cross-sectional structure of strategically-set fees. In particular, we adapt the retail competition model of Grice (2023), although Pellegrino's (2025) oligopolistic general equilibrium model is also related. Our contribution explicitly embeds investor portfolio choice and its interplay with fees charged by investment vehicles, thus opening up a class of tractable oligopoly models for studying competition between financial intermediaries in other contexts.[4] We also contribute to a wider literature on competition networks (e.g., Hoberg and Phillips, 2016; Dou and Wu, 2023).

Outline. The paper proceeds as follows. Section 2 describes the institutional setting and data. The equilibrium model of fee competition between fund families is presented in Section 3. We then calibrate the model and compare its predictions to observed fees in Section 4. Section 5 relates our modeled direct channel to captive and intermediated sales channels. Section 6 derives, calibrates and compares counterfactual fees. Section 7 analyzes the drivers of investor consideration. Section 8 concludes. The Appendix contains all proofs, figures, and tables. An Internet Appendix contains additional supporting analyses.

2 Data on Investor Consideration Patterns

Before proceeding to our model, we present some motivating evidence showing that investors consider a limited and heterogenous set of families prior to making their investment allocations.

We use the observed downloads of users of the SEC's EDGAR website to measure the sets of fund families that mutual fund investors consider and evaluate when making portfolio allocation decisions.[5] Internet Appendix B describes the EDGAR website, its usage logs and our data preparation steps, which include the removal of automated downloads from our sample.

We interpret downloading a prospectus filing from EDGAR as evidence that a potential investor is actively evaluating whether to invest in a mutual fund offered by the fund family that filed the document. This active evaluation is consistent with a late-stage, high-attention evaluation step of the investment decision process, where investors thoroughly evaluate a set of considered alternatives, determine the extent to which the alternatives meet their investment goals, and make the allocation decision. Therefore, the prospectuses an investor downloads reveals her final *consideration set* of alternatives she perceives to be comparable.

[4]The study of dispersion in expenses across employer-sponsored retirement plans by Loseto (2023) is a step in this direction.

[5]Our approach is somewhat similar to Lee, Ma, and Wang (2015)'s use of listed company filing downloads from EDGAR for the purpose of grouping these firms into industries. We differ by focusing on *individual* investors' considered alternatives, rather than in aggregate, as well as our focus on mutual fund prospectus data.

We now describe the patterns of investor consideration revealed by the EDGAR usage data. As well as providing a snapshot of a point in the investment decision process of mutual fund investors, these patterns provide a guide for the model we will introduce in Section 3 to characterize investor decisions and the competition they induce between the funds offered by different families. We will later calibrate our model to these observed patterns of investor consideration.

[Insert Figure I around here]

Limited Consideration. Despite the ease with which investors can acquire information via EDGAR, the number of fund families considered by individual users is very small. Figure I plots the distribution of the number of fund families for which prospectuses were downloaded by individual EDGAR users for each year from 2006 to 2016. The median consideration set size in every year is one. Indeed, the whole distribution is highly concentrated around the value one with a long right tail: 99% of users evaluate fewer than a dozen fund families in any year. Further, the distribution is stable from year to year across our sample, despite the substantial growth in EDGAR users over the same period seen in Figure IA.5. This behavior shows investors severely limit their consideration of the universe of investment opportunities before seeking out and evaluating the detailed fund information contained in prospectuses.

[Insert Figure II around here]

Heterogeneous Consideration. The EDGAR usage data also show investors limit their consideration to different fund families – no family is considered by all investors, yet every family is considered by some investor. Figure II plots the distribution of the proportion of EDGAR users by whom each fund family is evaluated for each year from 2006 to 2016. The average fund family is evaluated by only 0.2% of potential investors. The most widely-considered fund family in each year is evaluated by between 4% and 11% of users. Thus users must be evaluating different sets of fund families prior to making their investments. In other words, investor consideration is not only limited but also *heterogeneous*.

[Insert Figure III around here]

Furthermore, while the investors' heterogeneous consideration induces sparse co-consideration of pairs of fund families – consistent with the low probability of any single family being considered – there is structure in the sparsity showing EDGAR users do not download prospectuses at random. Figure III plots the co-consideration incidence among two different samples of fund families in 2016. Panel (a) plots the co-consideration of 100 randomly selected fund families. Most of the 100 sampled families are co-considered with at least one other family;

a minority are not. Among the families that are co-considered, many have 'co-considerees' in common. That is, transitivity in the co-consideration of families appears common. Panel (b) provides further evidence of this pattern. It plots the co-consideration incidence of a sample of fund families that are co-considered with the BlackRock family. The patterns indicate a high incidence of transitivity: two families that are both co-considered with BlackRock are more likely to be co-considered with one another. This commonality suggests investors do not randomly draw families to consider, but rather select fund families that are similar along some key dimensions.

Investor Attributes. As an additional check on the population of investors using EDGAR, Internet Appendix C uses their locations to proxy for sociodemographic characteristics. We find a wide variety of investors acquire their prospectuses from EDGAR. Later analyses in Section 7.2 indicate that less sophisticated investors appear to have higher search and evaluation costs, and therefore consider fewer alternatives, as expected.

3 A Model Connecting Consideration and Fund Fees

We now develop an analytical model of the demand for and pricing of mutual funds offered within an investment category and purchased directly by investors. Our objective with the model is to characterize how investors' limited consideration of fund families induces competition between funds and thus generates fee dispersion. To this end, we try to strike a balance between a model that is rich enough to capture the key mechanism linking consideration and fee setting, yet simple enough that the effect of this mechanism is not obscured by other forces.

We therefore focus on investor and fund manager behavior delimited by three assumptions.[6] First, investors select a small subset of families to consider based on family-level attributes, and perceive the funds offered by these families in any particular category to be imperfect substitutes with respect to their investment objectives. We do not model the selection of families into consideration, just the choice of funds conditional on consideration. Second, investors form beliefs about fund returns and incorporate fees into their decision-making *after* selecting families to consider and obtaining past performance and fee information about the funds they offer in the category of interest. Thus we model the effect of fees on fund choice, not the effect of fees on consideration nor the effect of non-fee fund attributes.[7] And third, the

[6]We believe these assumptions provide reasonable descriptions of investor and manager behavior in most differentiated financial intermediaries. Consequently, we believe the model is suitable for characterizing the impact of limited consideration on fee competition in a range of financial markets, not just that for mutual funds.

[7]As discussed in Section 4.3, the assumption that funds take consideration sets as given when setting fees means our subsequent empirical analyses of fees will provide conservative tests of our model.

time horizon over which funds set fees is shorter than the one over which they make decisions determining the position of the fund – such as deciding on the mandates determining their returns – and thus set fees given their current position. Fees are therefore the sole strategic decision made by funds in our model.

Our model takes as primitive the network of competitive interactions among funds implied by investors' consideration sets and maps this network into a vector of equilibrium fees for all funds in the category. This equilibrium fee vector can be interpreted as a network-based measure of competitive centrality: funds that are more centrally positioned in the competition network optimally choose higher fees. We begin by describing the demand-side of the model and then derive the fund fees set in response.

3.1 Investors' Fund Allocation Problem

3.1.1 Investor Consideration

The demand-side of the model consists of a unit mass of investors looking to invest in a particular investment category. Each investor considers one or more of the funds offered in the category by the families $i = 1, ..., N$, which are collectively denoted by the set $\mathcal{N}$. We assume every family offers only one fund in the category of interest, and thus use the subscript i to refer interchangeably to a family and its fund. All the investors who consider the same set of funds $\mathcal{N}_m \subseteq \mathcal{N}$ are grouped together into a segment m. We index these segments by $m = 1, ..., M$ and collectively denote them by the set $\mathcal{M}$.[8] The size of each of these segments is θ_m, with $\sum_{m=1}^{M} \theta_m = 1$. We refer to the collective structure of investor segments and their consideration sets as the category's *pattern of consideration*:

$$\mathcal{P} \triangleq \{(\mathcal{N}_1, \theta_1), ..., (\mathcal{N}_M, \theta_M)\}. \tag{1}$$

3.1.2 Segment Preferences

We assume the collective decisions of the investors in segment m can be represented by the choices of a segment-level representative investor who solves a portfolio allocation problem with respect to the segment's considered funds $\mathcal{N}_m$. This segment-m investor first observes fees f_i and forms beliefs about returns r_i for funds $i \in \mathcal{N}_m$ before deciding on her allocation. Specifically, she forms beliefs about the expectation of returns $\eta_i \triangleq \mathbb{E}[r_i]$ and their variance

[8] Note that segments are uniquely distinguished by their consideration sets. Thus there are no two segments $m, m' \in \mathcal{M}$, $m \neq m'$, such that $\mathcal{N}_m = \mathcal{N}_{m'}$. However, the set of families considered by m can overlap with that considered by m' ($\mathcal{N}_m \cap \mathcal{N}_{m'} \neq \emptyset$), and even be a subset ($\mathcal{N}_m \subset \mathcal{N}_{m'}$). Finally, we assume every family with a fund in the category is considered by at least one investor segment ($\forall\, i \in \mathcal{N}$, $\exists\, m \in \mathcal{M}$ s.t. $i \in \mathcal{N}_m$).

$\sigma_i^2 \triangleq \mathbb{V}\text{ar}(r_i)$. She then allocates an investment budget Y_m among her considered funds and an outside risk-free asset $i = 0$ with the goal of maximizing returns net of fees at the end of a common investment horizon.

To determine her allocation over her considered funds, the representative investor compares standardized ex-fee returns:

$$R_i \triangleq \frac{r_i - f_i}{\sigma_i} \tag{2}$$

where the levied fee f_i is computed on the same basis as the return r_i.[9] By using these standardized ex-fee returns, the investor partially accounts for fund riskiness when deciding her allocation. However, to fully incorporate the notions of risk and diversification, she must also take return covariances into account. We assume a constant correlation model is sufficient to capture her beliefs about covariances:[10]

$$\text{Cov}(r_i, r_j) \triangleq \rho \sigma_i \sigma_j, \tag{3}$$

for all $i, j \in \mathcal{N}_m$, $i \neq j$, with $\rho \in (0, 1)$.

The segment-m investor allocates her investment budget Y_m to maximize the following objective, which resembles a classic mean-variance utility function with unitary risk aversion:[11]

$$U_m(x_{m0}, \boldsymbol{x}_m) \triangleq r_0 x_{m0} + \boldsymbol{\mu}_m^T \boldsymbol{x}_m - \frac{1}{2} \boldsymbol{x}_m^T \boldsymbol{V}_m \boldsymbol{x}_m \tag{4}$$

subject to the full-allocation constraint:

$$x_{m0} + \mathbf{1}_m^T \boldsymbol{x}_m = Y_m, \tag{5}$$

where $\boldsymbol{x}_m$ is an $N_m \times 1$ vector of allocations x_{mi} for $i \in \mathcal{N}_m$, x_{m0} is the allocation to the risk-free asset, which has return r_0, and $\mathbf{1}_m$ is an $N_m \times 1$ vector of ones. The vector $\boldsymbol{\mu}_m$ is an $N_m \times 1$

[9] Standardizing in this way ensures both that the second moment of returns is fully taken into account by the investor, and – along with the constant-correlation assumption in Eqn. (3) – that the representative investor's portfolio optimization problem in Eqn. (11) has a closed-form solution. It is this closed-form solution in Eqn. (13) that enables us to aggregate demand across segments and ultimately characterize equilibrium fees in closed-form in Proposition 1.

[10] This specification could also be interpreted as investors explicitly adopting a constant correlation model, either due to bounded rationality (to avoid the cognitively-expensive task of estimating $N_m(N_m - 1)$ return covariances) or a rational desire to reduce the input sensitivity of the portfolio optimization process.

[11] We use the following notation throughout the paper: a is a scalar, $\boldsymbol{a}$ is a vector with transpose $\boldsymbol{a}^T$ and elements a_i, and $\boldsymbol{A}$ is a (square) matrix with elements a_{ij}.

vector of expected standardized ex-fee returns with elements:

$$\mu_i \triangleq \mathbb{E}[R_i] = \mathbb{E}\left[\frac{r_i - f_i}{\sigma_i}\right] = \frac{\eta_i - f_i}{\sigma_i} \tag{6}$$

for $i \in \mathcal{N}_m$, and $\boldsymbol{V}_m$ is an $N_m \times N_m$ covariance matrix of standardized ex-fee returns R_i for $i \in \mathcal{N}_m$, with leading-diagonal elements:

$$v_{ii} \triangleq \mathbb{V}\text{ar}(R_i) = \mathbb{V}\text{ar}\left(\frac{r_i - f_i}{\sigma_i}\right) = 1 \tag{7}$$

and off-diagonal elements:

$$v_{ij} \triangleq \text{Cov}(R_i, R_j) = \frac{\text{Cov}(r_i, r_j)}{\sigma_i \sigma_j} = \rho \tag{8}$$

for every $i \neq j$. Substituting the binding full-allocation constraint into Eqn. (4) yields the optimization problem:

$$\max_{x_{m0}, \boldsymbol{x}_m} \quad r_0 Y_m + \left(\boldsymbol{\mu}_m - r_0 \mathbf{1}_m\right)^T \boldsymbol{x}_m - \frac{1}{2} \boldsymbol{x}_m^T \boldsymbol{V}_m \boldsymbol{x}_m. \tag{9}$$

Note that by modeling investors' choices over their considered funds with the portfolio allocation problem of a representative investor, we are assuming these funds are perceived to be imperfect substitutes at the segment level. This perception is captured in the negative, non-unitary cross derivative of the representative investor's objective function with respect to fund allocations (see Vives, 2000):

$$\frac{\partial^2 U_m}{\partial x_i \partial x_j} = -\rho \in (-1, 0)\,. \tag{10}$$

Given the idiosyncratic nature of individual investor preferences, we explicitly prefer this segment-level assumption to one of perfect substitutability (i.e. a cross-derivative exactly equal to -1), which would instead imply the entire segment decides to allocate their investment budget to the same fund.

3.1.3 Segment Allocation

The segment-m investor's problem in Eqn. (9) is globally concave, and thus has a unique interior solution given by:

$$\boldsymbol{x}_m^* = \boldsymbol{V}_m^{-1}\left(\boldsymbol{\mu}_m - r_0 \mathbf{1}_m\right). \tag{11}$$

Therefore, the allocation by segment m to each fund can be fully characterized by:

$$x_{mi}^* = \begin{cases} h_{mii}(\mu_i - r_0) + \sum_{j\in\mathcal{N}_m, j\neq i} h_{mij}(\mu_j - r_0) & \text{for } i \in \mathcal{N}_m \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

for every $m \in \mathcal{M}$ and $i \in \mathcal{N}$, where:

$$h_{mii} = \frac{1}{\rho - 1} \cdot \frac{-\big((N_m - 2)\rho + 1\big)}{(N_m - 1)\rho + 1} > 0 \tag{13a}$$

are the leading-diagonal elements of the matrix $\boldsymbol{H}_m \triangleq \boldsymbol{V}_m^{-1}$, and the off-diagonal elements are:

$$h_{mij} = \frac{1}{\rho - 1} \cdot \frac{\rho}{(N_m - 1)\rho + 1} < 0 \tag{13b}$$

for every $i \neq j$. Note that h_{mii} and h_{mij} depend only on the number of funds considered by segment m, not the identities of the funds i and j. Thus the influence of a co-considered fund j's standardized ex-fee returns on segment m's allocation x_{mi} mechanically decreases with each additional fund segment m considers.

3.2 Category Demand and Competition Structure

The characterization of segment allocations in closed-form makes it possible to aggregate fund demand across segments analytically. The demand for a fund from all investor segments is given by the following category-level demand system:

$$\begin{aligned} x_i &= \sum_{m\in\mathcal{M}} \theta_m x_{mi}^* \\ &= \sum_{m\in\mathcal{M}} \theta_m \left(\mathbb{1}_{\{i\in\mathcal{N}_m\}} \sum_{j\in\mathcal{N}_m} h_{mij}(\mu_j - r_0) \right) \quad \forall\, i \in \mathcal{N}. \end{aligned} \tag{14}$$

Aggregate demand can contain cross-price terms $\frac{\partial x_i}{\partial f_j}$ that are zero, and thus can be sparse.

In expression (14), the demand for family i depends through the μ_j terms on the fees of all the funds also considered by the investor segments considering i. If there is no segment $m \in \mathcal{M}$ considering both i and j, then $\frac{\partial x_i}{\partial f_j} = 0$ and i's demand does not depend on j's fees (and vice versa). In that case i and j do not directly compete for any investment dollars. However their fee decisions may still not be independent. If there exists a family k such that a segment m considers i and k, and another segment considers k and j, then i's demand will indirectly depend on j's fees through their impact on k's choice of fees. In this case we say that i and j compete indirectly. Therefore competition between funds for investor demand has a network structure.

Competition Network. The network structure of fund competition is represented by the graph $\mathcal{E} \equiv \{\mathcal{N}, \boldsymbol{E}\}$ with nodes $i, j \in \mathcal{N}$ and adjacency matrix $\boldsymbol{E}$ with weighted edges:[12]

$$e_{ij} = \sum_{m \in \mathcal{M}} \mathbb{1}_{\{i \in \mathcal{N}_m\}} \mathbb{1}_{\{j \in \mathcal{N}_m\}} \theta_m h_{mij}. \tag{15}$$

An edge $e_{ij} < 0$ indicates funds i and j directly compete. That is, the demand of each is directly affected by the fee decisions of the other. If $e_{ij} = 0$, then i and j do not directly compete. However they will still compete indirectly if they belong to the same network component. That is, if there exists a sequence of directly competing funds connecting i and j, then a fee change by one will instigate a cascade of competitive fee adjustments that will ultimately affect the other's demand.[13] If i and j compete neither directly nor indirectly, they must occupy distinct components of the graph $\boldsymbol{E}$, which constitute separate submarkets of the mutual fund industry. Note that $\boldsymbol{E}$ also includes self-referring edges e_{ii}. These edges capture the fact that the demand for a family is affected by its own fee decisions. The adjacency matrix $\boldsymbol{E}$ is non-singular.[14]

[Insert Figure IV around here]

An example of the mapping from a market's pattern of consideration to the networked competition between funds is depicted in Figure IV. The example has 4 funds and 3 investor segments, assumed to be of equal size. The pattern of consideration is depicted on the left by nodes representing funds collected into circles representing segments' consideration sets. The consideration sets overlap, connecting the funds considered by one segment with funds considered by another through the family in the intersection. Thus although family i is

[12] From here on in we will refer to this graph simply by the adjacency matrix $\boldsymbol{E}$ for notational simplicity.

