

Common Ownership in America: 1980–2017*

Matthew Backus[†] Christopher Conlon[‡] Michael Sinkinson[§]

December 19, 2019

Abstract

We empirically assess the implications of the common ownership hypothesis from a historical perspective using the set of S&P 500 firms from 1980–2017. We show that the dramatic rise in common ownership in the time series is driven primarily by the rise of indexing and diversification and, in the cross-section, by investor concentration, which the theory presumes to drive a wedge between cash-flow rights and control. We also show that the theory predicts incentives for expropriation of undiversified shareholders via tunneling, even in the Berle and Means (1932) world of the widely held firm.

JEL Codes: L0, L21, L13, G34

*Thanks to seminar and conference participants at Bergen, Boston University, Brandeis University, CEPR IO, Cornell, Duke University, EARIE, Federal Trade Commission, Harvard-MIT, IIOC, MACCI - Mergers and Antitrust, Microsoft, NBER, Penn State, Purdue, Red Rock Finance Conference, Rice, Rochester, SITE, The London School of Economics, University of Chicago Booth, University of Pennsylvania, University of California at Santa Cruz, University of Texas at Austin, University of Georgia, University of Virginia, Utah WBEC, Wisconsin and Yale, and conversations with Steve Berry, Wouter Dessein, Wei Jiang, Charles Jones, Amit Khandelwal, Bruce Kogut, Nathan Miller, Fiona Scott Morton, Nancy Rose, Martin Schmalz, and Glen Weyl. All remaining errors are our own.

[†]Columbia University and NBER, matthew.backus@columbia.edu

[‡]New York University, cconlon@stern.nyu.edu

[§]Yale SOM and NBER, michael.sinkinson@yale.edu

1 Introduction

The near-universal assumption in economics is that firms take actions that maximize their own profits. Motivating the assumption, Friedman (1953) contends that investors will discipline firms that do not at least mimic profit-maximizing behavior. Investor’s interests, however, may be complicated by holdings in competing firms, which happens naturally when they seek the benefits of diversification. If firm decision-making is an expression of investor interests, and powerful investors have stakes in competing firms, then one might not expect the firm to simply maximize solely their own profits, yielding oligopoly outcomes, but instead to also value the profit of their competitors when making strategic decisions. The idea that large diversified owners imply nonzero “profit weights” among ostensibly competing firms is known as the *common ownership hypothesis*.

The theoretical framework of the common ownership hypothesis was first articulated in Rotemberg (1984), but it has recently become the subject of a lively public policy debate thanks to empirical work suggesting that the growth of large, diversified common owners may have caused prices to increase among banks and airlines (Azar et al., 2016, 2018).¹ Contemporaneously, Eekhout et al. (2018) argue that markups, economy-wide, have sharply increased since 1980. Combining these lines of work (see Shambaugh et al. (2018)) could go so far as to implicate common ownership in macro-level phenomena such as declining labor share and investment, the productivity slowdown, and diminished “dynamism” of the economy (Gutiérrez and Philippon, 2016).

However appealing this line of thought might be on the theory alone, there are myriad empirical gaps in the argument left to fill. Efforts to test it empirically have been narrowly focused on reduced-form correlations.² There, the null hypothesis is zero effect of common holdings on some outcome of interest, and the alternative — presumed to be due to common ownership — is any effect. This paper builds on that effort by precisely laying out the empirical implications of the common ownership hypothesis, taking seriously its theoretical foundations rather than loading it into an “alternative,” nonzero effect of common holdings in a reduced-form specification.

The payoff to this effort is threefold. First, it casts a light on the sources of variation in

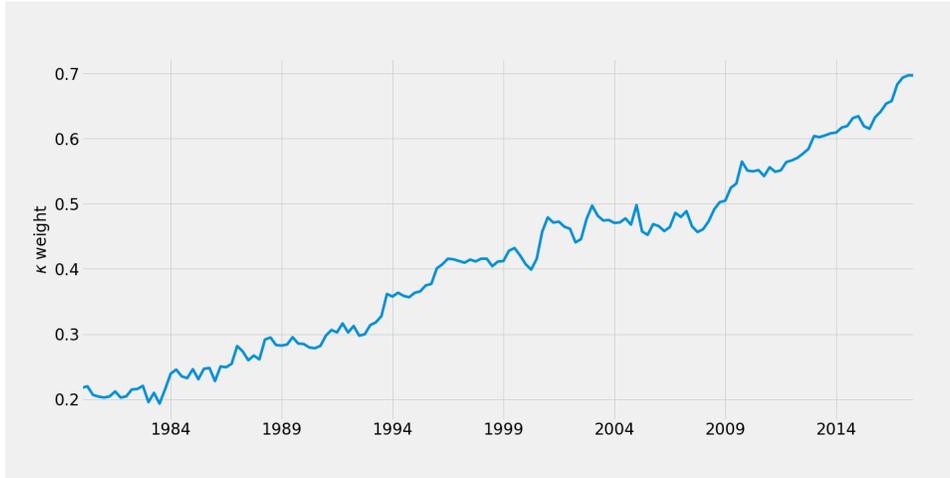
¹The latter paper had over 230 citations as of the date of this draft.

²Exceptions include Kennedy et al. (2017) and Backus et al. (2018).

prior empirical exercises. Much of prior work depends on aggregate measures of common ownership based on the so-called Modified Herfindahl-Hirschman Index, which is function of the common ownership profit weights we study and market shares. We show that between a third and a half of the variation in the profit weight measure comes not from overlapping ownership as many researchers assume, but instead from relative investor concentration, which we make precise in what follows. The role of relative investor concentration depends strictly on a model of corporate governance that defines the relationship between control rights and cash-flow rights, which has been previously unacknowledged in the common ownership literature. Second, taking the theory of common ownership seriously allows us to develop new testable implications. For one, modeling profit weights highlights the asymmetries that emerge, both within markets and even within pairs of firms. Already building on this observation, Boller and Scott Morton (2019) shows cumulative abnormal returns following the entry of a product-market competitor into the S&P 500 that are consistent with the asymmetric implications of the common ownership hypothesis. Additional implications include that the nature of competition should differ in the presence of private firms, and that seemingly-innocuous financial events such as being de-listed from market indices. Taking the theory to its logical conclusion, we show that it is possible for common ownership to create incentives for the “tunneling” of profits from one firm to another. This had been previously thought to be impossible in the Berle and Means (1932) world of the “widely-held firm” due to the absence of a controlling interest. Third, our empirical exercise offers some perspective on the plausibility of these implications. Taking the strict form of common ownership seriously, it could be used to micro-found a tremendous increase in markups between 1980 and 2017, and one might also conclude that over 10% of S&P 500 firms are engaging in tunneling behavior by 2017. We find these predictions to be unrealistically strong — rather, they suggest substantial gaps in our understanding of corporate governance and, in particular, the model of governance that underlies the common ownership hypothesis. In the discussion, we take these conclusions to motivate new directions for future work.