[13] That is, funds i and j indirectly compete if they are connected by a path through the network $\boldsymbol{E}$: funds $\{i_0, ..., i_n\}$ such that $i_0 = i$, $i_n = j$, and $e_{i_l i_{l+1}} < 0$ for all $l = 0, ..., n-1$.

[14] It can be shown that the adjacency matrix $\boldsymbol{E}$ is strictly diagonally dominant by construction, by a similar argument to Grice (2023, Lemma 2). Therefore $\boldsymbol{E}$ is non-singular by the Lévy-Desplanques theorem. A matrix $\boldsymbol{A}$ is strictly row-wise diagonally dominant if every row i satisfies the condition: $|a_{ii}| > \sum_{j \neq i} |a_{ij}|$.

only co-considered with family j by segment m, it is connected to the other two funds by the consideration of j. The pattern of consideration is converted into the corresponding competition network on the right. Edges between nodes depict direct competition between co-considered funds, with weights depicted by edge thickness (self-edges are not depicted for this example). Segments considering more funds correspond to thinner edges between the funds, because there will be less pairwise substitution when a segment allocates its budget over more funds. Thus thinner edges capture the fact that pairwise competition will be less intense between funds considered by segments with larger consideration sets. This inverse relationship between edge weight and consideration set size follows from the inverse relationship between h_{mij} and N_m in Equation (13).

The edges in Eqn. (15) can be used to parsimoniously re-express the fund demand in Eqn. (14), allowing the order of summation to be switched and the market's demand system to be expressed as:

$$\begin{aligned} x_i &= \sum_{m\in\mathcal{M}} \theta_m \left(\mathbb{1}_{\{i\in\mathcal{N}_m\}} \sum_{j\in\mathcal{N}_m} h_{mij}(\mu_j - r_0) \right) \\ &= \sum_{m\in\mathcal{M}} \theta_m \mathbb{1}_{\{i\in\mathcal{N}_m\}} \sum_{j\in\mathcal{N}} \mathbb{1}_{\{j\in\mathcal{N}_m\}} h_{mij}(\mu_j - r_0) \\ &= \sum_{j\in\mathcal{N}} \sum_{m\in\mathcal{M}} \mathbb{1}_{\{i\in\mathcal{N}_m\}} \mathbb{1}_{\{j\in\mathcal{N}_m\}} \theta_m h_{mij}(\mu_j - r_0) \\ &= \sum_{j\in\mathcal{N}} e_{ij}(\mu_j - r_0) \end{aligned} \tag{16}$$

for every $i \in \mathcal{N}$. This demand can also be written in matrix form using $\boldsymbol{E}$:

$$\boldsymbol{x} = \boldsymbol{E}(\boldsymbol{\mu} - r_0\mathbf{1}). \tag{17}$$

This re-expression reveals how edge weights summarize the pairwise co-consideration and substitution between funds:

$$\frac{\partial x_i}{\partial f_j} = e_{ij} \cdot \frac{\partial \mu_j}{\partial f_j} = -\frac{e_{ij}}{\sigma_j}. \tag{18}$$

The self-referring edges e_{ii} capture the overall propensity of the investors considering fund i to reallocate investment away from i to both other considered families and the risk-free asset when i's fee increases. The edges e_{ij} for $j \neq i$ capture the propensity of those investors to reallocate their investment from i to j. The primary determinant of this propensity will be the average h_{mij} in the segments considering funds i and j. As noted above, h_{mij} for a segment m is

itself primarily determined by N_m, the number of funds considered by that segment. Therefore, all else equal, the fewer other funds considered by the investors considering funds i and j, the greater will be the edge e_{ij}. In the extreme case that funds i and j are never considered in conjunction with other families, they will form a duopoly, and j will be i's only competitor for the allocations of its investors.

3.3 Funds' Profit Maximization Problem

On the supply-side of the model are the funds offered in the category of interest by the families $i = 1, ..., N$. As noted above, over the time horizon in which these funds set their fees, their families' positioning, their investment mandates, and the other attributes determining investor consideration are fixed. Thus the funds' only strategic decision is their fee, which they set knowing the size of the investor segments considering them and which other funds those segments consider.[15] That is, when setting fees, funds are aware of the demand they will face for any vector $\boldsymbol{f}$ of market fees.

We assume all the funds in the category face a common constant marginal cost, which we normalize to zero.[16] In addition, we assume the funds operate independently of any funds their families might offer in other categories, and so have the sole objective of optimizing the fee revenues earned from their own investors. Therefore, each fund $i \in \mathcal{N}$ sets its fee f_i to maximize profit:

$$\max_{f_i} \ \pi_i(f_i, \boldsymbol{f}_{-i}) \triangleq f_i \cdot x_i(\boldsymbol{f}). \tag{19}$$

This quadratic profit function is concave in f_i, and thus a fund i's profit maximizing choice of fee f_i^* is characterized by the following first order condition:

$$\frac{\partial \pi_i}{\partial f_i}\Big|_{f_i^*} = \frac{e_{ii}}{\sigma_i}\Big(\eta_i - 2f_i^* - \sigma_i r_0\Big) + \sum_{j \neq i} \frac{e_{ij}}{\sigma_j}\Big(\eta_j - f_j - \sigma_j r_0\Big) = 0. \tag{20}$$

[15] However, while funds know what proportion of their investors belong to each segment and what set of funds those segments consider – for example as the result of market research – they cannot know the segment to which any individual investor belongs. Therefore they are unable to price discriminate by charging different fees to investor segments and must instead charge all their prospective investors the same fee.

[16] It is possible to extend the model to incorporate non-uniform marginal costs. However because costs are not observed, they would become free parameters that would require structural estimation to recover. The model version we present characterizes demand-driven (not cost-driven) dispersion in fees, which is our focus in this paper.

Therefore funds' fees will be determined by the following system of linear best reply functions:

$$f_i^*(\boldsymbol{f}_{-i}) = \frac{1}{2}\Big(\eta_i - \sigma_i r_0 + \sum_{j\in\mathcal{N}, j\neq i} \frac{e_{ij}}{e_{ii}}\frac{\sigma_i}{\sigma_j}\Big(\eta_j - f_j - \sigma_j r_0\Big)\Big) \tag{21}$$

for every fund $i \in \mathcal{N}$.

These best reply functions imply the fees of directly competing funds are strategic complements (see Vives, 2000):

$$\frac{\partial f_i^*}{\partial f_j} = -\frac{1}{2}\frac{e_{ij}}{e_{ii}}\frac{\sigma_i}{\sigma_j} > 0. \tag{22}$$

Thus fund i best replies to a fee adjustment by competitor j by adjusting its fees in the same direction, and funds' fees are increasing in the fees of their direct competitors. If i and j are not direct competitors, fund i does not directly reply to a fee adjustment by j because its demand is not directly affected by j's fees (i.e., $e_{ij} = 0$).

3.4 Equilibrium Fund Fees

The best reply functions in Equation (21) characterize how direct competitors will respond to a fee change by fund i. However, as discussed in Section 3.2, the effect of such a fee change will not be limited to i's direct competitors. When a direct competitor j responds to i's change with a fee adjustment of their own, they will cause a reallocation of investment away from fund j to *its* direct competitors – including back to i – shifting their demand curves and thus the fees at which they maximize profits. The adjustments they then make in response will trigger yet another iteration of investment reallocations and fee adjustments, and so on. Therefore, i's initial fee change will trigger a cascade of fee adjustments traveling along the edges between direct competitors in the network $\boldsymbol{E}$ to also eventually affect all of i's indirect competitors.

The eventual effect of this fee-adjustment cascade is given by the equilibrium it converges to as it reverberates about the market. However, for the reverberations to peter out and eventually converge, the magnitude of the fee adjustments must decrease with each iteration of the cascade. To ensure that they do, we assume fund returns satisfy the following technical sufficient condition.

Assumption 1 (Strict diagonal dominance). *The returns of funds co-considered by investors are such that the adjacency matrix $\boldsymbol{E}$ is strictly (row-wise) diagonally dominant with weights $\frac{\sigma_i}{\eta_i - \sigma_i r_0}$.*

That is:

$$|e_{ii}|\frac{\eta_i - \sigma_i r_0}{\sigma_i} > \sum_{j \neq i} |e_{ij}|\frac{\eta_j - \sigma_j r_0}{\sigma_j} \tag{23}$$

for all $i \in \mathcal{N}$.

The requirement in Assumption 1 amounts to a restriction on the difference in standardized excess returns of funds (in this same category) co-considered by investors. As noted in Section 3.2, given unitary weights the adjacency matrix $\boldsymbol{E}$ is strictly diagonally dominant by construction. In addition, $e_{ij} = 0$ for every fund j not co-considered with fund i. Thus for $\boldsymbol{E}$ to also be strictly diagonally dominant with the weights in (23), a fund's standardized excess returns must not be too different to the standardized excess returns of the other funds with which it is considered by investors. In particular, the standardized excess returns of funds i and j must be closer the closer are i and j to a duopoly in which all the investors who consider i consider j and only j as an alternative to i (and vice versa).

Given Assumption 1, the adjustments funds make in each iteration of the fee-adjustment cascade will decay to zero in the limit and fees will settle at an equilibrium. In the following proposition we prove this equilibrium is unique, and explicitly characterize the associated equilibrium fund fees. The proof makes use of Banach's Fixed Point Theorem and is provided in the Appendix.

Proposition 1 (Nash equilibrium). *There exists a unique Nash equilibrium to the fee-setting game between funds characterized by non-zero fees:*

$$\boldsymbol{f}^* = \left(\mathbf{I} + \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\right)^{-1} \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\left(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\right), \tag{24}$$

where $\boldsymbol{\eta}$ *is a vector with elements* $\eta_i \triangleq \mathbb{E}[r_i]$, $\operatorname{diag}(\boldsymbol{S})$ *is the* $N \times N$ *diagonal matrix with the elements* σ_i *on the leading-diagonal and* 0 *on the off-diagonals, and* $\widetilde{\boldsymbol{E}} \triangleq \boldsymbol{E}\operatorname{diag}(\boldsymbol{S})^{-1}$ *is the* $N \times N$ *matrix with the elements* $\widetilde{e}_{ij} \triangleq \frac{e_{ij}}{\sigma_j}$, *and thus* $\operatorname{diag}(\widetilde{\boldsymbol{E}}) = \operatorname{diag}(\boldsymbol{E})\operatorname{diag}(\boldsymbol{S})^{-1}$.

Equilibrium Fees Depend on the Performance of Other Funds. One immediate implication of our equilibrium is that funds' fees are determined not just by funds' own performance (i.e. their excess returns and return riskiness), but also by the performance of the other funds in the market. Specifically, fund i's fees are a weighted sum of all funds' performance – including their own performance – where the weights are given by the ith row of the matrix $\left(\mathbf{I} + \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\right)^{-1} \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}$. These weights summarize each fund's position in the market's

networked competition structure, providing a measure of their centrality with a close resemblance to Katz-Bonacich centrality (Katz, 1953; Bonacich, 1972; Bonacich, 1987), which is a common measure of player importance in network games with strategic complementarities (Ballester, Calvó-Armengol, and Zenou, 2006). In the equilibrium expression in Eqn. (24), the relevant network is $\widetilde{\boldsymbol{E}}$, whose edges $\widetilde{e}_{ij} = -\frac{\partial x_i}{\partial f_j}$ capture the absolute responsiveness of families' demand to changes in their direct competitors' fees.

Therefore, we can understand the role of competition in determining fund fees as follows. When funds change their fees, say because of an industry-level shock to the demand for mutual funds, investors re-optimize their portfolio allocations. The resulting investment re-allocations, and competitive fee adjustments they prompt in response, will cascade through the market from competitor to competitor. Proposition 1 says these cascades will agglomerate as they reverberate through the market, concentrating on the most central funds in the networked competition structure as they peter out. Thus funds' equilibrium fees depend on not only the fees and performance of other funds with which they are co-considered, but on the ex-fee performance of all funds in the same component of the market's consideration-induced competition network.[17]

The expressions in Proposition 1 also allow us to characterize the equilibrium outcomes arising from competition between funds for investors' allocations under various frictions and intensities of competition and co-consideration. We derive such special cases in Section 6.

4 Main Calibration

In this section, we calibrate the model and compare its predicted fees to observed fees. We thus test our primary hypothesis that families' positions in the consideration-based competition network determine the fee-setting power of the funds that they offer.

Our calibration closely respects the following two key assumptions of the model: (1) Since we model competition between funds in a uniform category, our calibration will focus on S&P 500 index mutual funds – a uniform catgory of funds with quasi-identical portfolio attributes and strategies. (2) Since our model assumes investors are unconstrained in directing their allocations, we always condition on non-institutional share classes in these S&P 500 funds.

Crucially, our calibration employs *family* consideration set data from SEC EDGAR to try to explain the pricing of *funds* within a uniform category. Therefore, our calibration serves as a

[17]The one exception to this rule is when a fund enjoys a monopoly over its investors. As we derive in Section 6.1, in this case a fund's fee equilibrium will be fully determined by its performance alone.

test of the paper's principal argument that family positioning influences their offered funds' fee-setting decisions. Consistent with this hypothesis, we find the calibrated model's predictions are positively correlated with observed fees.

4.1 Calibration Data

Fees, Returns and Other Characteristics. We use the CRSP Survivor-Bias-Free US Mutual Fund Database to measure mutual fund fees and returns. Our fee calculations follow the extant literature, in particular the procedure laid out in Roussanov, Ruan, and Wei (2021, Internet Appendix, pp. 15-16). In brief, this procedure defines total fees as the sum of annualized front loads and expense ratios (which include management fees, among other expenses). Calculations are performed at the share class level, for the last available month in a calendar year. Front loads (which are levied at the point of purchase) are annualized by dividing by 7. If different front loads may be charged (by a pricing schedule), the largest front load is used. "Effective 12b-1 fees" are defined as the sum of front loads and actual 12b-1 fees; a ceiling of 1% is then set on effective 12b-1 fees. Expense ratios greater than 5% are assumed to be invalid. Missing or negative front loads are assumed to be zero. Finally, missing or negative expense ratios are replaced with their lifetime means. Note that CRSP expense ratios and management fees incorporate the effect of any fee waivers and rebates (Christoffersen, 2001).

We also follow the literature in calculating gross-of-fee monthly returns: these are computed by adding one-twelfth of the expense ratio to net-of-fee monthly returns ("mret").

As described earlier, we match share class records to SEC CIKs and fund Series IDs using the "CRSP_cik_map" table from the CRSP Mutual Fund Database. We aggregate fees and returns across share classes as a Total Net Assets (TNA)-weighted average, as in Roussanov, Ruan, and Wei (2021, Internet Appendix, pp. 16). In doing so, we use only share classes that are flagged by the CRSP database as non-institutional, since institutional fees may be set through a bilateral bargaining process.

S&P 500 Index Funds, and Aggregation. As mentioned, our calibration focuses on mutual funds that track the S&P 500 index. We aggregate fees and returns to the fund family level, weighting by TNA: the purpose of this aggregation is to ensure that the unit of analysis consistently remains at the family level, to avoid any potential econometric bias in our panel regression test.[18] Our procedure for identifying S&P 500 index funds, and a description of the

[18]Family-level quantities and fund-level quantities will be identical for those families that offer only a single S&P 500 tracker fund. Note that Massa (2003) and Khorana and Servaes (2012) similarly aggregate feees to the family level in their own empirical analyses.

cross-section of fees charged in this category, are described in Internet Appendix A.

4.2 Calibration Procedure

We calibrate the model and calculate the fees it predicts, $\widehat{f}$, by entering calibrated values into the right-hand side of Eqn. (24):

$$\widehat{f} = \left(\mathbf{I} + \text{diag}(\widetilde{E})^{-1}\widetilde{E}\right)^{-1} \text{diag}(\widetilde{E})^{-1}\widetilde{E}\left(\widehat{\eta} - \widehat{r}_0 \,\text{diag}(\widehat{S})\mathbf{1}\right), \tag{25}$$

where the vector $\widehat{\eta}$ contains investors' expectations about returns, $\widehat{\mathbb{E}}[r_i]$, the diagonal matrix $\text{diag}(\widehat{S})$ contains investors' perceptions about the standard deviation of returns, $\widehat{\sigma}_i$, the matrix $\widehat{E}$ contains the network ties between competitors emerging from investors' pattern of consideration and beliefs, with $\widetilde{E} \equiv \widehat{E}\,\text{diag}(\widehat{S})^{-1}$ by definition, and $\widehat{r}_0$ is the riskless rate of return. We calculate predicted fees for each year in our sample period of 2006 to 20016.

Calibrating Return Expectations. We calibrate expected returns at the fund family level to be consistent with the level of aggregation at which investors' consideration sets are observed. Expected returns $\widehat{\mathbb{E}}[r_i]$ and standard deviations $\widehat{\sigma}_i$ are calibrated as the sample analogues calculated from each family's monthly gross returns over the most recent year. We take as the riskless rate of return $\widehat{r}_0$ the estimate provided by Kenneth French, which is also expressed on a monthly basis.

We also calibrate investors' beliefs about the covariance of the returns of any two considered families, as defined in Eqn. (3). Return covariance beliefs $\widehat{\rho}$ are calibrated from monthly returns over the same lookback period as return expectations:[19]

$$\widehat{\rho} = \frac{1}{N(N-1)} \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{N}, j \neq i} \frac{\widehat{\text{Cov}}(r_i, r_j)}{\widehat{\sigma}_i \widehat{\sigma}_j}. \tag{26}$$

All calculations are repeated for every year in $t \in [2006, 2016]$.

Competition Network Calibration. The family-to-family competitive ties $\widehat{e}_{ij}$ in the adjacency matrix $\widehat{E}$ are calibrated through a weighted aggregation of investors' consideration sets according to the definition in Eqn. (15):

$$\widehat{e}_{ij} = \sum_{m \in \widehat{\mathcal{M}}} \mathbb{1}_{\{i \in \widehat{\mathcal{N}}_m\}} \mathbb{1}_{\{j \in \widehat{\mathcal{N}}_m\}} \widehat{\theta}_m \widehat{h}_{mij}, \tag{27}$$

[19] This approach to calculating ρ follows the empirical portfolio choice & asset pricing literature. Various studies in that literature, including Elton and Gruber (1973) and Elton, Gruber, and Urich (1978), have found this approach produces empirically accurate results.

where $\widehat{\mathcal{M}}$ is a segmentation of investors into groups based on the sets of fund families they consider, $\widehat{\mathcal{N}}_m$ is the consideration set common to all investors in segment m, $\widehat{\theta}_m$ is the proportion of investors belonging to segment m, and:

$$\widehat{h}_{mij} = \begin{cases} \frac{1}{\widehat{\rho}-1} \cdot \frac{-\left((\widehat{N}_m-2)\widehat{\rho}+1\right)}{(\widehat{N}_m-1)\widehat{\rho}+1} & \text{if } i = j \\ \frac{1}{\widehat{\rho}-1} \cdot \frac{\widehat{\rho}}{(\widehat{N}_m-1)\widehat{\rho}+1} & \text{if } i \neq j \end{cases} \tag{28}$$

where $\widehat{N}_m = |\widehat{\mathcal{N}}_m|$ is the number of families considered by segment m, and $\widehat{\rho}$ is investors' beliefs about return covariances.

The pattern of consideration $\widehat{\mathcal{M}}$ for each year t is formed from prospectus downloads by investors via the SEC's EDGAR website over the same lookback period with which we calibrate investors' beliefs about fund family returns. The weights $\widehat{\theta}_m$ are then calculated by dividing the number of downloaders with measured consideration set $\widehat{\mathcal{N}}_m$ by the total number of downloaders in year t.

Finally, to ensure the convergence condition in Assumption 1 is satisfied, we constrain each element of the vector $\text{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\left(\widehat{\boldsymbol{\eta}} - \widehat{r}_0 \,\text{diag}(\widehat{\boldsymbol{S}})\mathbf{1}\right)$ to be non-negative.

4.3 Main Calibration Results

Calibrated Competition Network. We begin with a family competition network constructed from the family-level consideration data described in Section 2. This calibrated network $\widehat{\boldsymbol{E}}$ forms a single network component, and is densely connected – of all possible edges, 93.7% are present in the calibrated network.[20] Further, 12.7% of families are directly connected to – and thus directly compete with – all other families. This density of the mutual fund industry's competition network is due to the aggregation of the sparse individual consideration sets documented in Section 2. The vast majority of investors consider investing in only a handful of fund families, and in 2016 no family is considered by more than 4% of investors. And yet the heterogeneity across investors in the families they consider connects the industry together into one large, densely-connected competition network. Note that this lack of sparsity does not imply uniform competition between fund families – the weights of the edges connecting families vary, corresponding to varying intensities of substitution and competition between direct competitors.

[20] Internet Appendix D includes visualizations of the calibrated competition network and its density, and reports key calibrated parameters.

Predicted vs. Observed Fees. We now show that the fees predicted by the calibrated model are significant predictors of observed fund fees. Our calibrated model makes a unique prediction for each year t from 2006 through 2016 and each family i. We compare these predicted fees $\widehat{f}_{it}$ with observed family average fees f_{it} by regressing the latter against the former. Table I presents the results of this post-calibration analysis.

[Insert Table I around here]

All specifications in Table I show the fees predicted by our calibration are positively and significantly associated with observed fees, even with standard errors clustered at both the year and the fund family level. A t-test of the coefficient estimates reveals these values are not significantly different to unity (p-values 0.67, 0.42 & 0.42, respectively); the observed cross-sectional fee dispersion thus matches our model-predicted cross-sectional fee dispersion in magnitude as well as direction. Time effects are included in all specifications to control for industry-wide fee shocks, and to focus on cross-sectional variation in our model output.

The baseline comparison in column (1) is augmented in column (2) with additional fund characteristics: log size (total net assets) is a control for scale effects that could influence the marginal cost of offering an S&P 500 fund, and the log number of unique prospectus downloaders in a year is a control for the level of recognition and awareness of the fund family, and its reach among the universe of potential consumers. The main estimated coefficient is little-changed by the inclusion of these controls.

Signing Potential Bias. One potential concern is that the positive association between model-implied fees and observed fees in Table I might be contaminated by feedback from fees into consideration sets. In particular, if investors use observed fees when deciding which families' prospectuses to download, then the pattern of EDGAR downloads could itself be influenced by current fees, and the regression shown in this table would implicitly allow observed fees to enter on both sides of the regression. Our maintained assumption in the model is that consideration sets are exogenous to fees, and furthermore we do not use fee information when calibrating our model.

Nevertheless, it is useful to think through the direction of any residual bias if some investors' consideration sets did in fact respond to fees. Suppose lower-fee funds attract more attention and are therefore included in more investors' consideration sets. To a first order, such highly-considered funds would be predicted to levy higher equilibrium fees in response to this demand. Because lower-fee funds would have higher model-predicted fees, this mechanism would push the slope coefficient downward, introducing a negative bias into our regression of observed

versus predicted fees. Therefore, such endogeneity would work against us finding a positive and significant association. While we cannot rule out such bias completely, the fact that the estimated coefficient is stable when we control explicitly for total downloads in column (2) suggests that any remaining bias is unlikely to be large.

Role of Family Attributes. Column (3) of Table I further extends the specification to include observable family attributes: the (log of) family total net assets,[21] the family's share of assets following an equity strategy, the family's share of assets following a passive strategy, the number of mutual funds offered by the family, the diversity of asset classes offered by the family (computed as the Herfindahl-Hirschman Index of assets per asset class at the family level) and finally the family-level diversity of their offered funds' returns.[22] In computing these measures, we use asset class and strategy flags from the CRSP Mutual Fund Database, and extract the number of offered funds from regulatory filings. Internet Appendix A.3 summarizes these fund- and family-level characteristics.

The main coefficient from specification (3) remains similar when family characteristics are controlled for. Family-level attributes do not drive out the predictive power of our calibration-implied fees; indeed, the opposite holds: the only relationship that survives is a weakly negative association with family-level (log) assets under management.

These results suggest that model-implied fees cannot be reduced to a simple observable characteristic of the fund family (e.g., size, age, or product mix). In our framework the calibrated vector of model-implied fees summarizes each family's centrality in the competition network implied by investors' revealed consideration sets, aggregating the influence of many observable and unobservable attributes on pricing power. The fact that adding a rich set of family-level observables barely affects the coefficient on $\widehat{f}$, and that these observables themselves explain little of the remaining variation, indicates that no single family characteristic is sufficient to capture this notion of centrality.

Indeed, the strength of our revealed-consideration approach is to allow data recording which families investors co-consider to discipline the measure of competitive position directly, without ex ante assumptions of what attributes *ought* to matter. The robust predictive power of the model-implied fees $\widehat{f}$ supports our argument that they capture any observable and unobservable

[21] Note that Sirri and Tufano (1998, pp. 1608) call this variable the "fund complex size" and interpret it as a proxy for search costs.

[22] This family-level attribute is calculated as the generalized variance (Wilks, 1932) of the returns of the individual funds offered by the family. Specifically, it is the log of the determinant of the variance-covariance matrix of the monthly returns of offered mutual funds over a rolling 5-year window. Higher values denote a greater generalized variance of fund-level returns offered by the fund family, and thus a more diverse product mix.

attributes in the eyes of investors that influence equilibrium fees. This includes unobservable attributes such as brand capital (Sialm and Tham, 2016; Belo, Lin, and Vitorino, 2014) that are challenging to measure as a standalone variable in the context of mutual fund families.

Altogether, our findings indicate our approach of combining investor consideration data with a portfolio allocation model to characterize demand successfully captures the competitive landscape moderating competition between mutual fund families. A model of family fee competition based on this characterization of demand predicts equilibrium fees that explain a significant proportion of the cross-sectional dispersion in family fees. Further, this prediction explains not just the extent of price dispersion as in Hortaçsu and Syverson (2004), but the actual cross-sectional structure. Thus the model predicts not just why some fund families can maintain a higher fee level in the face of competition, but which fund families.

4.4 Permutation Tests

One potential concern could be that non-consideration-related features of our model could be driving the explanatory power, leading us to unjustified conclusions. For example, if funds competed without the constraints imposed by family co-consideration, we might still expect the ex-fee returns of other funds in their category to play *some* role in their fee-setting. And because those returns enter our model through $\hat{\eta}$, $\hat{\sigma}$ and $\hat{\rho}$, the model's predictions may correlate with observed fees despite the observed pattern of consideration $\mathcal{P}^{obs}$ incorrectly constraining how our model maps $\hat{\eta}$, $\hat{\sigma}$ and $\hat{\rho}$ into a prediction about fees. Another (extreme) example of mis-specification would be the example of EDGAR-inferred consideration patterns being mere noise.

To address these concerns, we conduct three permutation tests to verify that our model's predictive power stems *specifically* from the observed pattern of family co-consideration, rather than other features of the data, such as fund returns, or the absolute level of attention revealed by prospectus download counts that is not based on the competition networks that we infer.

In each test, we recompute the model-implied equilibrium fees $\widehat{\boldsymbol{f}}$ 10,000 times using 10,000 randomly permuted consideration patterns that randomize over a particular feature of the data. Each of our three tests randomizes over a distinct feature of the data; the randomization effectively breaks the informativeness of this particular feature, and should decrease the explanatory power of the model if the feature was informative in the first place. By comparing the difference in explanatory power between permuted and non-permuted models, we thus gain an understanding of the importance of this feature for explaining observed fees.[23]

[23]This approach to hypothesis testing is inspired by non-parametric permutation tests of experimental hypothe-

Furthermore, permutation tests allow us to assess the statistical significance of any incremental improvement in explanatory power, thanks to the 10,000 permutations drawn in each test. Under the Null hypothesis that the particular feature we are randomizing over is mere noise, we can interpret any significant improvement as a joint test of our model and data to explain observed fund fees.

[Insert Figure V around here]

Figure V displays the results of our tests, together with descriptions of the three permutation procedures employed. Under each procedure, we draw 10,000 permutations of the consideration set data, and for each permutation, we compare model-implied fees $\widehat{\boldsymbol{f}}$ to observed fees $\boldsymbol{f}^{obs}$ and collect the R^2 of each of these 10,000 regressions. We plot these in a histogram, and overlay (in red) the R^2 of our baseline calibration's explanatory power. Finally, we report p-values for a nonparametric one-sided statistical test comparing permuted R^2 values to our baseline model's R^2 value.

In all three permutation tests, we find that the R^2 of our baseline model-implied fees significantly exceeds the distribution of R^2 values from models calibrated upon the permuted datasets. Baseline R^2 values lie in the extreme right tail of the permutation distributions, and the associated p-values are effectively zero. The significant improvement in explanatory power in each test confirms that the entire observed structure of family co-consideration is essential to best explain the cross-section of fund fees.

By holding fixed fund-return moments and the overall level of EDGAR activity while permuting over co-consideration patterns, our permutation analysis demonstrates that the predictive power of our model cannot be attributed simply to non-consideration-related features, such as fund returns. Second, because the permutations preserve the marginal distribution of downloads across investors and families, they provide a non-parametric benchmark for how much explanatory power could arise purely from noisy or mismeasured consideration sets. The statistically significant drop in explanatory power when permuting various features of the data speaks directly to concerns that EDGAR-based consideration might be too noisy or too idiosyn-

ses, which are conducted by re-calculating a statistic of interest (e.g., the mean outcome) under permuted labelings of treated and control observations (Pitman, 1937). The approach is closely related to those for measuring the "restrictedness" of models (e.g., Selten, 1991; Fudenberg, Gao, and Liang, forthcoming), and in particular that of Fudenberg, Gao, and Liang (forthcoming), which involves simulating synthetic data sets and measuring how well the model fits those data sets relative to a baseline model. The difference is that where those approaches seek to measure how much a model's assumptions restrict its flexibility to fit any data, we seek to test whether a model's fit to our data is due to the restrictions imposed by its assumptions, and thus to the model capturing structure specific to our data. Also related are permutation-based approaches developed by machine learning researchers to measure the importance of different variables for predictive performance across whole classes of models (e.g., Fisher, Rudin, and Dominici, 2019).

cratic to matter for pricing: the economically meaningful structure in the observed EDGAR network is essential for explaining the cross-section of fees, rather than being an artefact of noise or of other ingredients of the model.

4.5 Additional Tests

Expense Ratios. We verify that model-implied fees accurately predict expense ratios, specifically, in Internet Appendix E.

Established Funds. We verify that our main results are not qualitatively altered by limiting our sample to well-established funds that have at least $10 Million in assets and an age of 12 months. See Internet Appendix E.

5 Distribution Channels

This paper explicitly models the direct consideration channel, in which investors directly select families (and their funds), and thus a family's position in the consideration-based competition network shapes funds' pricing decisions. The current section discusses our mechanism in the context of other distribution channels that are present in the asset management industry.

5.1 Captive Investors

We first address the possibility that some investors are held captive within a single fund family, and thus do not consider mutual funds offered by other fund families when selecting an investment. Such a situation may arise when families that are affiliated to retirement plans restrict investment options to within the family (Elton, Gruber, and Blake, 2006; Pool, Sialm, and Stefanescu, 2016). Clearly, if *all* investors are held captive by families, funds outside the family have no ability to lure them, whether by lowering fees or otherwise.[24] Figure I shows the median consideration set size consists of a single family, which suggests some investors may be captive (although we cannot distinguish captivity from limited attention using this data).

Would the presence of a large number of captive investors influence our results? We argue it should not materially change our results, for two reasons. Conceptually, even if many investors are captive to a family, as long as a fringe of investors are able and willing to consider funds from multiple families then families retain an incentive to compete for this segment of investors. And practically, our calibrated model already allows for singleton consideration sets that could

[24] Such a situation would correspond to a monopoly. We specialize our model to a monopolistic market structure in Section 6.1.

result from captivity, highlighting our framework's flexibility and ability to account for complex market structures.

To demonstrate this, we rerun our calibration step, this time excluding every consideration set that consists of a single element (singletons). Summary statistics for the sub-sample of consideration sets that excludes singletons are reported in Internet Appendix F. As before, our calibration procedure yields a vector of model-implied fees $\widehat{f}$, but these may differ from our baseline calibration because they are generated by explicitly excluding any role for investors with singleton consideration sets. We compare the new model-implied fees $\widehat{f}$ to observed fund fees f^{obs} (which are unchanged).

[Insert Table II around here]

Table II columns (1) and (2) repeat previous results (from Table I) comparing model-implied fees from the baseline calibration in which no consideration sets are excluded. Columns (3) and (4) now report new comparisons of model-implied fees yielded by an alternative calibration in which singleton consideration sets are excluded before the competition network is calculated; the rest of the calibration steps are identical to the baseline calibrations.

Unsurprisingly, the associations between model-implied and observed fees are similar across both calibrations. The invariance of the results to excluding singletons highlights that our model's results are driven by marginal investors, whose consideration sets reveal are selecting among the funds of at least two families. This is precisely what one should expect for a competition-based mechanism such as ours.

5.2 Intermediated Distribution

We next examine the role played by distribution channels. Our proposed mechanism is intended to capture the direct channel of sales to investors, as this is where family positioning is most likely to influence investor choices. However, some investors' portfolio decisions could be mediated by brokers or other intermediaries (Bergstresser, Chalmers, and Tufano, 2008), whose recommendations may be shaped by compensation contracts or other agreements. For instance, Del Guercio and Reuter (2014) find that the flow-performance relationship weakens for broker-sold funds. When intermediaries play a role, the direct mapping from each investor's own consideration set to competitive pressure is likely to be weaker than in the direct channel we model. It is nonetheless useful to quantify our mechanism's importance when families acquire investors through an indirect channel.

[Insert Table III around here]

We use observed 12b-1 distribution fees as a proxy for reliance on indirect distribution among the S&P 500 funds in our sample. Table III shows our baseline model-implied fees most accurately predict observed fees for directly-sold funds: for funds charging no 12b-1 fees, the relationship between model-implied and observed fees (excluding 12b-1 fees and front loads) is nearly one-to-one. However, the relationship remains statistically and economically significant even for funds that do charge 12b-1 fees. This suggests that family positioning influences fund sales even when brokers and other intermediaries are involved in distribution.

6 Counterfactual Fees

This section uses our model to theoretically and empirically compute counterfactual fees. Our organizing framework remains the Nash equilibrium fees that we derived in Proposition 1. To highlight that our results depend directly on market structure induced by family identities, we re-express the result in Eqn. (24) in terms of the adjacency graph $\boldsymbol{E}$ explicitly:

$$\boldsymbol{f}^*(\boldsymbol{E}) = \Big(\mathbf{I} + \boldsymbol{Q} \odot \operatorname{diag}(\boldsymbol{E})^{-1}\boldsymbol{E}\Big)^{-1}\Big(\boldsymbol{Q} \odot \operatorname{diag}(\boldsymbol{E})^{-1}\boldsymbol{E}\Big)\Big(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\Big), \tag{29}$$

In Eqn. (29) above, $\odot$ connotes the Hadamard product (elementwise multiplication), and we define $\boldsymbol{Q}$ as the $N \times N$ matrix consisting of elements $q_{ij} \triangleq \frac{\sigma_i}{\sigma_j}$.

We now proceed with deriving a variety of counterfactual market structures $\boldsymbol{E}$, and then compute the resultant counterfactual fees $\boldsymbol{f}^*(\boldsymbol{E})$.

6.1 Derivation of the Monopolistic Benchmark

In the absence of competition – i.e., in a monopoly market structure – each fund captures a group of investors who do not consider any alternatives. Under our model, every fund thus forms its own distinct component of the network. This is illustrated by Figure VI(a). In this case, the corresponding adjacency matrix is $\boldsymbol{E}^{\text{Monopoly}} = \boldsymbol{I}$, so Eqn. (29) evaluates to:

$$\boldsymbol{f}^{\text{Monopoly}} \;=\; \boldsymbol{f}^*(\boldsymbol{E} = \boldsymbol{I}) \;=\; \frac{1}{2}\Big(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\Big) \tag{30}$$

Equivalently, each fund i's equilibrium fee f_i takes the following scalar closed-form:

$$f_i^{\text{Monopoly}} \;=\; \frac{1}{2}\big(\eta_i - \sigma_i r_0\big) \tag{31}$$

Eqn. (31) highlights that fund i's fee now no longer depends on the performance of other funds $j \neq i$, but only upon its own values of η_i and σ_i. This is due to the fact that in a monopolistic

market structure competitors simply do not exist. Equilibrium fees are thus set irrespective of the performance of other funds, since no investor considers them to be an alternative, removing any risk of an investor switching to a competitor.