The empirical setting for our exercise is the the full set of S&P 500 Index constituents, from 1980 through the end of 2017. For each pair of firms in each quarter, we compute the profit weights that each firm would place on the other, as implied by the common ownership hypothesis. The time series of the average pairwise profit weights paints a stark picture, depicted in Figure 1. For comparison: a weight of 0 corresponds to what we expect in a world of profit-maximizing firms, and a weight of 1 corresponds to the weight that a merged firm places on an acquired subsidiary business (or, equivalently, full collusion). We find that the average pairwise profit weights implied by the common ownership hypothesis more than

Figure 1: Common Ownership Profit Weights Over Time



Notes: This figure depicts the mean implied profit weight across all pairs of firms in the S&P 500 index by year, excluding own profit weights, which are normalized to 1. See Section 2 for an explicit formula for common ownership weights and derivation.

tripled among S&P 500 firms, from just over 0.2 in 1980 to almost 0.7 in 2017.

We are not the first to show that overlapping ownership is on the rise. Prior work has cast similar pictures in terms of the Modified Herfindahl-Hirschman Index (Gutiérrez and Philippon, 2016; Anton et al., 2016) or proposed altogether new measures, e.g. the measure of Gilje et al. (2019), which they name GGL). We eschew the MHHI index for a number of reasons, from the dependence on Cournot competition to the necessity of defining product markets, which we describe extensively in our other work on the topic (Backus et al., 2019, 2018). GGL offers an alternative, but this measure is particularly unsuited for empirical work on the implications of common ownership for market power. First, the model motivating the measure restricts attention to binary actions by managers in a setting with no strategic interactions, which rules out most models of market power from the start. In contrast, our profit weights approach is fully general as it is based on the firm’s objective function. Second, the GGL measure fails, by design, to weigh own profits and other-firm profits, and so it doesn’t actually convey anything about what firms will do when faced with a trade-off between own profits and competitor profits. All of these measures — profit weights, MHHI, and alternatives — agree on the broad trend in Figure 1. However, the profit weights approach, which starts with the objective function of the firm, is the only one that offers a fully general path forward for empirical study of the common ownership hypothesis. We emphasize that while we are the first to construct our measure — the common ownership profit weights — at this level of breadth, neither the innovation nor their use in empirical work is novel here. The theory goes back as far as Rotemberg (1984), is implicit in the MHHI

measure of Bresnahan and Salop (1986), has been applied to cross-ownership in O’Brien and Salop (2000), and has seen application in various tests of the common ownership hypothesis (Kennedy et al., 2017; Gramlich and Grundl, 2017; Boller and Scott Morton, 2019).

An additional contribution of this paper is a new dataset of institutional holdings of United States publicly traded firms. While most research to date in this area has used a commercial dataset of these holdings (Thomson Reuters), it has been frequently noted that this dataset has gaps in coverage and errors relative to the source documents. As a result, we collected all 13(f) filings from the SEC since electronic filing was made mandatory in 1999 through 2017 and extracted holdings of S&P 500 firms.³ We are making the code and output of this parsing exercise available to other researchers as our alternative dataset appears to provide more complete coverage, particularly during 2010-2014, as further discussed in Section 3. If one were to complete our exercise using only the commercial dataset, one would reach different qualitative and quantitative conclusions, as shown in Appendix Figure 20, which contrasts Figure 1 using the commercial dataset vs our novel dataset.

Our theoretical model also affords us perspective on some of the proposed policy answers to the common ownership hypothesis (Posner et al., 2017). We find that mergers and “break-ups” in the upstream space of institutional managers have a relatively minor effect on the average profit weight. Forcing these firms to abstain entirely from corporate governance would have a large effect on common ownership incentives, but may also have unintended consequences for owners’ abilities to monitor and discipline management. More substantial than either, however, in terms of dampening the expression of common ownership incentives, is the entry of a product market competitor with no overlapping ownership. In a calibrated example, we show that the presence of a “maverick,” e.g. a fully private or foreign-held firm, has a first-order effect on the price implications of the common ownership hypothesis. Between the rise of globalization and growing share of privately-held firms as a fraction of economic activity in the US, this offers some perspective on why the predictions of the common ownership hypothesis may be substantially dampened.

The literature on common ownership is flanked by two related literatures. In economics, it borrows its theoretical foundations from the literature on cross-ownership (Reynolds and Snapp, 1986; Bresnahan and Salop, 1986). These models assume that the firm fully internalizes the incentives of cross-ownership in strategic decision-making. A recent empirical

³A total of 318,038 quarterly filings by institutional investors, including amendments. The total size of the corpus is approximately 25GB.

contribution to this literature, Heim et al. (2019), shows how firms adopt cross-ownership positions in response to the introduction of leniency programs, arguing that this is an attempt to sustain collusive agreements. In finance, the common ownership hypothesis mirrors a large body of work documenting the internalization of cross-incentives implied by holdings of institutional investors. In an early example of tunneling, which we discuss in Section 5.1, Matvos and Ostrovsky (2008) show that institutional investors vote in favor of mergers that seem to damage their own share value, when these interests are offset by gains to holdings in the target firm. Moreover, there is a growing body of work suggesting that when institutional managers hold both debt and equity in a firm, they use the control rights implied by their equity holdings in favor of debtor-friendly policies (Jiang et al., 2010; Keswani et al., 2019). In some sense, the common ownership hypothesis sits at the join of these two literatures, leaning both on the internalization of such incentives by institutional managers as well as the belief that they are communicated from owners to decision-makers within the firm. A careful assessment of the theoretical implications of common ownership is a necessary first step to evaluating that claim.

The structure of the paper is as follows: In Section 2 we outline the theory of common ownership, the derivation of the common ownership profit weights, and finally highlight some novel mathematical features of those weights. In section 3 we describe our data sources as well as the advantages of our scraped dataset over the Thompson Reuters s34. In Section 4 we offer our main descriptive evidence on profit weights from the S&P 500. Section 5 discusses the economic implications of the implied common ownership profit weights through the lens of tunneling and through simulation. We also consider policy remedies. Robustness considerations to various assumptions are addressed in Section 6, and Section 7 concludes.

2 Theoretical Foundations

We begin with a generic setup: a firm f makes a strategic choice x_f and earns profits given by $\pi_f(x_f, x_{-f})$, which depend on their rivals' choices x_{-f} as well. In the standard framework, the profit function is the objective function of the firm, and in this way economists have modeled behavior ranging from pricing to entry to research and development. This framework is motivated by the claim that the firm answers to its investors, who will withdraw capital should the firm fail to at least mimic profit maximization (Friedman, 1953). So, the firm behaves in a way that maximizes π_f because that maximizes shareholder value. This is the

point of departure for the common ownership hypothesis. In a world with common owners, maximizing shareholder value yields a different objective function.