[Insert Figure VI around here]

This theoretical result adds an important nuance to the argument made in prior research that *competition* between funds will lead to an equilibrium in which funds' fees equal their expected excess returns (e.g., Berk and Green, 2004; Gil-Bazo and Ruiz-Verdú, 2009). Berk and Green (2004) argue this equality will arise because investors will rationally reallocate their investment towards funds offering positive net-of-fee returns, which will cause excess returns to drop back into equality with fees due to the decreasing economies of scale in fund management.[25] Gil-Bazo and Ruiz-Verdú (2009) argue the equality could also arise from competitive fee-adjustment by funds. However, in classic terms, Eqn. (31) describes a situation in which a firm successfully prices to extract the entire surplus it creates by trading with its customers. Such a situation is not competitive, but rather entails the firm having substantial market power.

6.2 Specifying Competitive Benchmarks

To supplement the monopoly fee counterfactual (above), we construct additional counterfactuals for the opposite case in which the market structure is highly competitive. Our approach remains the same: we define variants of $\boldsymbol{E}$ that correspond to different patterns of consideration, in order to then compute the ensuing equilibrium fee vector $\boldsymbol{f}^*(\boldsymbol{E})$. We evaluate two highly-competitive market structures.

The Uniform Competition Benchmark. In our first competitive counterfactual, which we term the *Uniform* benchmark, each investor considers a subset of available alternatives, but every possible pairwise consideration set exists and has an equal likelihood of occurring. We implement this market structure by assuming each investor considers a pair of funds: $N_m = 2$ in Eqns. (13a) & (13b); this corresponds to the empirical mean consideration set size (see Figure I). Out of all possible N funds, a total of $\frac{N(N-1)}{2}$ pair-sized consideration sets exist, and so the (uniform) mass of investors matched to each is $\theta_m = \frac{2}{N(N-1)}$ in Eqn. (15). This is illustrated by Figure VI(b). One can think of this scenario as capturing investors with uniform, homogeneous

[25] Berk and Green (2004) study a dynamic model in which the equality of fees and expected excess returns defines an equilibrium, and competition between funds is absent. In contrast, we study a static model in which competition is present, but the mechanism proposed by Berk and Green to equate fees and returns – the dependence of fund performance on investor allocations – is absent. Thus the fee-return relationship between the two models is not directly comparable.

search costs. As a result of these parameters, the Uniform market structure takes this form:

$$\boldsymbol{E}^{\text{Uniform}} = \frac{2}{N(\rho^2-1)} \begin{bmatrix} \ddots & \cdots & \frac{\rho}{N-1} \\ \vdots & -1 & \vdots \\ \frac{\rho}{N-1} & \cdots & \ddots \end{bmatrix} \tag{32}$$

The Universal Competition Benchmark. In our second competitive counterfactual, which we term the *Universal* benchmark, each investor is aware of and considers all available funds. This is illustrated by Figure VI(c). There is thus a universal, common consideration set containing all N available funds, so $N_m = N$ in Eqns. (13a) & (13b). Because there is only a single consideration set, $\theta_m = 1$ in Eqn. (15). Thus the Universal market structure is as follows:

$$\boldsymbol{E}^{\text{Universal}} = \frac{1}{(1-\rho)[(N-1)\rho+1]} \begin{bmatrix} \ddots & \cdots & -\rho \\ \vdots & (N-2)\rho+1 & \vdots \\ -\rho & \cdots & \ddots \end{bmatrix} \tag{33}$$

6.3 Calibration and Comparison

We have previously calibrated our model to observed consideration sets, which effectively specified an observed competition network $\boldsymbol{E}^{\text{Observed}}$, and then used this to produce model-predicted fees $\widehat{\boldsymbol{f}} = \boldsymbol{f}^*(\boldsymbol{E}^{\text{Observed}})$ in Section 4.3. Our current goal is to compare these factual fees to various counterfactuals to assess the degree of competition among mutual fund families.

The counterfactuals we will examine are the Monopoly benchmark $\boldsymbol{f}^{\text{Monopoly}} \triangleq \boldsymbol{f}^*(\boldsymbol{E}^{\text{Monopoly}})$, the Universal competition benchmark $\boldsymbol{f}^{\text{Universal}} \triangleq \boldsymbol{f}^*(\boldsymbol{E}^{\text{Universal}})$, and the Uniform competition benchmark $\boldsymbol{f}^{\text{Uniform}} \triangleq \boldsymbol{f}^*(\boldsymbol{E}^{\text{Uniform}})$. The competition network parameters $\boldsymbol{E}^{\text{Monopoly}}$, $\boldsymbol{E}^{\text{Universal}}$ and $\boldsymbol{E}^{\text{Uniform}}$ are as defined above. Other calibrated parameters (ρ, N, and expected returns) are unchanged from our baseline calibration (Section 4).

We proceed with the three counterfactual calibrations to obtain counterfactual fees $\boldsymbol{f}^{\text{Monopoly}}$, $\boldsymbol{f}^{\text{Uniform}}$ and $\boldsymbol{f}^{\text{Universal}}$. Figure VII displays a comparison of the equilibrium fee vectors to a common yardstick of the monopoly counterfactul fees $\boldsymbol{f}^{\text{Monopoly}}$: to be precise, the chart displays regression coefficients of each fee vector upon the corresponding vector of monopoly fees $\boldsymbol{f}^{\text{Monopoly}}$ (as the common independent variable in the OLS regressions).

[Insert Figure VII around here]

The comparison presented in Figure VII reveals that baseline equilibrium fees are approx-

imately one-third the level of monopoly fees. By contrast, the counterfactual fees that would have obtained in a highly-competitive market are closer to a tenth of monopoly fees, and a third of the baseline equilibrium fees. This counterfactual analysis shows that the baseline degree of competition, while far from monopolistic, remains imperfect: the family component of fund fees is approximately three times as high as the highly-competitive benchmarks.

7 Analyzing Investor Consideration Sets

Having determined that family-level consideration patterns influence fund-level fee-setting, we now investigate the drivers of consideration set formation by investors.

7.1 Co-consideration Likelihood Increases in Fund Family Similarity

We first examine the relationship between fund family similarity and co-consideration by investors. We do so by comparing the aggregate frequency of co-consideration for each pair of fund families to the similarity of their attributes.

Frequency of co-consideration is measured by the number of EDGAR downloaders considering both families within a given calendar year. Similarity of attributes is measured by the absolute distance between the percentiles of two families in the cross-sectional distribution for each attribute. Thus each variable takes values in the range $[0, 1]$, where 0 corresponds to two families having the same attribute value, and 1 corresponds to the families having the minimum and maximum values for the attribute. Finally, we compare co-consideration frequency to attribute similarity by estimating Poisson regressions with fixed effects (Cohn, Liu, and Wardlaw, 2022).

[Insert Table IV around here]

We find that fund families which are co-considered more frequently are indeed more similar in their observable attributes. Table IV presents the results of the Poisson regressions. Specifications (1)-(2) include additive fixed effects for each pair of companies i, j to control for the mechanical effect of an often-considered family i or j increasing the frequency of the pair being co-considered. Specification (2) adds year fixed effects, while specification (3) includes interacted family-year fixed effects to control for family attributes that vary over time, including the time-varying frequency with which each family is considered. Including these fixed effects isolates the relationship between pairwise attribute similarity and co-consideration. Triple-clustered standard errors are used to assess coefficient significance. All significant coefficients are negative in magnitude: similarity in fund family attributes is thus associated with being

co-considered by more investors.

The two family attributes in which similarity is most strongly associated with co-consideration are the share of Total Net Assets (TNA) following a passive indexing strategy and the number of offered funds. The estimated coefficient on Passive Share Distance of -0.6691 implies two fund families with 0% and 100% of their TNA following a passive indexing strategy can expect with be co-considered by $1-e^{-0.6691} = 49\%$ fewer investors than if they had equal proportions of their TNA following passive indexing strategies. Similarly, two families offering a number of funds placing them in the 1st and 100th percentiles among all families are co-considered $1-e^{-0.3953} = 33\%$ less often than equivalent families offering the same number of funds. The importance of similar fund variety for co-consideration implied by this second result is consistent with the importance of brand assortment for overall sales and perceived brand quality established by the marketing literature (Borle et al., 2005; Berger, Draganska, and Simonson, 2007).

These findings on the similarity in attribute space of co-considered families could be consistent with investors selecting mutual funds by first directing their attention toward a few fund families offering similar assortments of funds, and thus serving as imperfect substitutes for the fulfillment of their investment goals. The implication for mutual fund competition is that fund families, not just individual funds, compete for investor attention – and thus investor dollars – with competitors located near to them in the space of observable fund family attributes.

7.2 Consideration Set Size Increases With Investor Sophistication

We next investigate the determinants of consideration set size. Since a larger consideration set reveals investors are potentially selecting among a wider range of alternatives, one should expect this to correlate with investor sophistication, for two reasons. First, more sophisticated investors are likely to incur a lower cost of search and evaluation. And second, more sophisticated investors are less likely to fall prey to behavioral biases and heuristics.

We measure the size of an investor's consideration set as the number of unique CIKs that identify all the prospectus documents that have been acquired by that investor. We also collect a number of sociodemographic characteristics that proxy for investors' level of financial sophistication; the characteristics and the process by which they are collected are described in Internet Appendix C.

We find significant covariation between sophistication-proxying demographic variables and consideration set size aggregated to the county level, which can be interpreted as describing

relationships between the attributes of a county-level representative investor.[26] Table V shows the results of OLS and panel regressions of the county-level average consideration set size for 2006 through 2016. The covariates are county-level demographic variables standardized as z-scores. Thus each estimated coefficient measures the difference in average consideration set size associated with a value of the covariate one standard deviation above the mean. Therefore, the estimated coefficients measure the extent to which the same degree of variation in each covariate explains variation in the average size of consideration sets. Recall that the means and standard deviations of the un-standardized covariates are reported in Table IA.2.

[Insert Table V around here]

The estimated relationship between sophistication-proxying demographics and consideration set size is in the direction we would expect from the literature on investor sophistication. Larger average consideration sets are significantly associated with younger, wealthier, more educated and more urban investors. The significance and magnitudes of these associations are little-changed when time effects are included. The estimated magnitudes of the coefficients measure the difference in average consideration set size from the mean of 1.56 associated with a one standard deviation difference in the covariate. Thus a county whose residents' median household income that is one standard deviation above the county-mean can be expected to have investors with consideration sets that include 0.08 more funds.

We include political affiliation as a control in these regressions because it is an observable proxy for the models of the world that investors use when forming beliefs (Meeuwis, Parker, Schoar, and Simester, 2022), investors' non-pecuniary preferences (Hong and Kostovetsky, 2012), and a variety of other sociodemographic attributes (Griffin, Frey, and Teixeira, 2019)

The relationships we find between sophistication-proxying demographics and consideration set size are consistent with more sophisticated investors considering a greater number of mutual funds. Agarwal, Driscoll, Gabaix, and Laibson (2009) show older investors make more financial mistakes, which they argue is due to changes in cognitive ability; we find older counties have investors who consider fewer funds on average. Von Gaudecker (2015) finds more financially literate households hold more diversified portfolios. If financial literacy is associated with college education, then our finding of counties with a greater share of college graduates having larger average consideration sets is consistent with more financially literate investors selecting a wider set of funds to consider. Differing levels of financial sophistication most likely explain

[26] Our results – which are based on county-level means and medians – could also be interpreted as approximations for individual investor behaviour, though we do not emphasise or rely on such an interpretation in our discussion. County-level characteristics have previously been used as proxies for individual characteristics in the literature (e.g., Meeuwis, Parker, Schoar, and Simester, 2022).

our findings on the relationship between household income and average consideration set size (Calvet, Campbell, and Sodini, 2009), and the urban-rural gap in consideration set size (Klapper and Panos, 2011).

8 Conclusion and Discussion

This paper shows that the relative positioning between fund *families* is an important force shaping *fund* fees in a seemingly homogeneous category of quasi-identical index funds.

Our contributions include a network-based theory of fund fee competition that produces closed-form equilibrium fee predictions, and is sufficiently general that it can be applied to study price dispersion in other financial contexts.

Our conclusion that fund families' identities influences the fees that their offered funds can charge poses two important questions for further research. First, what actions do fund families take in order to improve their relative positioning in the competitive landscape? Second, to what extent are fee decisions centralized at the family level?[27] The evidence on the latter question is so far suggestive: Christoffersen (2001, pp. 1137-1138) argues the fees of some funds might be explained by "loss-leader" pricing strategies, by which families set low fees on some funds to attract investors into the overall complex. Iannotta and Navone (2012) document a substantial family-level component in fees, which could be due to centralized fee-setting. Khorana and Servaes (2012) argue that families pass along savings from economies of scale as lower fees. The mechanism that we document in the present study does not rely upon explicitly centralized decision-making at the family-level, but would also reinforce it if it did exist. Further research on within-family decision-making would be valuable in this regard.

Our paper also introduces a rich and unusually comprehensive source of data on investors' consideration patterns. We encourage future work to likewise use investors' prospectus acquisition decisions on EDGAR to study other aspects of investor behaviour and the equilibrium responses of other investment vehicles.

Relatedly, our finding that investors' consideration patterns vary with financial sophistication proxies suggests a further examination of the supply side of the mutual fund industry (similar to that of Ru and Schoar (2016)) could shed light on potential welfare implications for investors.

Finally, the prevalence of sparse and heterogeneous investor consideration sets that we document shows this is an important feature of investor decision making. Future work could

[27]Because fund families are known to centralize various decisions, Ben-David, Franzoni, Moussawi, and Sedunov (2021) argue they are "granular" in their behavior.

explore the importance of limited and heterogeneous consideration for households' decisions about other financial products, as well as incorporate limited consideration into an asset pricing demand system (Koijen and Yogo, 2019).

References

Abis, S., and Lines, A. (2024). Broken promises, competition, and capital allocation in the mutual fund industry. *Journal of Financial Economics*, *162*, 103948.

Agarwal, S., Driscoll, J. C., Gabaix, X., and Laibson, D. (2009). The age of reason: Financial decisions over the life cycle and implications for regulation. *Brookings Papers on Economic Activity*, *2009*(2), 51–117.

Ballester, C., Calvó-Armengol, A., and Zenou, Y. (2006). Who's who in networks. Wanted: The key player. *Econometrica*, *74*(5), 1403–1417.

Barber, B. M., and Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, *21*(2), 785–818.

Belo, F., Lin, X., and Vitorino, M. A. (2014). Brand capital and firm value. *Review of Economic Dynamics*, *17*(1), 150–169.

Ben-David, I., Franzoni, F., Moussawi, R., and Sedunov, J. (2021). The granular nature of large institutional investors. *Management Science*, *67*(11), 6629–6659.

Berger, J., Draganska, M., and Simonson, I. (2007). The influence of product variety on brand perception and choice. *Marketing Science*, *26*(4), 460–472.

Bergstresser, D., Chalmers, J. M., and Tufano, P. (2008). Assessing the costs and benefits of brokers in the mutual fund industry. *The Review of Financial Studies*, *22*(10), 4129–4156.

Berk, J. B., and Green, R. C. (2004). Mutual fund flows and performance in rational markets. *Journal of Political Economy*, *112*(6), 1269–1295.

Bonacich, P. (1972). Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, *2*(1), 113–120.

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology*, *92*(5), 1170–1182.

Bonelli, M., Buyalskaya, A., and Yao, T. (2023). Financial product incentives to differentiate: Evidence from mutual funds. *Available at SSRN*, *3939239*.

Borle, S., Boatwright, P., Kadane, J. B., Nunes, J. C., and Galit, S. (2005). The effect of product assortment changes on customer retention. *Marketing Science*, *24*(4), 616–622.

Brown, Z. Y., Egan, M. L., Jeon, J., Jin, C., and Wu, A. A. (2023). Why do index funds have market power? Quantifying frictions in the index fund market. *NBER Working Paper No. 31778*.

Calvet, L. E., Campbell, J. Y., and Sodini, P. (2009). Measuring the financial sophistication of households. *American Economic Review: Papers & Proceedings*, *99*(2), 393–98.

Christoffersen, S. E. (2001). Why do money fund managers voluntarily waive their fees? *Journal of Finance*, *56*(3), 1117–1140.

Cohn, J. B., Liu, Z., and Wardlaw, M. I. (2022). Count (and count-like) data in finance. *Journal of Financial Economics*, *146*(2), 529–551.

DeHaan, E., Song, Y., Xie, C., and Zhu, C. (2021). Obfuscation in mutual funds. *Journal of Accounting and Economics*, *72*(2-3), 101429.

Del Guercio, D., and Reuter, J. (2014). Mutual fund performance and the incentive to generate alpha. *Journal of Finance*, *69*(4), 1673–1704.

Dou, W. W., and Wu, W. (2023). Competition network and predictable industry returns. *Available at SSRN*, *4535226*.

Elton, E. J., and Gruber, M. J. (1973). Estimating the dependence structure of share prices – Implications for portfolio selection. *Journal of Finance*, *28*(5), 1203–1232.

Elton, E. J., Gruber, M. J., and Blake, C. R. (2006). The adequacy of investment choices offered by 401(k) plans. *Journal of Public Economics*, *90*(6-7), 1299–1314.

Elton, E. J., Gruber, M. J., and Busse, J. A. (2004). Are investors rational? Choices among index funds. *Journal of Finance*, *59*(1), 261–288.

Elton, E. J., Gruber, M. J., and Urich, T. J. (1978). Are betas best? *Journal of Finance*, *33*(5), 1375–1384.

Fisher, A., Rudin, C., and Dominici, F. (2019). All models are wrong, but many are useful: Learning a variable's importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, *20*(177), 1–81.

Fudenberg, D., Gao, W., and Liang, A. (forthcoming). How flexible is that functional form? Quantifying the restrictiveness of theories. *Review of Economics and Statistics*.

Gallaher, S., Kaniel, R., and Starks, L. T. (2015). Advertising and mutual funds: From families to individual funds. *CEPR Discussion Paper*, (10329).

Gil-Bazo, J., and Ruiz-Verdú, P. (2009). The relation between price and performance in the mutual fund industry. *Journal of Finance*, *64*(5), 2153–2183.

Grice, R. (2023). Chain-linked markets.

Griffin, R., Frey, W., and Teixeira, R. (2019). States of change: How demographic change is transforming the Republican and Democratic parties, Brookings Institution.

Gueciloueur, A. (2023). Soothing investors: The impact of manager communication on mutual fund flows. *Available at SSRN, 5465254*.

Hoberg, G., Kumar, N., and Prabhala, N. (2018). Mutual fund competition, managerial skill, and alpha persistence. *Review of Financial Studies*, *31*(5), 1896–1929.

Hoberg, G., and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, *124*(5), 1423–1465.

Hong, H., and Kostovetsky, L. (2012). Red and blue investing: Values and finance. *Journal of Financial Economics*, *103*(1), 1–19.

Hortaçsu, A., and Syverson, C. (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics*, *119*(2), 403–456.

Iannotta, G., and Navone, M. (2012). The cross-section of mutual fund fee dispersion. *Journal of Banking & Finance*, *36*(3), 846–856.

Investment Company Institute. (2025). Investment company factbook.

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, *18*, 39–43.

Khorana, A., and Servaes, H. (2012). What drives market share in the mutual fund industry? *Review of Finance*, *16*(1), 81–113.

Klapper, L., and Panos, G. A. (2011). Financial literacy and retirement planning: The Russian case. *Journal of Pension Economics & Finance*, *10*(4), 599–618.

Koijen, R. S., and Yogo, M. (2019). A demand system approach to asset pricing. *Journal of Political Economy*, *127*(4), 1475–1515.

Kostovetsky, L., and Warner, J. B. (2020). Measuring innovation and product differentiation: Evidence from mutual funds. *Journal of Finance*, *75*(2), 779–823.

Lee, C. M., Ma, P., and Wang, C. C. (2015). Search-based peer firms: Aggregating investor perceptions through internet co-searches. *Journal of Financial Economics*, *116*(2), 410–431.

Loseto, M. (2023). The structure of investment advisors' markups: Evidence from retirement allocations.

Loseto, M., and Mainardi, F. (2023). Market power, fund proliferation, and asset prices. *Available at SSRN, 4364189*.

Massa, M. (2003). How do family strategies affect fund performance? When performance-maximization is not the only game in town. *Journal of Financial Economics*, *67*(2), 249–304.

Meeuwis, M., Parker, J. A., Schoar, A., and Simester, D. (2022). Belief disagreement and portfolio choice. *Journal of Finance*, *77*(6), 3191–3247.

Nanda, V., Wang, Z. J., and Zheng, L. (2004). Family values and the star phenomenon: Strategies of mutual fund families. *Review of Financial Studies*, *17*(3), 667–698.

Pellegrino, B. (2025). Product differentiation and oligopoly: A network approach. *American Economic Review*, *115*(4), 1170–1225.

Pitman, E. J. G. (1937). Significance tests which may be applied to samples from any populations. *Journal of the Royal Statistical Society Supplement*, *4*, 119–130.

Pool, V. K., Sialm, C., and Stefanescu, I. (2016). It pays to set the menu: Mutual fund investment options in 401(k) plans. *Journal of Finance*, *71*(4), 1779–1812.

Roussanov, N., Ruan, H., and Wei, Y. (2021). Marketing mutual funds. *Review of Financial Studies*, *34*(6), 3045–3094.

Ru, H., and Schoar, A. (2016). Do credit card companies screen for behavioral biases? *NBER Working Paper No. 22360*.

Selten, R. (1991). Properties for a measure of predictive success. *Mathematical Social Sciences*, *21*, 153–167.

Sialm, C., and Tham, T. M. (2016). Spillover effects in mutual fund companies. *Management Science*, *62*(5), 1472–1486.

Sirri, E. R., and Tufano, P. (1998). Costly search and mutual fund flows. *Journal of Finance*, *53*(5), 1589–1622.

Vives, X. (2000). *Oligopoly pricing: Old ideas and new tools*. MIT Press.

Von Gaudecker, H.-M. (2015). How does household portfolio diversification vary with financial literacy and financial advice? *Journal of Finance*, *70*(2), 489–507.

Wahal, S., and Wang, A. Y. (2011). Competition among mutual funds. *Journal of Financial Economics*, *99*(1), 40–59.

Wilks, S. S. (1932). Certain generalizations in the analysis of variance. *Biometrika*, *24*(3/4), 471–494.

Appendix

Proof of Proposition 1

Proposition 1 (Nash equilibrium). *There exists a unique Nash equilibrium to the fee-setting game between funds characterized by non-zero fees:*

$$\boldsymbol{f}^* = \left(\mathbf{I} + \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\right)^{-1} \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1}\widetilde{\boldsymbol{E}}\left(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\right), \tag{24}$$

where $\boldsymbol{\eta}$ is a vector with elements $\eta_i \triangleq \mathbb{E}[r_i]$, $\operatorname{diag}(\boldsymbol{S})$ is the $N \times N$ diagonal matrix with the elements σ_i on the leading-diagonal and 0 on the off-diagonals, and $\widetilde{\boldsymbol{E}} \triangleq \boldsymbol{E}\operatorname{diag}(\boldsymbol{S})^{-1}$ is the $N \times N$ matrix with the elements $\widetilde{e}_{ij} \triangleq \frac{e_{ij}}{\sigma_j}$, and thus $\operatorname{diag}(\widetilde{\boldsymbol{E}}) = \operatorname{diag}(\boldsymbol{E})\operatorname{diag}(\boldsymbol{S})^{-1}$.

Proof. The proof proceeds in two steps. The first step is to show the linear mapping described by the system of best reply functions in Eqn. (21) is a contraction mapping in $\mathbb{R}^N$. The second step is to show that there exists a convex, compact set $\Delta \subset \mathbb{R}^N_{++}$ over which the contraction mapping operates, and to then apply Banach's Fixed Point Theorem to establish existence of a unique fixed point within Δ. This fixed point is the Nash equilibrium. The proof concludes with a characterization of the fixed point.

Let us re-express the best reply functions in Eqn. (21) as

$$\widetilde{f_i^*}(\widetilde{\boldsymbol{f}_{-i}}) = \frac{1}{2}\left(1 + \sum_{j \neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f}_j\right)\right)$$

where

$$\widetilde{f}_i \equiv \frac{f_i}{\eta_i - \sigma_i r_0}$$

for all $i \in \mathcal{N}$. Denote by $T : \mathbb{R}^N \to \mathbb{R}^N$ the linear mapping described by these best reply functions. T is a contraction mapping over $\mathbb{R}^N$ equipped with the L_1-norm if:

$$|T(\widetilde{\boldsymbol{f}_1}) - T(\widetilde{\boldsymbol{f}_0})| < |\widetilde{\boldsymbol{f}_1} - \widetilde{\boldsymbol{f}_0}|$$

for $\widetilde{\boldsymbol{f}_0}, \widetilde{\boldsymbol{f}_1} \in \mathbb{R}^N$, which would be implied by:

$$|T_i(\widetilde{\boldsymbol{f}_1}) - T_i(\widetilde{\boldsymbol{f}_0})| < |\widetilde{\boldsymbol{f}_{1i}} - \widetilde{\boldsymbol{f}_{0i}}| \quad \forall i \in \mathcal{N}$$

where $T_i(\cdot)$, $\widetilde{\boldsymbol{f}_{1i}}$ and $\widetilde{\boldsymbol{f}_{0i}}$ are the ith rows of the respective vectors.

We can establish the following inequality:

$$
\begin{aligned}
|T_i(\widetilde{\boldsymbol{f}}_1) - T_i(\widetilde{\boldsymbol{f}}_0)| &= \left| \frac{1}{2}\left(1 + \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f_{1j}}\right)\right) - \frac{1}{2}\left(1 + \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f_{0j}}\right)\right)\right| \\
&< \left| \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(\widetilde{f_{1j}} - \widetilde{f_{0j}}\right)\right| \\
&\leq \sum_{j\neq i} \left| \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\right| \left|\widetilde{f_{1j}} - \widetilde{f_{0j}}\right|
\end{aligned}
$$

where the last inequality follows from absolute homogeneity and the triangle inequality for the L_1-norm. From the strict diagonal dominance of $\mathbf{E}$ with weights $\frac{\sigma_i}{\eta_i - \sigma_i r_0}$ given by Assumption 1 we have:

$$
\sum_{j\neq i} \left| \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\right| < 1 \quad \forall i \in \mathcal{N},
$$

and thus T is a contraction mapping.

We next want to show that there exists a convex, compact set $\Delta \subset \mathbb{R}^N_{++}$ for which $T(\Delta) \subseteq \Delta$. Let $\Delta \equiv [\epsilon, 1-\epsilon]^N$ for some ϵ satisfying

$$
0 < \epsilon < \frac{1}{2}\left(1 - \max_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}, j\neq i} \frac{|e_{ij}|}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\right).
$$

From $\mathbf{E}$ with weights $\frac{\sigma_i}{\eta_i - \sigma_i r_0}$ strictly diagonally dominant we have $\epsilon \in (0, \frac{1}{2})$, which implies Δ is nonempty, convex, closed and bounded (and thus compact) in $\mathbb{R}^N$, and bounded away from the origin. It remains to establish:

$$
\begin{aligned}
& T_i(\widetilde{\boldsymbol{f}}) \in [\epsilon, 1-\epsilon] \\
\Leftrightarrow\ & \frac{1}{2}\left(1 + \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f}_j\right)\right) \in [\epsilon, 1-\epsilon] \\
\Leftrightarrow\ & \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f}_j\right) \in [-(1-2\epsilon), 1-2\epsilon]
\end{aligned}
$$

for all $i \in \mathcal{N}$ and all $\widetilde{f}_j \in [\epsilon, 1-\epsilon]$. Thus it is necessary and sufficient to show:

$$
\left| \sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1 - \widetilde{f}_j\right)\right| \leq 1 - 2\epsilon
$$

for all $i \in \mathcal{N}$ and all $\widetilde{f}_j \in [\epsilon,\, 1-\epsilon]$. This can be established as follows:

$$\begin{aligned}
\left|\sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\left(1-\widetilde{f}_j\right)\right| &< \left|\sum_{j\neq i} \frac{e_{ij}}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\right| \\
&= \sum_{j\in\mathcal{N}, j\neq i} \frac{|e_{ij}|}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0} \\
&\leq \max_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}, j\neq i} \frac{|e_{ij}|}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0} \\
&= 1 - 2\left(\frac{1}{2}\left(1 - \max_{i\in\mathcal{N}} \sum_{j\in\mathcal{N}, j\neq i} \frac{|e_{ij}|}{e_{ii}} \frac{\sigma_i}{\sigma_j} \frac{\eta_j - \sigma_j r_0}{\eta_i - \sigma_i r_0}\right)\right) \\
&< 1 - 2\epsilon.
\end{aligned}$$

Thus we have shown that T is a contraction mapping on Δ, a non-empty complete metric space. By Banach's Fixed Point Theorem there exists a unique fixed point $\mathbf{x}^* \in \Delta$ such that $T(\widetilde{\boldsymbol{f}^*}) = \widetilde{\boldsymbol{f}^*}$. Therefore, there exists a unique Nash equilibrium of the pricing game between firms $i \in \mathcal{N}$ with equilibrium prices $f_i^* \in (0,\, \eta_i - \sigma_i r_0)$ for every $i \in \mathcal{N}$.

The equilibrium outcomes can be characterized from the first order conditions to firms' profit maximization problem in Eqn. (19) as follows:

$$\begin{aligned}
\frac{\partial \pi_i}{\partial f_i} &= x_i + \frac{\partial x_i}{\partial f_i} f_i \\
&= x_i - \frac{e_{ii}}{\sigma_i} f_i \\
\Rightarrow \boldsymbol{x}(\boldsymbol{f}^*) &= \operatorname{diag}(\boldsymbol{E}) \operatorname{diag}(\boldsymbol{S})^{-1} \boldsymbol{f}^* \\
\Leftrightarrow \boldsymbol{E}\left(\operatorname{diag}(\boldsymbol{S})^{-1}(\boldsymbol{\eta} - \boldsymbol{f}^*) - r_0 \mathbf{1}\right) &= \operatorname{diag}(\boldsymbol{E}) \operatorname{diag}(\boldsymbol{S})^{-1} \boldsymbol{f}^* \\
\Leftrightarrow \widetilde{\boldsymbol{E}}\left(\boldsymbol{\eta} - \boldsymbol{f}^* - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\right) &= \operatorname{diag}(\widetilde{\boldsymbol{E}}) \boldsymbol{f}^* \\
\Leftrightarrow \boldsymbol{f}^* &= \left(\operatorname{diag}(\widetilde{\boldsymbol{E}}) + \widetilde{\boldsymbol{E}}\right)^{-1} \widetilde{\boldsymbol{E}}\left(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\right) \\
\Leftrightarrow \boldsymbol{f}^* &= \left(\mathbf{I} + \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1} \widetilde{\boldsymbol{E}}\right)^{-1} \operatorname{diag}(\widetilde{\boldsymbol{E}})^{-1} \widetilde{\boldsymbol{E}}\left(\boldsymbol{\eta} - r_0 \operatorname{diag}(\boldsymbol{S})\mathbf{1}\right)
\end{aligned}$$

where the fourth equality uses the direct demand in Eqn. (17). This characterization completes the proof. □

Figures

Figure I: Investor Consideration Set Size

This chart summarizes the distribution of the number of fund families included in each potential investor's consideration set. A potential investor's consideration set is measured using her prospectus downloads from the SEC EDGAR website each year. The y-axis represents the number of fund families, in log-scale.

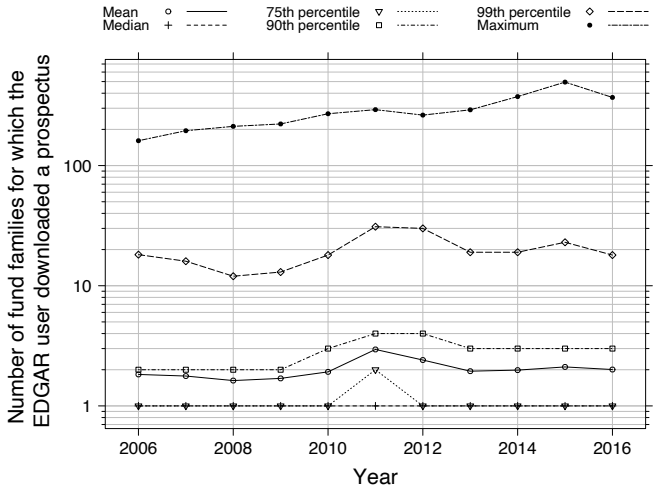


Figure II: Probability of a Fund Family to be Included in a Consideration Set

This chart summarizes the distribution of the frequency of inclusion of each fund family in any potential investor's consideration set. A potential investor's consideration set is measured using her prospectus downloads from the SEC EDGAR website each year. The y-axis represents the probability $\in [0, 1]$ of consideration set membership, in log-scale.

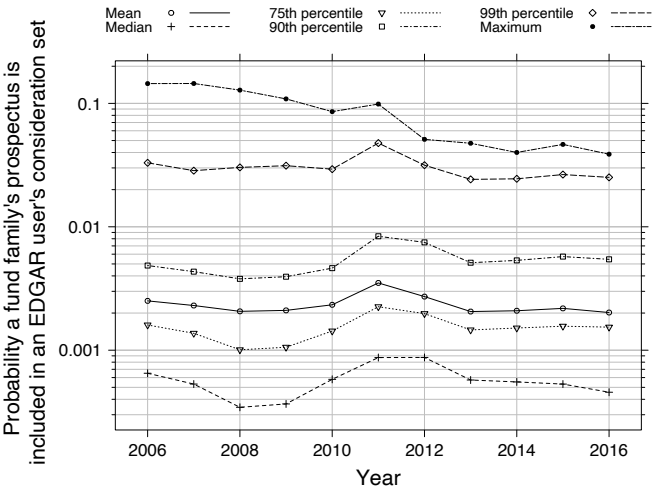
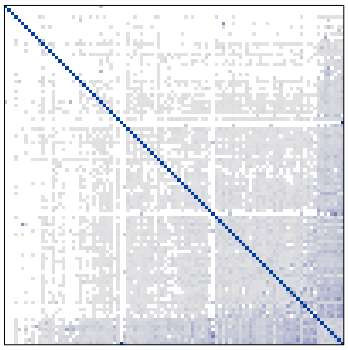


Figure III: Example Co-Consideration Patterns

Aggregate co-consideration patterns for two samples of 100 fund families in 2016. Each row represents a fund family, and each column also represents a fund family. Cell colors denote the (normalized) frequency of co-consideration: white cells have value 0 and indicate the row and column fund families were never co-considered during that period. Rows are normalized by the diagonal cells, which therefore take the maximum value 1. Sample (a) is a uniform sample. Sample (b) is conditioned on fund families being co-considered with the BlackRock family and then uniformly sampled, as an illustrative example.

(a) Uniform sample



(b) Co-considered with BlackRock

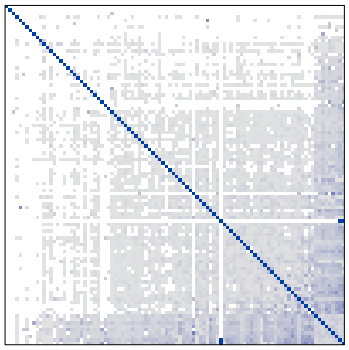


Figure IV: Consideration Sets Inducing a Competition Network

This figure illustrates an example in which the consideration sets of 3 investor segments (left) induce a network of weighted competitive ties between 4 funds (right). Self ties are not depicted.