Consider the payoffs of an investor — for our purposes, a shareholder of a publicly traded company. We assume that shareholder s has cash-flow rights denoted β_{fs} , equal to the fraction of firm f that they own. Assumption 1 is that the profit of the shareholder, $\tilde{\pi}_s$, is given by the sum of profits over their portfolio of investments weighted by cash-flow rights,

$$\tilde{\pi}_s = \sum_{\forall g} \beta_{gs} \pi_g. \quad (\text{A1})$$

We call an investor a common owner if $\beta_{fs} > 0$ for multiple firms.

The following derivation is not novel: it follows directly from the objective function proposed by Rotemberg (1984); here we use the notation and formulation of O'Brien and Salop (2000). What it shows is that maximization of shareholder value does not, in general, imply own-firm profit maximization as the standard model assumes.

In the framework of Rotemberg (1984), a firm acts to maximize the profits of shareholders. However, because their portfolios differ, investors will disagree about the optimal strategy. Assumption 2 is that firm f resolves this as a social choice problem, by placing Pareto weights γ_{fs} on the profits of investor s and maximizing the Pareto-weighted sum of their investors' profits. Letting Q_f denote the proposed objective function of the firm, we can derive the weight, κ_{fg} , that firm f places on its competitors g 's profits, π_g , as follows:

$$\begin{aligned} Q_f(x_f, x_{-f}) &= \sum_{\forall s} \gamma_{fs} \cdot \tilde{\pi}_s(x_f, x_{-f}) & (\text{A2}) \\ &= \sum_{\forall s} \gamma_{fs} \cdot \left(\sum_{\forall g} \beta_{gs} \cdot \pi_g(x_f, x_{-f}) \right) \\ &= \sum_{\forall s} \gamma_{fs} \beta_{fs} \pi_f + \sum_{\forall s} \gamma_{fs} \sum_{\forall f \neq g} \beta_{gs} \pi_g \\ &\propto \pi_f + \sum_{g \neq f} \underbrace{\left(\frac{\sum_{\forall s} \gamma_{fs} \beta_{gs}}{\sum_{\forall s} \gamma_{fs} \beta_{fs}} \right)}_{\equiv \kappa_{fg}(\gamma_f, \beta)} \pi_g. \\ &= \pi_f + \sum_{g \neq f} \kappa_{fg}(\gamma_f, \beta) \cdot \pi_g \end{aligned}$$

The second line substitutes in Assumption (A1), and the third rewrites the objective function in terms of own and other firms' profits. Finally, it is useful to normalize by $\sum_{\forall s} \gamma_{fs} \beta_{fs}$, as we do in the second to last line. Implicitly, κ_{ff} is normalized to one $\forall f$, so that κ_{fg} can be interpreted as the value of a dollar of profits accruing to firm g , relative to a dollar of profits for firm f , in firm f 's maximization problem. These are the profit weights that are the object of interest in this paper.

Our notation nests a range of behavioral models. For instance, own-firm profit maximization results if $\kappa_{fg} = 0 \forall f \neq g$. A large literature in Industrial Organization treats mergers as changing $k_{fg} = k_{gf} = 0 \rightarrow 1$ (see, e.g., Bresnahan (1987); Nevo (2001)). Common ownership offers a framework for $\kappa_{fg} > 0$. This occurs when $(\gamma_{fs}, \beta_{fs}, \beta_{gs}) > 0$, in other words, when at least one investor which f pays attention to ($\gamma_{fs} > 0$) has cash-flow rights in both the firm f and the rival g .⁴

Most objections to the common ownership hypothesis can be mapped back to objections to either (A1) or (A2). However, a model of common ownership must specify the Pareto weight a firm places on each of its shareholders, sometimes called the *control weight*. Any formulation of γ is implicitly a model of corporate governance, and one where theory offers precious little guidance. Absent an obvious alternative, much of the literature assumes $\gamma_{fs} = \beta_{fs}$. This assumption is sometimes motivated by intuitive appeals to *proportional control*—the “one share one vote” rule which characterizes most publicly traded firms in the US. We caution that there is no formal link between this parameterization and any micro-founded voting game that we are aware of.

For the main derivations that follow, we will follow the literature in assuming proportional control. However, we will at times relax this assumption and allow for $\gamma_{fs} = f(\beta_{fs})$. There are two desirable properties that we would like to retain: first, that $f(\cdot)$ be monotonically increasing and continuous in holdings, and second, that $f(0) = 0$.⁵ A convenient choice is $f(\beta_{fs}) \propto (\beta_{fs})^\alpha$, which satisfies both.⁶ By varying α we can modify the convexity of the control weights, with a larger value of α leading to more weight on the largest investors.

⁴It is difficult to rationalize the conventional model of own-profit maximization in this framework, in the presence of diversified investors. Implicitly, one needs to motivate the assumption that $\gamma_{fs} = 0$ for common owners (including all investors with diversified portfolios), and $\gamma_{fs} > 0$ for entirely undiversified investors.

⁵As an example where these features may fail, consider κ in the case where $\gamma = 1$ for all shareholders of firm f , i.e. the firm maximizes their shareholders' portfolio value. This model introduces a potentially large discontinuity when a new investor with a large portfolio purchases a single share of a firm.

⁶We write \propto rather than $=$ because we can always scale the $S \times 1$ vector $\gamma_{\cdot s}$ by a scalar, and this is because it appears in both numerator and denominator of $\kappa_{fg} = \frac{\langle \gamma_{f, \beta_g} \rangle}{\langle \gamma_{f, \beta_f} \rangle} = \frac{\langle \alpha \gamma_{f, \beta_g} \rangle}{\langle \alpha \gamma_{f, \beta_f} \rangle}$.

We will show that most of our results are qualitatively insensitive to the choice of α . For example, Figure 13 shows that the trends in Figure 1 are broadly the same across different values of α , and is discussed in Section 6.

2.1 Decomposing κ

Next, we highlight an additional mathematical property of κ to set the stage for our empirical exercise. Starting from the definition of κ_{fg} in (A2), letting $\gamma_{fs} = \beta_{fs}$ (proportional control), and letting β_f denote a vectors over s , then κ_{fg} can be expressed as a ratio of inner (dot) products $\frac{\langle \beta_f, \beta_g \rangle}{\langle \beta_f, \beta_f \rangle}$. And, from the geometric definition of an inner product, $\langle x, y \rangle = \cos(x, y) \|x\| \|y\|$, with $\cos(x, y)$ the cosine distance (i.e. the cosine of the angle between vectors x and y) and $\|x\|$ the L_2 norm $\sqrt{\sum_i x_i^2}$. Substituting, we obtain a useful decomposition of κ_{fg} :

$$\kappa_{fg}(\beta) = \underbrace{\cos(\beta_f, \beta_g)}_{\text{overlapping ownership}} \cdot \underbrace{\sqrt{\frac{IHHI_g}{IHHI_f}}}_{\text{relative IHHI}}. \quad (1)$$

Here, $IHHI_f \equiv \|\beta_f\|^2$. Because β_{fs} represents the fraction of firm f owned by s , then $\|\beta_f\|^2 = \sum_{s=1}^S \beta_{fs}^2$ is the Herfindahl–Hirschman Index (HHI) for the *investors* in firm f , which we label the $IHHI_f$.