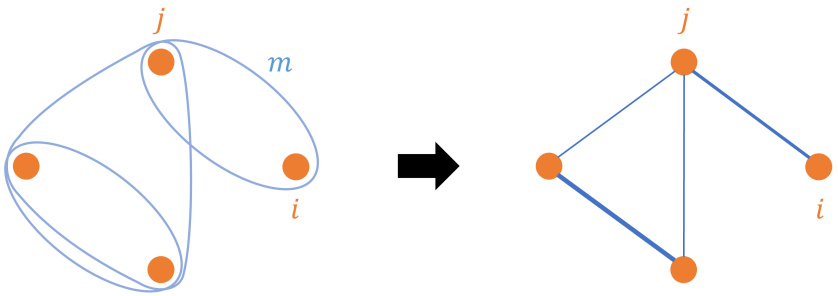
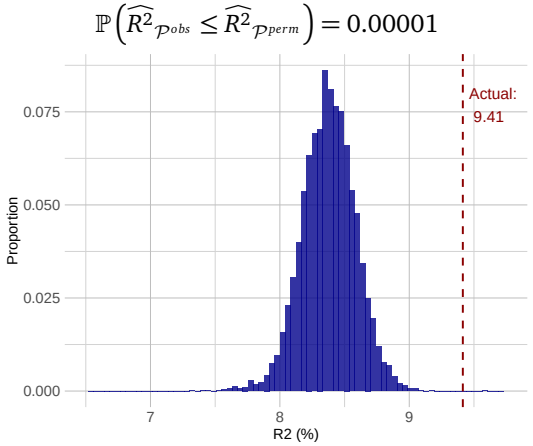


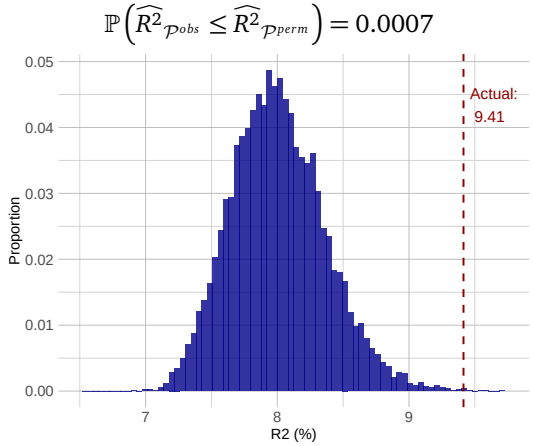
Figure V: Distribution of R^2 Values Across Permuted Patterns of Consideration

These sub-figures show the distribution of values of the R^2 statistic from regressions of f_{it} on $\widehat{f}_{it}$ across various permutations of $\mathcal{P}$. For each permutation $\mathcal{P}^{perm}$, the predicted fees $\widehat{f}_{it}$ are calculated as described in Section 4.2, and observed fees f_{it} are regressed on these predictions using the specification in Column (2) of Table I. A different pattern of permutation is used in each test. In (a), the pattern of consideration is permuted by re-drawing the identities of the families in each investor's consideration set each year from the families offering S&P 500 index funds that year, according to the empirical CDF. In (b), the pattern is permuted by swapping family identities. In (c), the pattern is permuted by re-drawing the identities of the investors considering each family each year from the investors considering families offering S&P 500 index funds that year, according to the empirical CDF. In each case, we draw 10,000 permutations. The vertical line marks the value of the baseline R^2 statistic for the observed pattern of consideration $\mathcal{P}^{obs}$.

(a) Re-drawing families per investor



(b) Shuffling family identities



(c) Re-drawing investors per family

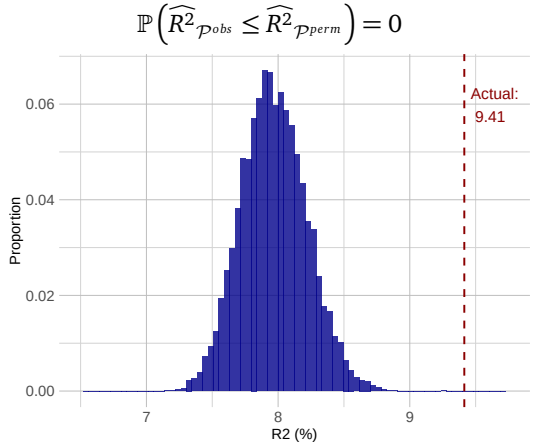


Figure VI: Illustrations of Counterfactual Market Structures

For an example scenario of $N = 3$ funds in total, the three benchmark market structures are illustrated in the following sub-figures. Orange points represent funds, and blue shapes represent unique investor consideration sets (potentially shared by multiple investors).

(a) Monopoly **(b)** Uniform **(c)** Universal

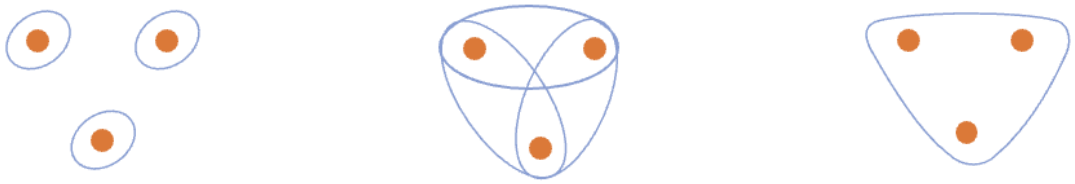
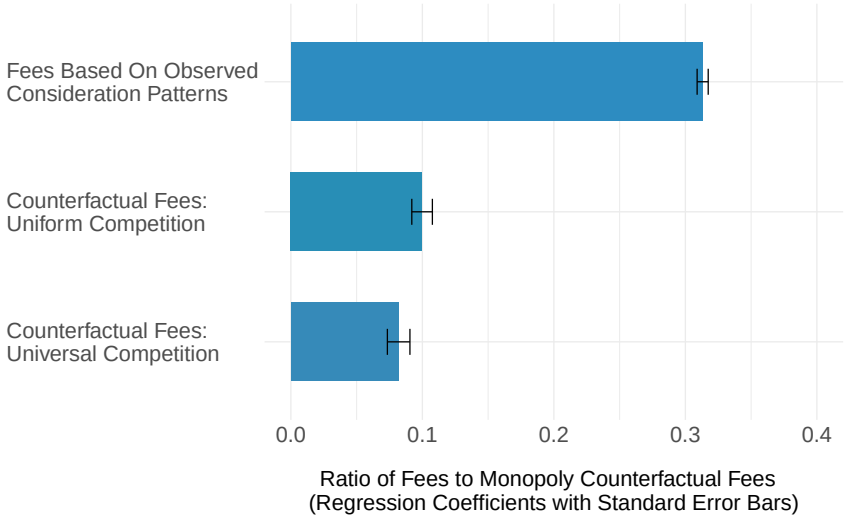


Figure VII: Comparison of Calibrated Fees

This chart shows OLS regression coefficients (and standard error bars) for separate regressions of counterfactual monopoly fees $\boldsymbol{f}^{\text{Monopoly}}$ (the independent variable) against other equilibrium fees predicted by different calibrations of our model: $\boldsymbol{f}^{\text{Observed}}$, $\boldsymbol{f}^{\text{Uniform}}$ and $\boldsymbol{f}^{\text{Universal}}$ (the dependent variables). These model-implied quantities are defined in Section 6.3.



Tables

Table I: S&P 500 Index Fund Observed Fees vs. Model-Implied Fees

This table displays panel regressions of observed total fees f_{it}^{obs} against calibration-predicted fees $\widehat{f}_{it}$ over the period 2006-2016. All fees vary with time (year t) and are aggregated to the same unit of observation within the same regression specification for consistency: observed fees f_{it}^{obs} are for non-institutional share classes of S&P 500 index mutual funds aggregated up to the level of the fund family i (in case a family offers more than one index fund share class), and calibration-predicted fees $\widehat{f}_{it}$ are produced by the baseline calibration conducted using family-level consideration patterns. Column (2) controls for fund attribute, and column (3) also controls for family attributes, where available.

Dependent Variable:	f_{it}^{obs}		
	(1)	(2)	(3)
$\widehat{f}_{it}$	1.168**	1.283***	1.270***
	(0.3949)	(0.3539)	(0.3301)
Log Size$_{it}$		0.0002	0.0005
		(0.0003)	(0.0003)
Log Downloaders$_{it}$		-0.0007	-0.0002
		(0.0004)	(0.0006)
Log Family Size$_{it}$			-0.0009*
			(0.0005)
Family Equity Share$_{it}$			-0.00003
			(0.00002)
Family Passive Share$_{it}$			-0.0005
			(0.00003)
Family Fund Count$_{it}$			1.16×10^{-6}
			(8.52×10^{-6})
Family Asset Class Diversity$_{it}$			-3.27×10^{-7}
			(3.29×10^{-7})
Family Return Diversity$_{it}$			1.21×10^{-8}
			(3.87×10^{-8})
Year FEs	✓	✓	✓
N	693	693	475
R^2	0.07200	0.09414	0.14886

Double-clustered (Family & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table II: Recomputing Calibrated Fees With Singleton Consideration Sets Excluded

This table contains OLS & panel regressions of observed total fees f_{it}^{obs} against calibration-predicted fees $\widehat{f}_{it}$ over the period 2006-2016. All fees vary with time (year t) and are aggregated to the same unit of observation within the same regression specification for consistency: observed fees f_{it}^{obs} are for non-institutional share classes of S&P 500 index mutual funds aggregated up to the level of the fund family i (in case a family offers more than one index fund share class), and calibration-predicted fees $\widehat{f}_{it}$ are produced by the baseline calibration conducted using family-level consideration patterns. Columns (1) and (2) use predicted fees from the baseline calibration, in which all consideration sets are employed. Columns (3) and (4) use predicted fees from a separate calibration wherein singleton consideration sets are excluded when computing the competition network.

Dependent Variable:	f_{it}^{obs}			
Singletons:	Included		Excluded	
	(1)	(2)	(3)	(4)
$\widehat{f}_{it}$	1.168**	1.283***	1.260***	1.347***
	(0.3949)	(0.3539)	(0.3726)	(0.3340)
Log Size$_{it}$		0.0002		0.0002
		(0.0003)		(0.0003)
Log Downloaders$_{it}$		-0.0007		-0.0007
		(0.0004)		(0.0004)
Year FEs	✓	✓	✓	✓
N	693	693	693	693
R^2	0.07200	0.09414	0.07660	0.09778

Double-clustered (Family & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table III: Interactions with Indicators for Pure Direct Distribution

This table contains panel regressions of the observed residual expense ratio,

$$\text{Residual Expense Ratio}_{i,t} = \text{Expense Ratio}_{i,t} - \text{Actual 12b-1 Fee}_{i,t}, \tag{34}$$

against calibration-predicted fees $\widehat{f}_{it}$ over the period 2006-2016. All fees vary with time (year t) and are aggregated to the same unit of observation within the same regression specification for consistency: observed fees f_{it}^{obs} are for non-institutional share classes of S&P 500 index mutual funds aggregated up to the level of the fund family i (in case a family offers more than one index fund share class), and calibration-predicted fees $\widehat{f}_{it}$ are produced by the baseline calibration conducted using family-level consideration patterns. Specifications (2) and (3) include an interaction with an indicator for whether the family pays no distribution fees (thus following a purely direct distribution channel), with distribution fee measured as either the actual 12b-1 fee (column 2) or an effective 12b-1 fee (column 3; see Roussanov, Ruan, and Wei, 2021).

Dependent Variable:	Observed Residual Expense Ratio$_{i,t}$		
	(1)	(2)	(3)
$\widehat{f}_{it}$	0.8363***	0.6218**	0.6110**
	(0.2347)	(0.2486)	(0.2556)
Log Size$_{it}$	0.00004	0.00004	0.00004
	(0.0001)	(0.0002)	(0.0002)
Log Downloaders$_{it}$	-0.0008**	-0.0009**	-0.0009***
	(0.0003)	(0.0003)	(0.0003)
$\mathbb{1}\{\text{Purely Direct Distribution}\}_{i,t}$		-0.0018**	-0.0018**
		(0.0008)	(0.0008)
$\widehat{f}_{it} \times \mathbb{1}\{\text{Purely Direct Distribution}\}_{i,t}$		0.2998***	0.3484***
		(0.0096)	(0.0134)
Distribution Fee Definition		Actual 12b-1	Effective 12b-1
Year FEs	✓	✓	✓
N	693	693	693
R^2	0.11660	0.15026	0.15120

Double-clustered (Family & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table IV: Frequency of Co-Consideration vs. Pairwise Attribute Distance

This table contains Poisson regressions of the number of unique EDGAR users who considered both fund family pairs i, j during a calendar year t, with additive pairwise fixed effects, over the period 2006-2016. Each covariate represents the distance between the cross-sectional attribute percentiles x of the two fund families that published the prospectus pairs i, j. Self-joins are excluded ($i \neq j$). The distance between the pair of values x_i, x_j for some attribute x is defined to be $|x_i - x_j|$. Attribute percentiles x_i, x_j and the distances between them $|x_i - x_j|$ are all in the range $[0, 1]$. The "size" attribute is the log of the Total Net Assets (TNA) of the fund family. The "share" attributes measure the share of TNA invested in a given asset class (equity, fixed income, mixed), or sold as a share class to retail investors ("retail share"), or following a passive indexing strategy ("passive share").

Dependent Variable:	Co-considerations$_{i,j,t}$		
	(1)	(2)	(3)
Total Fee Distance$_{i,j,t}$	-0.0774***	-0.0864***	-0.0609***
	(0.0243)	(0.0200)	(0.0161)
Age Distance$_{i,j,t}$	-0.2671***	-0.1650***	-0.1587***
	(0.0386)	(0.0179)	(0.0177)
Size Distance$_{i,j,t}$	-0.3106***	-0.2054***	-0.2091***
	(0.0571)	(0.0311)	(0.0267)
Fund Count Distance$_{i,j,t}$	-0.5880***	-0.5524***	-0.3953***
	(0.0534)	(0.0397)	(0.0389)
Equity Share Distance$_{i,j,t}$	-0.1226***	-0.0212	-0.0070
	(0.0417)	(0.0291)	(0.0117)
Fixed Income Share Distance$_{i,j,t}$	0.0596	0.0244	-0.0126
	(0.0418)	(0.0293)	(0.0189)
Mixed Assets Share Distance$_{i,j,t}$	-0.0824	-0.1248**	-0.0831**
	(0.0587)	(0.0543)	(0.0359)
Retail Share Distance$_{i,j,t}$	-0.1040***	-0.0724***	-0.0973***
	(0.0217)	(0.0156)	(0.0123)
Passive Share Distance$_{i,j,t}$	-0.4789***	-0.6423***	-0.6691***
	(0.1450)	(0.1140)	(0.1061)
Family i FEs	✓	✓	
Family j FEs	✓	✓	
Year t FEs		✓	
Family i–Year t FEs			✓
Family j–Year t FEs			✓
N	702,472	702,472	702,472
R^2	0.71280	0.81200	0.93206

Clustered (Family i & Family j & Year t) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table V: County-Level Determinants of Consideration Set Size

This table presents regression results analyzing the predictors of investor consideration set sizes. The unit of analysis is at the geographic county level. All covariates are z-scores (standardized to have zero means and unit standard deviations) so that relative magnitudes are comparable; the non-standardized covariates are summarised in Table IA.2.

Dependent Variable:	Consideration set size (mean)	
	(1)	(2)
(Intercept)	1.560***	
	(0.0530)	
Age (median)	-0.0240**	-0.0284**
	(0.0103)	(0.0126)
Household income (median)	0.0611**	0.0773**
	(0.0234)	(0.0249)
Home value (median)	-0.0205	-0.0172
	(0.0248)	(0.0230)
College graduates (share)	0.0854***	0.0711**
	(0.0196)	(0.0241)
Urban residents (share)	0.0552***	0.0423***
	(0.0128)	(0.0122)
Democrat voters (share)	0.1650***	0.1514***
	(0.0383)	(0.0309)
Year FEs		✓
N	18,209	18,209
R^2	0.01952	0.02876

Clustered (Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Internet Appendix for "Mutual Fund Families, Market Structure, and Fee Competition"

A Fees and Characteristics of S&P 500 Trackers 2
A.1 Identifying Index Funds 2
A.2 Fees 3
A.3 Characteristics of Funds and Their Families in Post-Calibration Analyses . 5
B SEC EDGAR 7
C County-Level Investor Attributes 13
D Calibrated Competition Network and Parameters 14
E Additional Calibration Tests 17
F Consideration Patterns When Singleton Consideration Sets Are Excluded 20
References in Internet Appendix 22

A Fees and Characteristics of S&P 500 Trackers

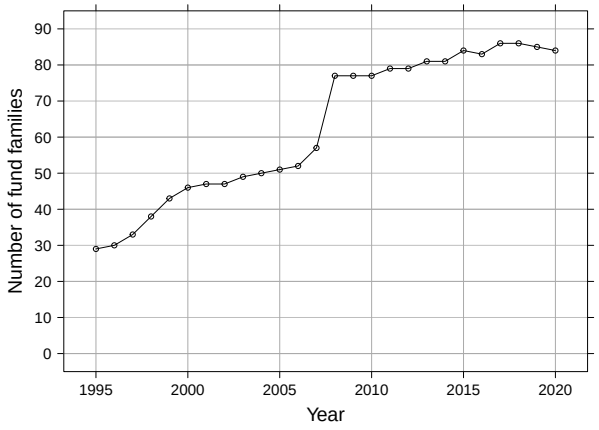
A.1 Identifying Index Funds

To identify mutual funds that track the S&P 500 Index, we compare funds' historical returns to those of the S&P 500 Index, and first select funds whose returns have a correlation of 0.98 or higher with those of the index. We then exclude funds whose names indicate they do not follow a passive tracking strategy, including funds with tilts, screens and mixed asset classes.

Since our unit of analysis in this paper is at the family level, Figure IA.1 plots the number of fund families offering S&P 500 index funds over time.

Figure IA.1: Fund Families Offering S&P 500 Index Funds

This chart displays the number of mutual fund families in our sample that offer S&P 500 index funds. These fund families file prospectuses with the SEC, and can be matched to the CRSP Mutual Fund Database.



A.2 Fees

We now document that the considerable fee dispersion of S&P 500 index funds first observed by Hortaçsu and Syverson (2004) and Elton, Gruber, and Busse (2004) has persisted over time to the present day.

We compute fees at the fund family level, aggregating by TNA to allow for the possibility a fund family can offer multiple S&P 500 index funds. Figure IA.2 presents cross-sectional dispersions of total fees, and shows that fee dispersion remains substantial today, far beyond the end of Hortaçsu and Syverson (2004)'s sample period. Although we have a different procedure for aggregating fees than Hortaçsu and Syverson (2004) do, the statistics for our overlapping period of 1995-2000 are similar. Other studies likewise document the persistence of fee dispersion to the present day (with varying methodologies): Cooper, Halling, and Yang (2021) focus on S&P 500 funds, and Brown et al. (2023) focus on a broad class of index funds.