What is helpful about this expression in (1) is that it decomposes profit weights into two economically meaningful components: overlapping ownership and relative IHHI or relative investor concentration.

Overlapping Ownership: The first important term in (1) is the cosine of the angle between the positions which investors hold in f and those which investors hold in g . So long as all investors hold long positions in both (f, g) we have that $\cos(\beta_f, \beta_g) \in [0, 1]$. As the investor positions become more similar, the angle between those portfolios shrinks and $\cos(\beta_f, \beta_g) \rightarrow 1$. This suggests a link between indexing strategies, e.g. investing in the “market portfolio,” and common ownership profit weights, which explore further in our empirical exercise.

Overlapping ownership is what, in general, the literature construes to be “common owner-

ship.” It is the origin of the incentive to internalize the profits of another firm. However, as we will show, it only makes up a little over half of the empirical variation in common ownership profit weights. The remainder comes from variation in the ability of common owners to exert control, implicitly modeled as a function of investor concentration.

Relative Investor Concentration: This is the less understood source of variation in common ownership profit weights, having earned no mention in the literature so far — it ties the theory of common ownership to the notion that investor concentration drives a wedge between control rights and cash flow rights. Typically, the discussion of these two hinges on institutional structures that divorce them, e.g. “golden shares” in the hands of founders, or business groups that centralize control (Porta et al., 1999). In the objective function defined by (A2), the mechanisms are different: since the numerator of κ_{fg} depends on the product of γ and β , and both are increasing in the size of an investor’s stake, investor concentration plays a major role.

Relative IHHI has intuitive comparative statics. All other things being equal, firms with concentrated investors will place more weight on their own profits and less weight on competitor profits, because $IHHI_f$ appears in the denominator. Holding all else fixed, if firm g has fewer, larger investors then $IHHI_g$ will be large, control rights relatively expensive, and κ_{fg} smaller; if firm f has many small investors, $IHHI_f$ will be small, control rights relatively cheaper, and κ_{fg} larger. However, if a diversified investor increases its positions in both firms f and g , this may not change the ratio $\frac{IHHI_g}{IHHI_f}$.

It is entirely possible for $\sqrt{\frac{IHHI_g}{IHHI_f}}$ to be greater than one, or even greater than two or three, which makes it possible that $\kappa_{fg} > 1$ (a firm places more weight on its competitors’ profits than their own), despite the fact that the cosine similarity is never greater than one. Finally, note that since $\cos(\beta_f, \beta_g) = \cos(\beta_g, \beta_f)$, relative investor concentration is responsible for all asymmetry between profit weights κ_{fg} and κ_{gf} .

2.2 Examples of the Math of Common Ownership

The following examples maintain the proportional control assumption of $\gamma_{fs} = \beta_{fs}$.

Example 1: Consider a market with three firms. Firm 1 is privately held, in its entirety, by

Table 1: Example 1 Ownership Structure

	Firm 1	Firm 2	Firm 3
Investor 1	100%	-	-
Investor 2	-	20%	-
Investor 3	-	-	20%
Investor 4	-	20%	20%
Retail Share	-	60%	60%

Notes: This table presents investor holdings in three firms for Example 1.

an undiversified investor. Firms 2 and 3 have the following identical ownership structure: 60 percent of each is held by small, undiversified retail investors. 20 percent of each are held, respectively, by two large, undiversified investors. The final 20 percent of each is held by a single, diversified investor. This ownership pattern is summarized in Table 1.

This yields the following set of profit weights:

$$\kappa = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 1/2 \\ 0 & 1/2 & 1 \end{bmatrix}.$$

To see how this calculation is done, denote column j of Table 1 as β_j (excluding the bottom row). Then, the profit weight firm f has on firm g 's profit is $\kappa_{fg} = (\beta_f' \cdot \beta_g) / (\beta_f' \cdot \beta_f)$. This example highlights that the profit weights can be quite large with a modest amount of common ownership. An important factor here is the large retail share, which at 60% corresponds to the average retail share (i.e. non-institutional share) among S&P 500 firms in the early 1980s (see Figure 4 below).

Example 2 Now consider an alternative market with just two firms. The vast majority of both firms are held by a large set of undiversified retail investors. A boundedly small fraction of both firms is held by a finite set of N symmetric, diversified investors who each hold 1 percent of firm one and x percent of firm two, and we assume $N \cdot x < 100$. This ownership pattern is summarized in Table 2.

Then, we would have the following κ matrix of profit weights:

Table 2: Example 2 Ownership Structure

	Firm 1	Firm 2
Investor 1	1%	$x\%$
Investor 2	1%	$x\%$
\vdots	\vdots	\vdots
Investor N	1%	$x\%$
Retail Share	$(100 - N) \%$	$(100 - N \cdot x)\%$

Notes: This table presents investor holdings in two firms for Example 2. Note that $N \cdot x < 100$.

$$\kappa = \begin{bmatrix} 1 & x \\ 1/x & 1 \end{bmatrix}.$$

The calculation follows in the same manner as Example 1. This example highlights a few points about profit weights. Notice that the profit weights do not depend directly on N . Letting $x = 1$, we have that an arbitrarily small share of ownership has led to monopoly behavior. If x is 2%, then the first firm will value \$1 of the competitor’s profit as \$2 of their own. Therefore firm 1 would, if it could, divert profits directly to firm 2. This raises concerns around tunneling (Johnson et al., 2000), which we discuss in Section 5.1.

3 Data on Common Ownership

The empirical component of this paper depends on computing profit weights for S&P 500 firms for the period 1980–2017. These profit weights depend upon β , the cash-flow rights of institutional investors, which we observe as the ratio of shares held to total shares outstanding.

Our first data source for investor holdings is Thomson Reuters (TR) S34 database, which consolidates the “13(f)” filings required by the SEC for all investment managers with over \$100 million in holdings among a list of “13(f) securities.”⁷ The filings are quarterly and mandatory. These data are available to researchers through Wharton Research Data Services (WRDS) and span the period from 1980 to 2017. There are some documented data issues in the S34 database, particularly in later years.⁸ We augment this ownership data by scraping

⁷The SEC publishes a quarterly list of 13(f) securities whose holdings must be reported.

⁸Recently, WRDS and some researchers (Ben-David et al. (2018)) noticed data quality issues regarding

the data ourselves from the SEC filings. These data are available from 1999 onwards (when the SEC started requiring electronic filing), though they are much more reliable beginning in mid-2013 when the filings were required to be in XML format.⁹ We also gather data on prices and shares outstanding from the Center for Research in Security Prices (CRSP).