Figure IA.3 shows that the rank-autocorrelation of the cross-section of fees is very high, which captures the fact that *relative* fee levels (within the entire cross-section) tend to persistent over time.

Figure IA.2: S&P 500 Total Fee Dispersion Over Time

This chart displays annual cross-sectional standard deviation of S&P 500 index mutual fund total fees (expressed as an annualized percentage).

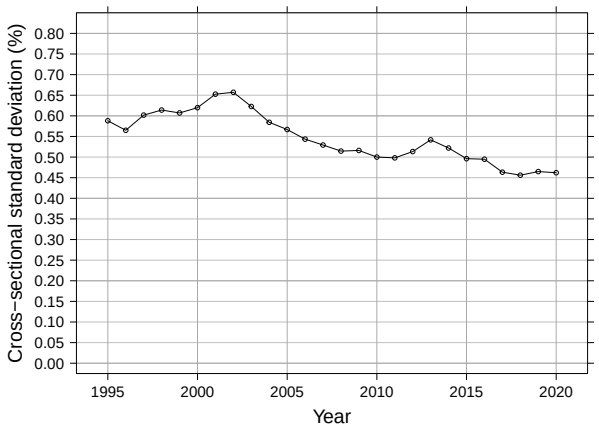
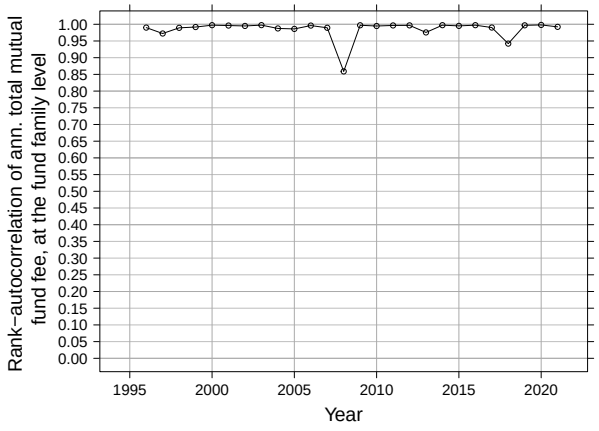


Figure IA.3: S&P 500 Total Fee Rank-Autocorrelation Over Time

This chart displays annual cross-sectional rank-correlation of S&P 500 index mutual fund total fees in year $t-1$ and t.



A.3 Characteristics of Funds and Their Families in Post-Calibration Analyses

Table IA.1 summarizes fund and family characteristics used in our post-calibration tests (i.e. Tables I, II, and III).

As in our tests, we aggregate characteristics to the family level (using total net assets as weights).

Table IA.1 also display summary statistics for fees, at the same level of aggregation. As in our main tests, we aggregate these to the family level, and we restrict our focus to non-institutional share classes when computing fees. Note that Hortaçsu and Syverson (2004, Figure 2) found non-institutional share classes have higher fees than institutional ones.

Additional Details on Computing Family-Level Characteristics. CRSP MFDB flags are used in computing family asset class-based characteristics.

To compute family passive shares, we proxy for whether individual funds are passive using the simple name-based procedure of Ben-David, Li, Rossi, and Song (2022); our resulting measure is thus a lower-bound on the true degree of indexing in a family's fund offering.

There are fewer observations for the fund family count due to this variable being merged in from the SEC's "Investment Company Series and Class Information" dataset, which only begins in 2010. We use this dataset to count the number of funds because it is not limited to open-ended mutual funds, unlike the CRSP Mutual Fund Database.

Table IA.1: S&P 500 Fund and Family Characteristics

This table summarizes the characteristics of the S&P 500 funds and their families that we employ in our post-calibration tests.

Statistic	N	Mean	S.D.	Min.	Lower quartile	Median	Upper quartile	Max.
Total Fee (%)	693	0.47	0.51	0.00	0.01	0.26	0.91	2.08
Expense Ratio (%)	693	0.37	0.39	0.00	0.01	0.21	0.60	1.67
Effective 12b-1 Fee (%)	693	0.17	0.27	0.00	0.00	0.02	0.18	1.00
Actual 12b-1 Fee (%)	693	0.07	0.12	0.00	0.00	0.01	0.09	1.00
Size ($ Mln)	693	5,613.54	19,952.50	0.10	186.50	918.30	3,096.50	233,096.40
Unique Downloaders	693	351.81	510.91	4	64	155	387	3,760
Family Size ($ Bln)	693	138.62	349.64	0.08	8.98	42.41	113.79	3,949.19
Family Equity Share (%)	693	56.59	20.91	3.83	45.05	57.61	69.61	100.00
Family Passive Share (%)	693	15.31	21.38	0.00	0.07	4.76	21.32	100.00
Family Fund Count	486	88.53	78.64	2	33	68	124	456
Family Asset Class Diversity (HHI)	693	2,271.69	1,450.48	32.72	1,289.49	1,755.25	2,664.15	8,180.88
Family Return Diversity (log-det)	680	12.10	1,250.58	−9,874.30	−115.55	−1.14	100.86	8,941.34

B SEC EDGAR

This appendix describes the SEC EDGAR website, trends in its usage, and a geographic characterization of its users; this geographic distribution is broad-based and tracks population patterns on observable dimensions.

Mutual Fund Prospectuses. A mutual fund's prospectus is a description of the fund as a product. Namely, it describes the mandate defining the fund's investment strategy (e.g. investment universe, ability to engage in short-selling, take on leverage or trade derivatives), the fund's past performance, the fees charged by the fund, and other key fund attributes. The prospectus must be made available to investors on request, and must also be filed in the SEC's EDGAR system, after which it is publicly available via the associated EDGAR website.

The SEC encourages investors to read and compare the prospectuses of mutual funds before making an investment. For example, the SEC (2019) advises, *"You can research a mutual fund or ETF by reading its prospectus carefully to learn about its investment strategy and the potential risks. [...] Investors can — and should — request and read the mutual fund's prospectus before making an investment decision ."*

Fund Families File Prospectuses. The SEC requires fund prospectuses to be filed as (part of) one of three forms: Form N-1A must be filed by all open-ended funds upon initial registration, and one of either Form 485APOS or 485BPOS must be filed annually to keep the prospectus up to date. The forms ensure fund descriptions are provided in a consistent, electronic format, which facilitates comparisons across funds by investors.

A fund prospectus is filed by the fund family, rather than an individual fund. Since fund families typically offer a variety of mutual fund products, their prospectus filings (N-1A, 485APOS, 485BPOS, and their amendments) may contain multiple individual mutual fund prospectuses, collated as part of the same document. Conversely, each document is associated with only one fund family.

Defining Mutual Fund Families. All SEC EDGAR filings, including fund prospectuses, are tagged with a Central Index Key (CIK) that identifies the legal entity making the filing. To allow for the possibility that multiple CIKs may exist under the umbrella of the same fund family, we follow the mutual fund literature and use parsed family names in order to explicitly identify the fund families (e.g. Dannhauser and Spilker, 2023). First, we look up the name of the legal entity identified by the CIK, and additionally match CIKs to entries in the CRSP Survivor-Bias-Free US Mutual Fund Database using the "CRSP_cik_map" table. Then, we determine and hand-validate the mutual fund family responsible for filing each prospectus by parsing the

names of the legal entity (recorded by the SEC) and the management family (recorded by CRSP).

Popularization of the EDGAR Website. Prior to the SEC launching the EDGAR website, it was not easy for investors to search for and obtain mutual fund prospectuses. A survey by the Investment Company Institute (1997) found that only 51% of fund families had a website at that point in time, and that fund prospectuses were available for download on only 32% of those websites. Thus for the majority of fund families, an investor would either have to directly contact the family and request a prospectus for one of its funds or indirectly obtain a copy from an intermediary.

However on 12 June 2006 investors' ability to search for and compare prospectuses greatly improved. On that date the SEC introduced two new tools on its free and publicly-accessible EDGAR website: a tool for searching for and downloading regulatory filings by mutual fund names and tickers, and a tool for searching through the text of mutual fund prospectuses (and other types of filing).[28] Screenshots of these two search tools are displayed in panels (a) and (b) of Figure IA.4. The indexing of prospectuses for these search tools also made prospectuses easily accessible via general search engines such as Google and Bing. For example, the screenshot Figure IA.4(c) shows the results returned by a generic Google search for the prospectus of a BlackRock equity fund. A link to the EDGAR-hosted prospectus document is listed just after a link to the fund's own product page. Therefore, as a consequence of the SEC introducing these EDGAR search tools in June 2006, the effort and time investors have to spend to obtain fund prospectuses has been substantially reduced.

EDGAR Download Logs. This paper exploits data on the usage of the SEC EDGAR website to measure the extent to which investors compare fund families. The full public dataset comprises timestamped records of all documents downloaded from the EDGAR website during the period 14 February 2003 to 30 June 2017. We focus our analysis on the sample from 2006 – when mutual fund search tools became available – to the end of 2016, the latest full calendar year. Each download by a website visitor is associated with an identifier for the visitor's Internet Protocol (IP) address. The dataset allows individual IP addresses to be tracked through time, thereby identifying downloads with a downloader. All downloaders are potential mutual fund investors.

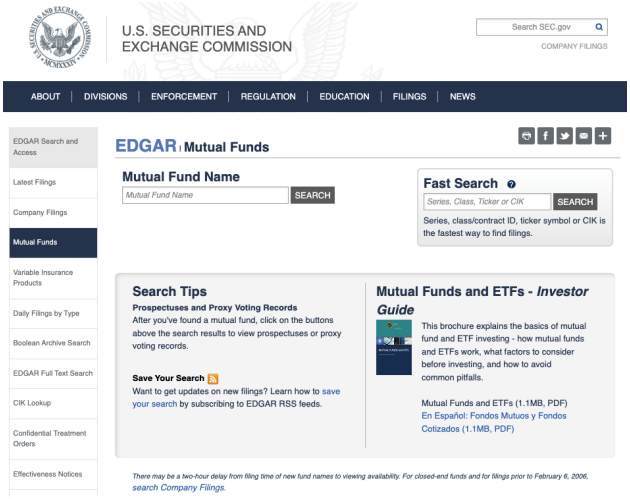
We measure downloads on an annual basis. Each user who downloads at least one prospectus

[28] See SEC press release No. 2006-93: "SEC Announces That Over a Million Corporate and Mutual Fund Reports Are Now Fully Searchable via Web", `https://www.sec.gov/news/press/2006/2006-93.htm`, accessed on 9th July 2022. Note also that the introduction of these two tools followed the adoption of Rule 33-8590, which instituted a system for identifying individual mutual funds on EDGAR.

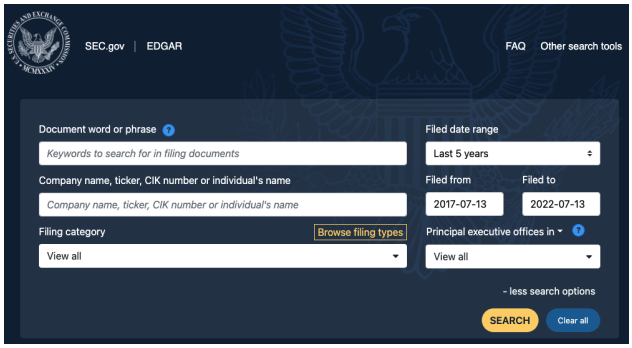
Table IA.1: Searching for Prospectuses

Sub-figures (a) & (b) are screenshots of the two search tools added to the SEC's EDGAR website on 12 June 2006. The mutual fund search tool is accessible at `https://www.sec.gov/edgar/searchedgar/mutualsearch.html`, and the text search tool is accessible at `https://www.sec.gov/edgar/search/`. Sub-figure (c) is a screenshot of results from a Google search for a BlackRock fund prospectus, with a link to an EDGAR-hosted prospectus highlighted. (Search conducted on 18 April 2023.)

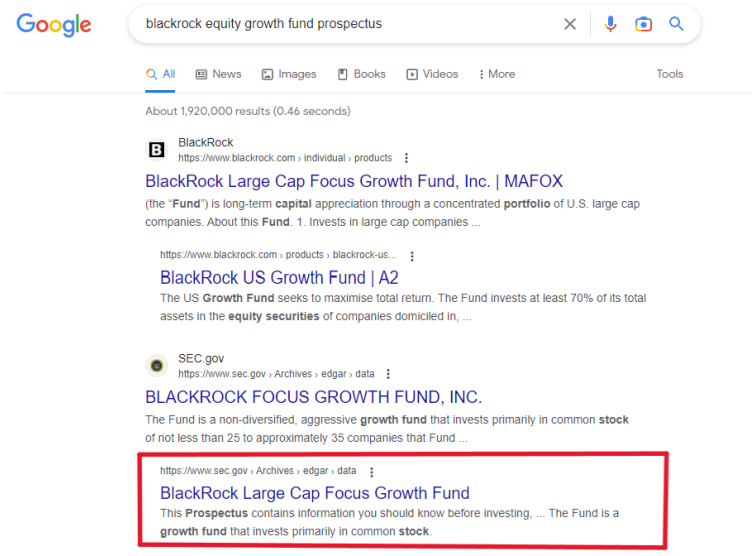
(a) EDGAR mutual fund search tool



(b) EDGAR text search tool



(c) Example Google search result linking to EDGAR



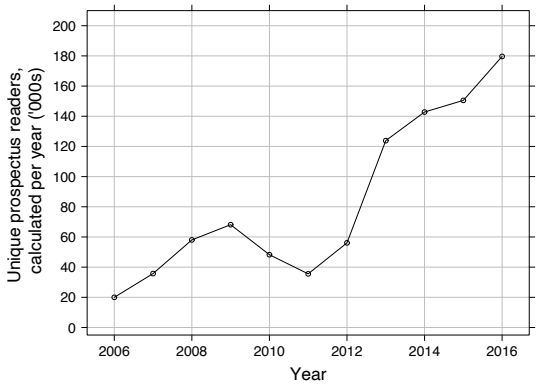
document (Form N-1A, 485APOS or 485BPOS) during a given year is included in our sample. An individual's *download set* consists of the unique fund families for which she downloads a prospectus during that calendar year.

The presence of very large download sets representing systematic downloads by robo-downloaders could introduce bias into our analysis, because only the download patterns of individual humans measure fund comparisons by investors. Therefore we follow Lee, Ma, and Wang (2015) in identifying and removing robot users from the sample. First, we remove any record that the SEC's systems have classified as being a web crawler. Then, we count the number of entities per day for which each user has downloaded any Form type (including non-mutual fund related Forms) and, if this exceeds 50, discard these records from our sample.

User Count. Examining aggregate usage statistics for the downloaders in our cleaned sample, we find substantial growth in usage of EDGAR to access mutual fund prospectuses since 2006. Figure IA.5 plots the number of unique users of the SEC EDGAR website in each year from 2006 through to 2016 who downloaded at least one fund prospectus document. The annual number of users multiplied tenfold in that decade from 20,000 in 2006 to nearly 200,000 in 2016. The vast majority of that growth occurred after 2011. This growth suggests EDGAR has become a go-to source of mutual fund prospectuses.

Figure IA.5: SEC EDGAR Usage for Prospectus Downloads is Widespread

This figure displays the annual number (in thousands) of unique EDGAR users who downloaded at least one prospectus from the EDGAR website during that year. EDGAR users are uniquely identified based on their IP addresses.



Breadth of Usage. To understand who uses EDGAR to obtain prospectuses, we begin by measuring their geographic location. We de-anonymize individual IP addresses in the EDGAR

logs following the procedure of Chen et al. (2020). We then map these IP addresses to US counties using the historical geolocation tables provided by MaxMind. These historical tables account for the possibility of IP addresses changing their physical location over time.

We now compare the geographic distribution of EDGAR users with the national population. Figure IA.6(a) plots the number of unique EDGAR users in each US state over the period 2006 to 2016. The figure reveals EDGAR users are more prevalent in more populous states, such as California, New York, Texas, Florida and Pennsylvania. The correlation between these EDGAR user counts and state populations according to the US Census Bureau is 0.89 (the correlation at the county level is 0.73). This correlation with the distribution of the US population is inconsistent with EDGAR users being financial professionals because they would be highly concentrated in New York and Massachusetts where many investment companies are headquartered (Kim, Wang, and Wang, 2022). Figure IA.6(b) visualizes how county-level correlations between unique EDGAR users and the total population vary from state to state. All such correlations are positive, and almost uniformly high. Therefore, on observable geographic dimensions, EDGAR prospectus downloaders are widely dispersed and track population patterns closely. This evidence is reassuring for external validity, and is consistent with empirical evidence by Liaukonytė and Žaldokas (2022) linking EDGAR search activity to Google search activity and retail trading activity.

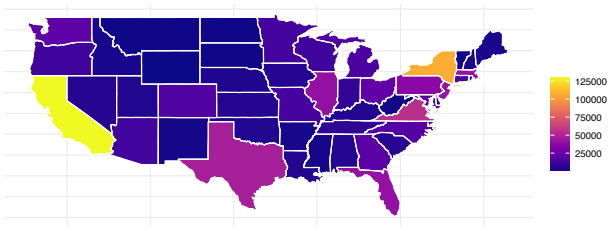
If EDGAR is also used by investment advisors who provide recommendations to retail investors (or by other intermediaries), the resulting download sets remain informative about which fund families are compared in the course of retail investment decisions, because recommendations and due diligence ultimately rely on prospectus disclosures. This is the case because the clients of the investment advisors obtaining prospectuses from EDGAR will inherit the download sets of their advisors' downloads via their advisors' recommendations.

We also argue any bias in the size of EDGAR-user download sets will be to the higher side, even though it is possible for investors to entirely bypass EDGAR when obtaining fund prospectuses (or be ignorant of its existence). Given the physical costs of searching for an additional prospectus are effectively zero once an investor is already an EDGAR user, and likely higher when the investor is instead sourcing prospectuses from fund websites or brokers, EDGAR users are likely to obtain prospectuses for more fund families than the average investor not using EDGAR. Yet as we document in Section 2, the majority of EDGAR website users download prospectus documents that have been filed by a single fund family, and 90% of users limit their prospectus acquisitions to only 3-4 different families in total; therefore, the magnitude of any such bias is likely to be limited.