We use our scraped data on 13(f) holdings from 2000 onwards, and the S34 database for filings from 1980-1999. We provide additional details on dataset construction and comparisons of the two databases in Appendix B. We show that our scraped data seem to have better coverage than the Thomson Reuters database from 1999-2017 in Figure 2. Our sample of S&P 500 firms does not always include all 500 firms in each period. Because of our focus on profit weights that arise from overlapping investors, it is inappropriate to calculate these from financial holdings when there are controlling shareholders or multiple share classes. Therefore we exclude companies with controlling shareholders or special share classes with enhanced (or no) voting rights, such as Alphabet (Google) or Facebook.¹⁰ We also exclude firms where the US listing is an ADR of a stock primarily traded on a foreign exchange. The result is what we call our “restricted” sample.

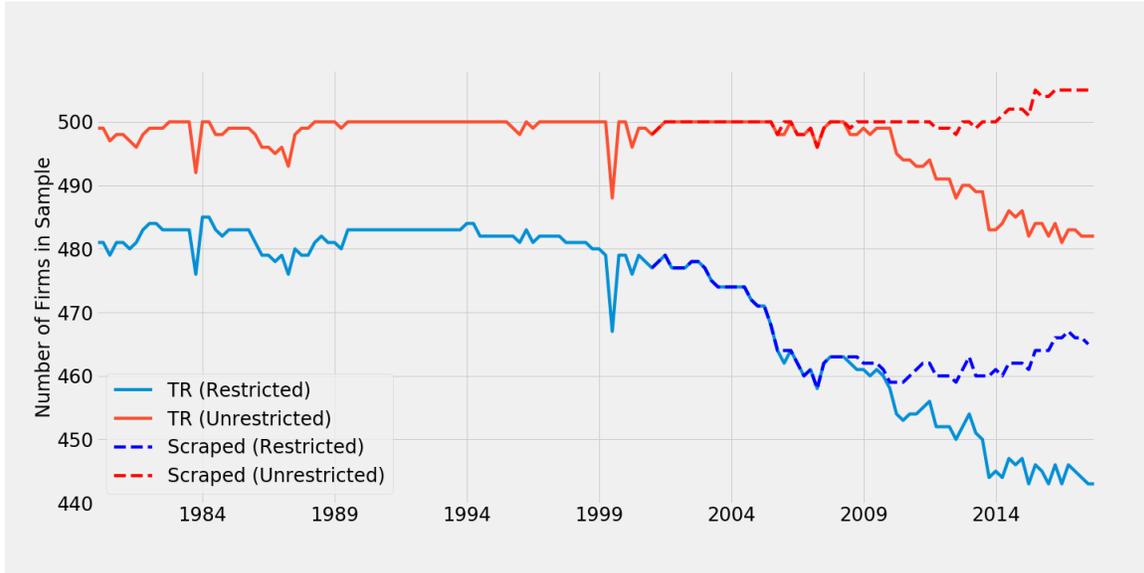
We also document the number of 13(f) managers holding S&P 500 constituents in Figure 3. The number of managers rises from around 500 in 1980 to around 4000 by 2017. In part, this rise is driven by the fact that the reporting threshold of \$100 million in 13(f) securities is nominal rather than indexed to inflation. Both the Thomson Reuters and our scraped data indicate similar numbers of 13(f) managers. We also compute the share of each firm owned by 13(f) managers and report the straight average over index constituents in Figure 4. This share has been rising from below 40% in 1980 to more than 80% by 2017, in part driven by the increasing number of 13(f) filers from Figure 3. Around 2010, the Thomson Reuters data indicates a sharp decline in the 13(f) share, while we observe no such decline

the TR dataset, and they have worked to resolve these issues. We use the July 2018 update provided by WRDS below. We consolidate all BlackRock entities. Data quality issues are discussed in more depth in Backus et al. (2019) and in Appendix B, where we document problems that remain after the 2018 update.

⁹A highly critical report from the SEC’s Inspector General in 2010 noted a number of shortcomings in how 13(f) filings were treated, prompting a number of changes to 13(f) reporting. See Securities and Inspector General (2010).

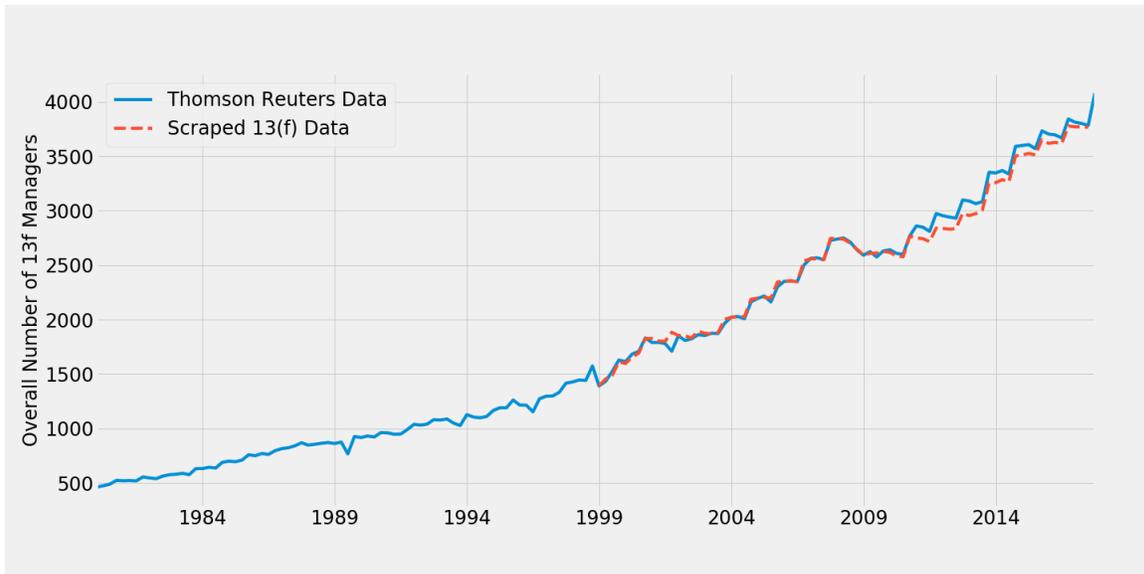
¹⁰Occasionally, these controlling shareholders are inside or retail investors, e.g. the Walton family, in violation of our theoretical assumption that retail investors are atomistic. We have excluded known examples here, however it is possible to use data from SEC Forms 4,5, 6, and 144, available from the Thomson Reuters Insider holdings database through WRDS, in order to construct industry holdings where available. These data are impractical to clean for analysis at the aggregate level, however it is feasible and important to do so for case studies of particular industries as, e.g., Azar et al. (2018) do when they compute the profit weights for airlines.

Figure 2: Number of Firms in The S&P 500 Sample



Notes: We report the Thomson Reuters in solid lines and our scraped sample in dashed lines. We report two sets of firms for each sample: (Red) an unrestricted sample consisting of all firms in the dataset (Blue) a restricted sample which drops firms with multiple share classes unlikely to satisfy control assumptions. The S&P 500 Index can contain fewer than 500 securities on a particular date (if the end of a quarter occurs on a weekend), and more recently has included over 500 securities as multiple classes of shares for the same company are included and deemed to count as one constituent (ie: BRKA and BRKB).