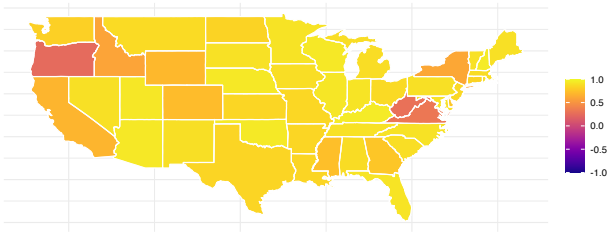
Figure IA.6: SEC EDGAR Prospectus Downloaders Are Broadly Distributed Geographically.

Both charts visualize statistics derived from unique EDGAR user counts over the sample period 2006-2016. Users are geolocated based on their IP addresses. Panel (a) displays the number of unique EDGAR prospectus downloaders per state; lighter states have more unique users, colored according to the legend. These EDGAR user counts are highly correlated with total population counts. Panel (b) displays the correlations of unique downloader counts with county population counts, for each state. Lighter states exhibit higher correlations, colored according to the legend.

(a) State-level unique user counts.



(b) Within-state user-to-population correlations.



C County-Level Investor Attributes

This appendix describes the sociodemographic characteristics of EDGAR users who downloaded prospectuses, at a county-level.

Our ability to geolocate EDGAR website users allows us to proxy for investor attributes via aggregated demographic data. We use the geographic counties in which downloaders are physically located (based on their IP addresses) to merge in county-level demographic data from the US Census Bureau's annual American Community Survey, as well as Presidential election data and the Census Bureau's Rural/Urban Classifications.

Table IA.2 displays the means, standard deviations, and extreme values of these proxies for investors' sociodemographic characteristics. The summary statistics indicate considerable variation.

Table IA.2: County-level Demographic Variables

This table summarizes the county-level covariates that are first standardized and then used in the regressions of Table V. Counties without any prospectus downloads are excluded. Each variable is updated annually except for Democrat voter share (which is updated every 4 years, following the Presidential election cycle).

Variable	N	Mean	St. Dev.	Min	Max
Consideration set size (mean)	19,869	1.544	1.879	1.000	67.000
Age (years, county median)	19,869	39.459	4.931	21.700	66.400
Household income (USD'000s, county median)	19,869	52.481	13.696	19.264	129.588
Home value (USD'000s, county median)	19,869	165.091	98.915	31.400	1,000.001
College graduates (%, share of county)	19,869	23.563	10.028	4.900	78.800
Urban residents (%, share of county)	19,869	56.720	28.182	0.000	100.000
Democrat voters (%, share of county)	19,869	41.378	14.091	5.773	92.457

D Calibrated Competition Network and Parameters

Section 4.3 of the paper presents results from calibrating family-to-family competitive ties using investor's prospectus downloads from the SEC's EDGAR website in 2016. This appendix reports additional details about the calibrated competition network.

Calibrated Competition Network. The calibrated competition network is densely connected. Figure IA.7 plots the network. The dense connections are clearly visible, but edges connecting individual pairs of fund families are obscured when all plotted together. Therefore, Figure IA.8 plots the distribution of network degrees for all fund families. Network degree is measured by the number of edges directly connecting a family to other families in the competition network. The distribution is skewed, with the vast majority of families being directly connected to most other families. These high degrees for most fund families correspond to the overall network's density. In total, 12.7% of families directly compete with all other families, while only 22.0% of families directly compete with fewer than 90% of other families.

Other Calibrated Parameters. Calibrated values of $\widehat{\rho}$ are as follows:

Year	$\widehat{\rho}$
2006	0.973
2007	0.950
2008	0.977
2009	0.986
2010	0.992
2011	0.995
2012	0.995
2013	0.986
2014	0.969
2015	0.979
2016	0.990
2017	0.980

Figure IA.7: Calibrated Competition Network

This figure visualizes the competition network connecting fund families calibrated from investors downloads of family prospectus documents from the SEC's EDGAR website in 2016.

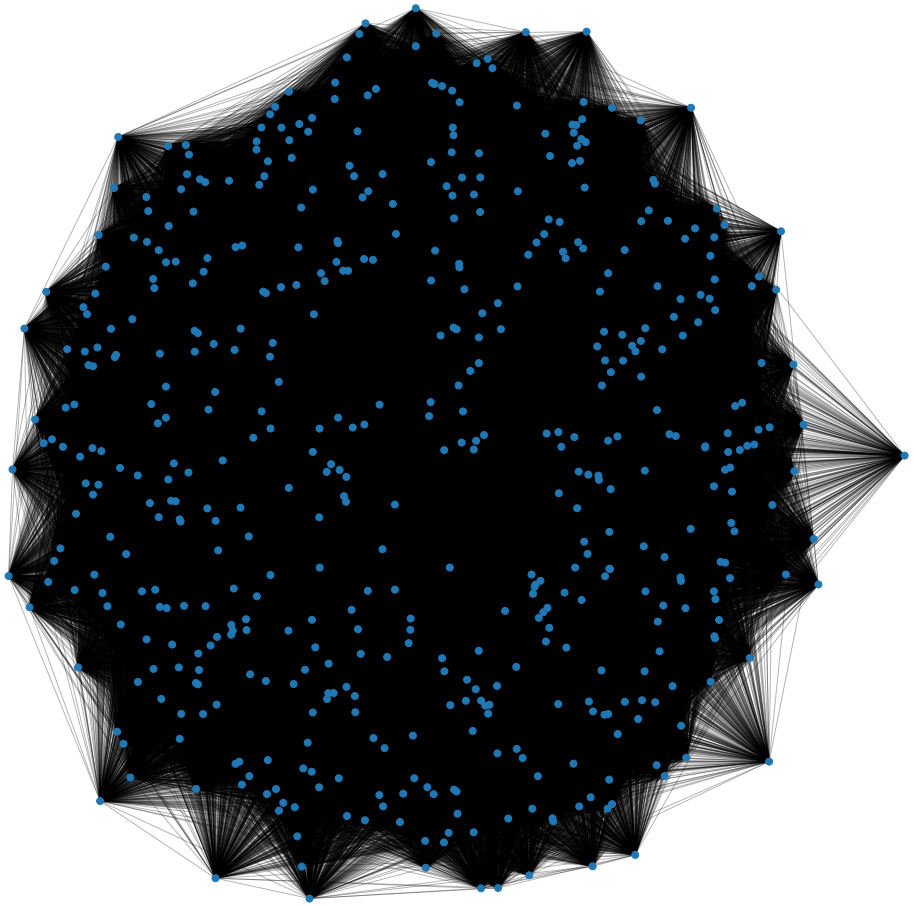
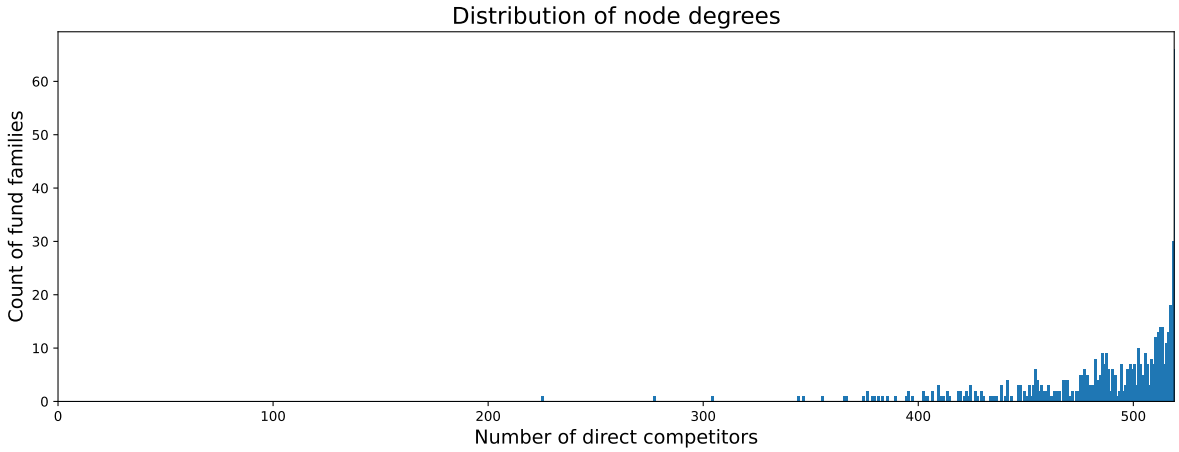


Figure IA.8: Distribution of Family Network Degrees

This chart is a histogram of node degree of the fund families in our calibrated competition network. Node degree is a measured by the number of edges connecting a node to other nodes in the network.



E Additional Calibration Tests

This appendix verifies the robustness of our conclusions based on the main calibration tests, using two variants of our baseline specification.

Table IA.3 focuses on observed expense ratios; i.e., excluding front loads from the fee calculation.

Table IA.4 focuses only on established funds; i.e., those with at least $10 Million in assets and an age of 12 months.

Table IA.3: S&P 500 Index Fund Observed Expense Ratios vs. Model-Implied Fees

This table displays panel regressions of observed expense ratios f_{it}^{obs} against calibration-predicted fees $\widehat{f}_{it}$ over the period 2006-2016. All fees vary with time (year t) and are aggregated to the same unit of observation within the same regression specification for consistency: observed fees f_{it}^{obs} are for non-institutional share classes of S&P 500 index mutual funds aggregated up to the level of the fund family i (in case a family offers more than one index fund share class), and calibration-predicted fees $\widehat{f}_{it}$ are produced by the baseline calibration conducted using family-level consideration patterns. Column (2) controls for fund attribute, and column (3) also controls for family attributes, where available.

Dependent Variable:		$\widehat{f}_{it}$	
	(1)	(2)	(3)
$\widehat{f}_{it}$	0.8990***	0.9915***	1.017***
	(0.2776)	(0.2460)	(0.2091)
Log Size$_{it}$		0.00008	0.0004*
		(0.0002)	(0.0002)
Log Downloaders$_{it}$		-0.0007*	-0.00004
		(0.0003)	(0.0004)
Log Family Size$_{it}$			-0.0012**
			(0.0004)
Family Equity Share$_{it}$			-0.00003
			(0.00002)
Family Passive Share$_{it}$			-0.00002
			(0.00002)
Family Fund Count$_{it}$			6.2×10^{-6}
			(6.54×10^{-6})
Family Asset Class Diversity$_{it}$			-1.6×10^{-7}
			(2.37×10^{-7})
Family Return Diversity$_{it}$			-1.29×10^{-8}
			(2.33×10^{-8})
Year FEs	✓	✓	✓
N	693	693	475
R^2	0.07774	0.11198	0.21738

Double-clustered (Family & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table IA.4: Established S&P 500 Index Fund Observed Fees vs. Model-Implied Fees

This table displays panel regressions of observed total fees f_{it}^{obs} against calibration-predicted fees $\widehat{f}_{it}$ over the period 2006-2016. All fees vary with time (year t) and are aggregated to the same unit of observation within the same regression specification for consistency: observed fees f_{it}^{obs} are for non-institutional share classes of S&P 500 index mutual funds aggregated up to the level of the fund family i (in case a family offers more than one index fund share class), and calibration-predicted fees $\widehat{f}_{it}$ are produced by the baseline calibration conducted using family-level consideration patterns. Column (2) controls for fund attribute, and column (3) also controls for family attributes, where available. Only established funds, with at least $10 Million in TNA and an age of at least 12 months, are included in the sample.

Dependent Variable:	f_{it}		
	(1)	(2)	(3)
$\widehat{f}_{it}$	1.240***	1.305***	1.362***
	(0.3645)	(0.3419)	(0.3576)
Log Size$_{it}$		-0.0001	0.0003
		(0.0003)	(0.0004)
Log Downloaders$_{it}$		-0.0006	-0.0002
		(0.0005)	(0.0006)
Log Family Size$_{it}$			-0.0008
			(0.0005)
Family Equity Share$_{it}$			-0.00003
			(0.00003)
Family Passive Share$_{it}$			-0.00005*
			(0.00003)
Family Fund Count$_{it}$			1.82×10^{-6}
			(1.23×10^{-5})
Family Asset Class Diversity$_{it}$			-4.32×10^{-7}
			(1×10^{-5})
Family Return Diversity$_{it}$			-3.09×10^{-8}
			(1×10^{-5})
Year FEs	✓	✓	✓
N	635	635	446
R^2	0.08844	0.11060	0.15042

Double-clustered (Family & Year) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

F Consideration Patterns When Singleton Consideration Sets Are Excluded

This appendix demonstrates that the stylized facts about investors' patterns of consideration (described in Section 2) continue to hold when singleton consideration sets are eliminated from the sample.

Figure IA.9 shows that consideration sets remain sparse, confirming that investors exhibit limited consideration in the present sample.

As before, limited consideration goes hand-in-hand with heterogeneous consideration: Figure IA.10 shows that the probability of consideration of any one entity remains very low.

Figure IA.9: Investor Consideration Set Size, Excluding Singletons

This chart summarizes the distribution of the number of fund families included in each potential investor's consideration set. A potential investor's consideration set is measured using her prospectus downloads from the SEC EDGAR website each year. Singleton consideration sets are excluded from the sample. The y-axis represents the number of fund families, in log-scale.

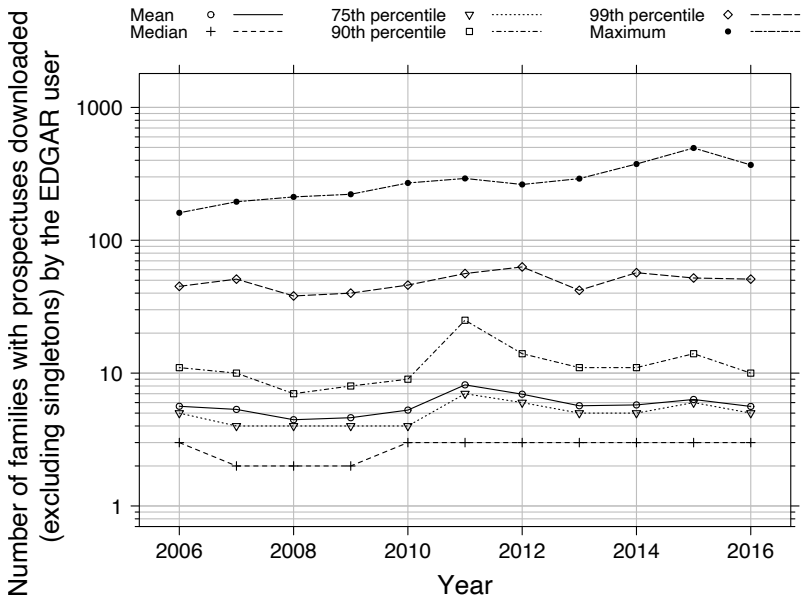
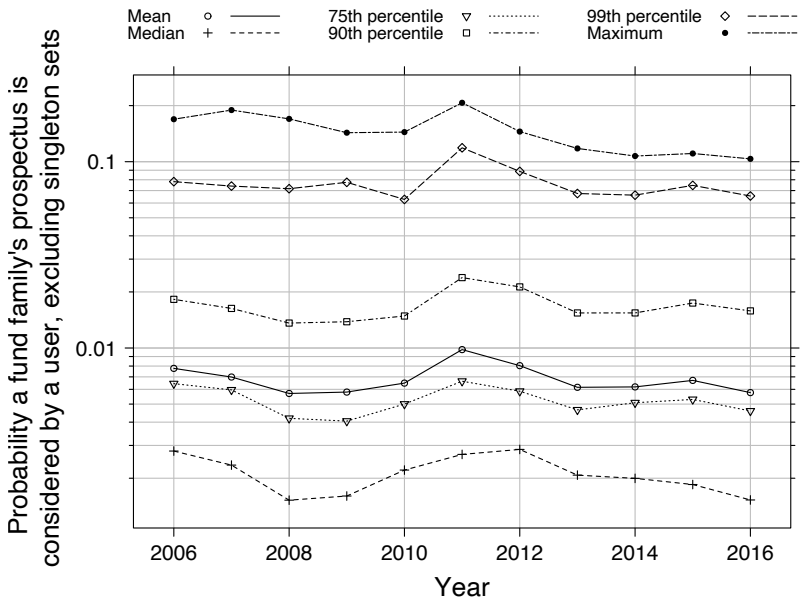


Figure IA.10: Probability of a Fund Family to be Included in a Consideration Set, Excluding Singletons

This chart summarizes the distribution of the frequency of inclusion of each fund family in any potential investor's consideration set. A potential investor's consideration set is measured using her prospectus downloads from the SEC EDGAR website each year. Singleton consideration sets are excluded from the sample. The y-axis represents the probability $\in [0, 1]$ of consideration set membership, in log-scale.



References in Internet Appendix

Ben-David, I., Li, J., Rossi, A., and Song, Y. (2022). What do mutual fund investors really care about? *Review of Financial Studies, 35*(4), 1723–1774.

Brown, Z. Y., Egan, M. L., Jeon, J., Jin, C., and Wu, A. A. (2023). Why do index funds have market power? Quantifying frictions in the index fund market. *NBER Working Paper No. 31778*.

Chen, H., Cohen, L., Gurun, U., Lou, D., and Malloy, C. (2020). IQ from IP: Simplifying search in portfolio choice. *Journal of Financial Economics*, *138*(1), 118–137.

Cooper, M. J., Halling, M., and Yang, W. (2021). The persistence of fee dispersion among mutual funds. *Review of Finance*, *25*(2), 365–402.

Dannhauser, C. D., and Spilker, H. D. (2023). The modern mutual fund family. *Journal of Financial Economics*, *148*(1), 1–20.

Elton, E. J., Gruber, M. J., and Busse, J. A. (2004). Are investors rational? Choices among index funds. *Journal of Finance*, *59*(1), 261–288.

Hortaçsu, A., and Syverson, C. (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics, 119*(2), 403–456.

Investment Company Institute. (1997). How mutual funds use the internet.

Kim, D., Wang, Q., and Wang, X. (2022). Geographic clustering of institutional investors. *Journal of Financial Economics*, *144*(2), 547–570.

Lee, C. M., Ma, P., and Wang, C. C. (2015). Search-based peer firms: Aggregating investor perceptions through internet co-searches. *Journal of Financial Economics*, *116*(2), 410–431.

Liaukonytė, J., and Žaldokas, A. (2022). Background noise? TV advertising affects real-time investor behavior. *Management Science*, *68*(4), 2465–2484.

SEC. (2019). Mutual funds and ETFs: A guide for investors. *SEC Pub. 182 (12/16)*.