Figure 3: Number of 13(f) Managers holding S&P 500 Constituents

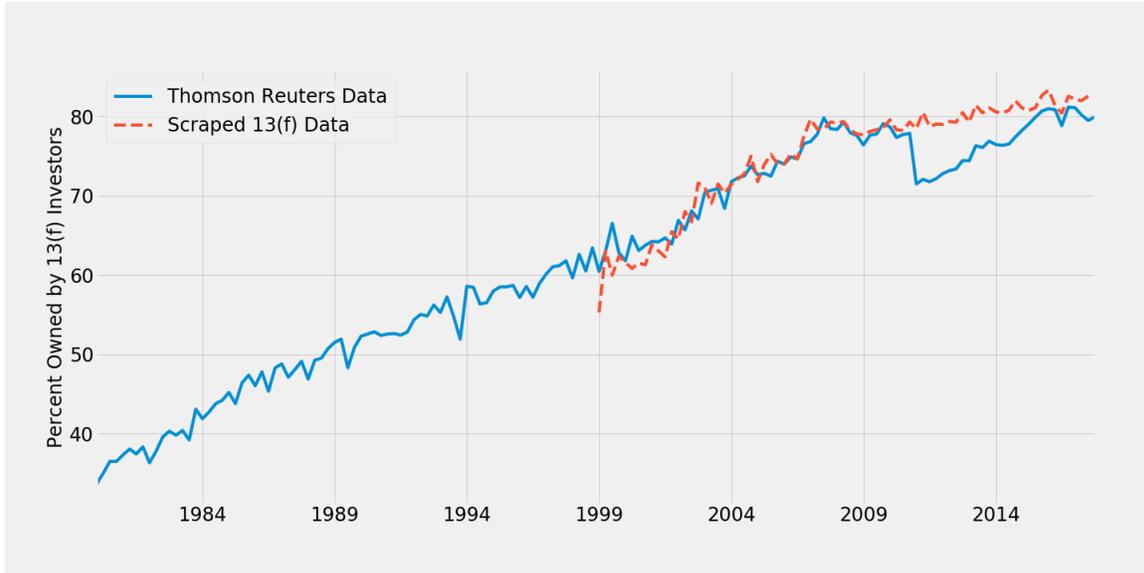


Notes: This figure depicts the number of managers filing 13(f) reports by year. For the scraped dataset, a manager is a Central Index Key (CIK). In the Thomson Reuters data, a manager is identified by a “mgrno”.

in our scraped data.¹¹

¹¹This is one of the documented issues with the S34 database see Ben-David et al. (2018).

Figure 4: Share of S&P 500 Owned by 13(f) Managers



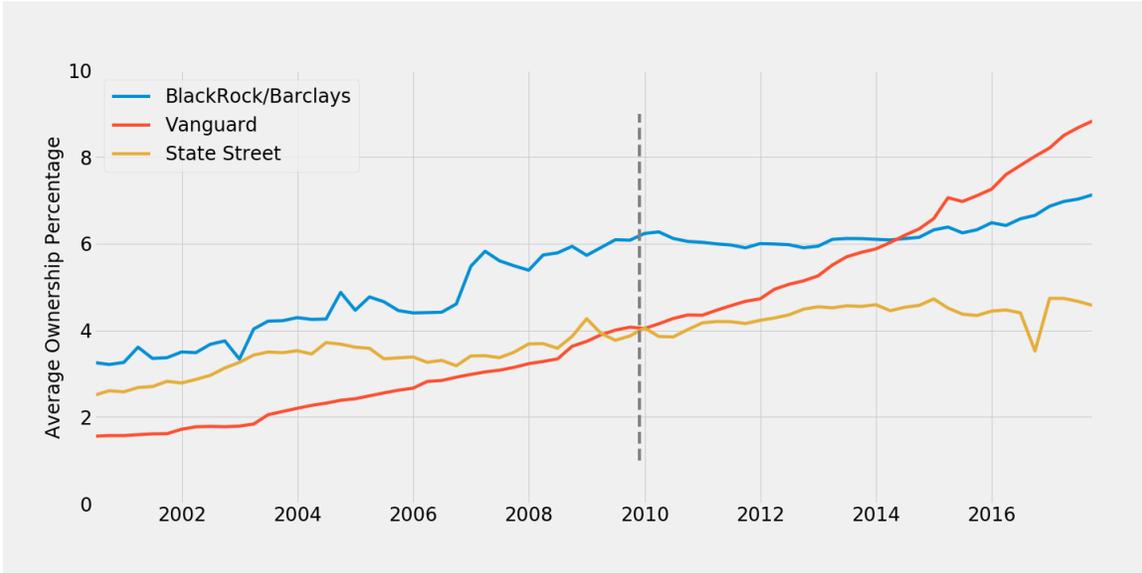
Notes: This figure depicts the average total share of a firm that is owned by managers filing form 13(f). This corresponds to the institutional ownership share of the firm, and one hundred minus this number corresponds to what we are calling the retail share. We report the straight average across index constituents rather than a weighted average.

We document a number of additional discrepancies between our scraped dataset and the Thomson Reuters S34 dataset in Appendix B. In particular, Appendix Figure 18 shows the distribution of the number of owners reported for S&P 500 constituents over time in the TR dataset, as well as our scraped and parsed sample. In TR, up to 10% of firms have fewer than 50 reported shareholders in some periods, while in our data, the numbers are more consistent over time. To further highlight this coverage issue, Appendix Figure 19 shows how much of the ownership of three particular, large firms is reported in the TR dataset versus what we find in our dataset. There is an inexplicable drop in reported ownership in the TR data, while our dataset produces a smooth series for each firm. Finally, Appendix Figure 20 shows that if one were to create Figure 1 using only the TR dataset, one would get a very different time-series, with average profit weights doubling in some time periods. Our novel dataset is available to interested researchers.

4 Trends and Patterns in Common Ownership

While there is broad agreement that common ownership is on the rise — under the premise that there is growing concentration among highly diversified institutional investors — little

Figure 5: Share of Typical Firm Owned by Big Three Institutional Owners



Notes: This figure depicts the holdings of the three large asset managers over time, combining BlackRock and Barclays. The vertical line denotes the acquisition of the Barclay’s Global Investors iShares business by BlackRock. The source data are the authors’ own scraped 13(f) dataset.

is known about the magnitude of the trend or patterns therein. Which types of firms seem most exposed to common ownership? And, what is it that drives the heterogeneity?

Discussions of common ownership are often linked to the rise in concentration among a firm’s investors, and the “Big Three” (BlackRock, Vanguard, and State Street) in particular. These three institutional investors collectively manage over \$13 Trillion at present.¹² Figure 5 highlights holdings by these “Big Three” managers. The plot shows that these firms holdings in an average S&P 500 constituent has increased substantially over time, to between 4% and 9% of a typical S&P 500 firm in 2017. Most of that rise happened after the year 2000; combined, the “Big Three” owned approximately 6% of the average firm in 2000, and 21% percent of the average S&P 500 firm by the end of 2017. While this rise is staggering, Figure 1 indicates that much of the rise in common ownership incentives predates it; indeed the average pairwise κ rose from 0.2 to 0.5 from 1980-1999, and 0.5 to 0.7 from 1999-2017. Here we turn to decomposing the variation in profit weights and their primary sources in turn. Finally, in section 4.4 below, we show that once these are accounted for, the holdings of the big three are in fact negatively correlated with common ownership profit weights.

¹²Fichtner et al. (2017) maps the historic rise of the “Big Three” and raises concerns for their role in corporate governance.

We compute common ownership profit weights (κ values) among all firms in the S&P 500 for the period 1980–2017, excluding a relatively small set of firms that use dual-class shares to separate control rights from cash-flow rights.¹³ We use the S&P 500 as it is designed to reflect the broader US economy; it consists of widely held firms, and many investment funds offer products tied to the constituent firms in one way or another.

4.1 Variation in Profit Weights

Recall from equation (1) that the profit weight κ can be mathematically decomposed into the product of two elements: overlapping ownership and relative investor concentration. Taking logs, these sources of variation are additively separable, and so we can attribute the variance due to each component:

$$\begin{aligned} Var(\log \kappa_{fg}) = & Var(\log \cos(\beta_f, \beta_g)) + Var\left(\log \sqrt{\frac{IHHI_g}{IHHI_f}}\right) \\ & + 2 \cdot Cov\left(\log \cos(\beta_f, \beta_g), \log \sqrt{\frac{IHHI_g}{IHHI_f}}\right) \end{aligned} \quad (2)$$

These are observable objects, and so the decomposition helps us to understand the sources of variation in the common ownership profit weights. Results are reported in Table 3 for the raw sample, the cross-section (residualized on quarter fixed effects), the time-series (residualized on ordered pair fixed effects), and the panel (residualized on both quarter and directional pair fixed effects).

We learn two things from Table 3. The first is that Relative investor concentration makes up a surprisingly large fraction of the variation in common ownership profit weights across all three specifications, never less than 30%. This highlights the critical role that the model of corporate governance plays in these weights. While $\cos(\beta_f, \beta_g)$ captures the overlapping ownership between firms f and g , investors' ability to use those holdings to divert profits depends on the wedge between control rights and cash flow rights, which is amplified when

¹³We exclude a total of 49 firms for using dual-class shares throughout our sample. These tend to be relatively recent entrants, which in our sample falls somewhat more steeply below 500 constituents in later years, as seen in Figure 2.

Table 3: Decomposition of Variance of $\log \kappa$

	Overlapping Ownership	Relative IHHI	Covariance
Raw	68.67%	31.32%	.01%
Cross-Section	54.88%	45.11%	.01%
Time-Series	61.04%	30.71%	8.24%
Panel	53.71%	36.97%	9.31%

This table describes the attribution of variance according to the decomposition in (2). The raw sample is unmodified; the cross-section is residualized on quarter FE; the time-series is residualized on firm pair FE, and finally the panel case is residualized on both quarter and pair FE. Note that pairs are directional.

the firm’s investor holdings are relatively unconcentrated. In other words, the most severe distortions of corporate conduct, according to the common ownership hypothesis, come about when there is overlapping ownership as well as relatively less investor concentration in one of the two firms, allowing investors in that firm relatively cheaper control rights.

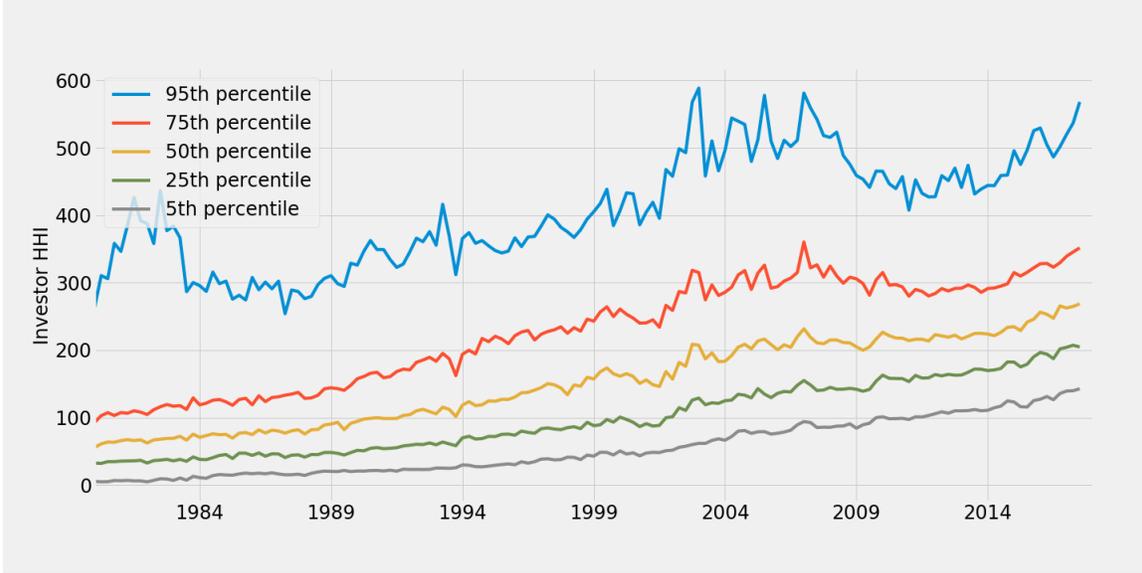
Second, we learn that the role of relative investor concentration is greatest in the cross-section, over 45% of the variation. This is easily reconciled — although investor concentration is on the rise in the time series as retail share shrinks (as evident in Figure 4) what appears in the numerator for κ_{fg} appears in the denominator of κ_{gf} . Therefore it is the increase in overlapping ownership, driven by indexing behavior, which explains the lions share of the rise of common ownership in the time series.

4.2 Relative Investor Concentration

Given the role of relative investor concentration, we next consider the question: how concentrated are the set of investors in a typical S&P 500 constituent? We can calculate the investor HHI: $IHHI_f = \sum_s \beta_{fs}^2$ and interpret this measure in terms of equivalent symmetric investors as $\frac{1}{IHHI_f}$. We report the quantiles of investor concentration (multiplied by 10,000 as is common practice) in Figure 6. What we see is that investor concentration has grown dramatically since 1980. In 1980, the median firm’s investor concentration was around 50 points (or approximately 200 symmetric investors), and today it has an $IHHI \approx 250$, or around 40 symmetric investors. For the most concentrated firms (95th percentile of investor concentration), the $IHHI \approx 500$, which would represent around 20 equally-sized investors.¹⁴

¹⁴Note that by antitrust standards, investors are not very concentrated at all. For example, the DOJ and FTC consider product markets to be highly concentrated only when $HHI > 2500$, and consider markets to be moderately concentrated when $HHI \in [1500, 2500]$. We caution that there is no reason to think antitrust

Figure 6: Investor β Concentration (IHHI)



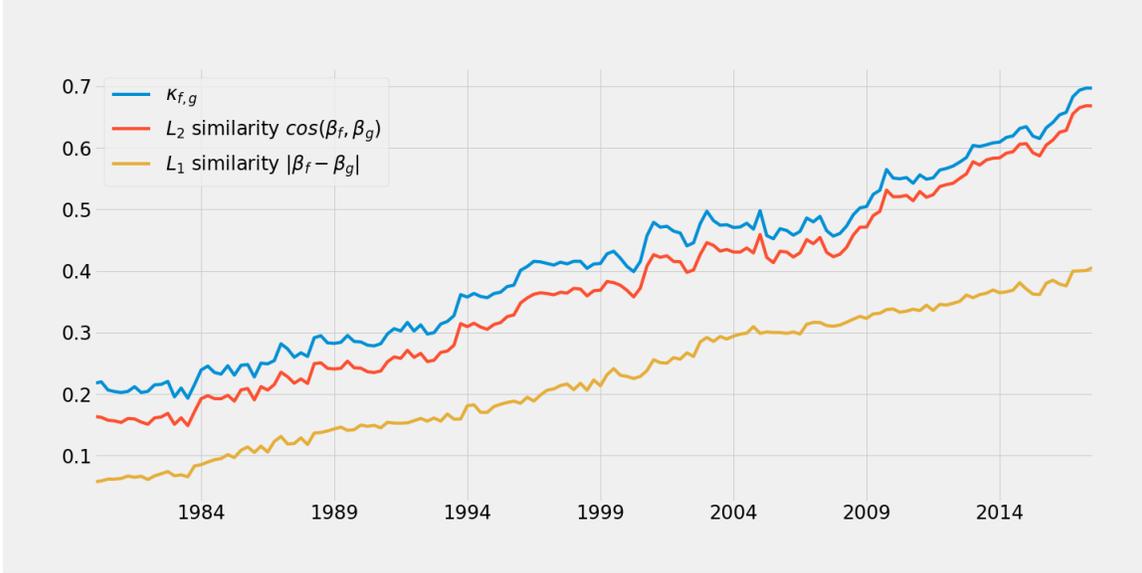
Notes: This figure plots quantiles of the firm distribution of IHHI, which is given by $\sum_s \beta_{fs}^2$. As is common in antitrust we report $10,000 \times IHHI_f$.

What has driven the rise in $IHHI$ over time? Note that $\sum_s \beta_{fs}$ is not guaranteed to be one; rather, it sums to the institutional investor share, or one minus the retail share (defined here to be the fraction of shares held by investors who do not file a 13(f) form). Therefore $IHHI_f$ is inversely related to r_f , the retail share of firm f . Also recall that the typical retail share (Figure 4) has fallen from around 60% in 1980 to around 20% today. Thus part of this trend is about 13(f) filers taking larger positions, such as the rise of the “Big Three”, while part is driven by the rise in 13(f) filers overall.

However, the theoretical relationship between investor concentration and profit weights is not straightforward. Recall equation (1) which showed that $\kappa_{fg} = \cos(\beta_f, \beta_g) \sqrt{\frac{IHHI_g}{IHHI_f}}$, or that profit weights depend on relative investor concentration. Holding all else equal, as firm f 's own investors become more concentrated we expect them to put *less* weight on other firms' profits. But a general rise in IHHI will appear in both the numerator and the denominator, so the effect is ambiguous. So, though $IHHI$ has been rising since 1980, relative investor concentration cannot be rising for all pairs of firms simultaneously, and therefore rising investor concentration this cannot fully explain the rise over time in κ . Rather, as Table 3 reflects, its role is largest in the cross-section.

guidelines for product markets are appropriate to apply to investors.

Figure 7: Cosine Similarity Among Investors



Notes: We report average similarity measures across all pairs of firms in each period. Note that similarity here is across firms (the vector β_f) rather than investors (the vector β_s), where the latter appears in Figure 8 below.

4.3 Overlapping Ownership and Indexing

Besides relative investor concentration, overlapping ownership, or the cosine similarity of vectors β_f and β_g , is the other element determining profit weights in (1). Cosine similarity is an L_2 measure, and measures how similar the investors' positions in firm f are to those in firm g . For long-only portfolios it ranges from $[0, 1]$ and is maximized when the vector of investor shares in firm f can be expressed as a scalar multiple of the investor positions in firm g . This can arise if all of the investors agree on all of the portfolio weights for their investments but have differently sized portfolios.¹⁵ To be explicit we can write:

$$L_2(\beta_f, \beta_g) = \cos(\beta_f, \beta_g) = \frac{\sum_s \beta_{fs} \beta_{gs}}{\|\beta_f\| \|\beta_g\|}.$$

One potential criticism of L_2 measures of similarity is that they put additional weight on the largest investors, and may therefore conflate investor similarity and investor concentration. To better get at investor similarity directly, we can construct an L_1 measure. The core of this measure is $1 - \sum_s |\beta_{fs} - \beta_{gs}|$. It is largest when all investors hold the same fraction of both firms (f, g) so that $\beta_{fs} = \beta_{gs}$. Assuming no short positions are allowed, it is largest

¹⁵As an example: assume that all investors have different sizes to their overall portfolio but allocate a portfolio share of β_{fs} to firm f and β_{gs} to firm g . If we can write $\frac{\beta_{fs}}{\beta_{gs}} = a$ for all investors s then $\cos(\beta_f, \beta_g) = 1$

when investors hold either a position in firm f or in firm g , and thus are not common owners. We construct a L_1 measure of similarity which varies from $[0, 1]$:¹⁶

$$L_1(\beta_f, \beta_g) = \frac{1}{2} \sum_s (\beta_{fs} + \beta_{gs} - |\beta_{fs} - \beta_{gs}|). \quad (3)$$

This is not our preferred measure, as it does not correspond to a profit weight of an objective function, but it may help us quantify the extent to which firms (f, g) have owners in common. In Figure 7 we depict this relationship; we find that the average (across pairs of firms) cosine similarity almost perfectly tracks the average profit weight κ . We also see that the L_1 measure of overlapping investors is also increasing though it doesn't line up as directly with the profit weights.

Both of our L_1 and L_2 measures focus on pairs of firms, and tell us that positions held in firm f look more similar to those in firm g over time. Perhaps the most important phenomenon from 1980–2017 is the rise of index investors. Instead of looking at pairs of firms, we might want to focus the extent to which investors pursue indexed strategies. For each period we can construct a set of $\bar{w}_f = \frac{\sum_s \beta_{fs}}{\sum_{f,s} \beta_{f,s}}$'s which represent the market portfolio.¹⁷ We can then compare the normalized portfolio weights $w_{fs} = \frac{\beta_{fs}}{\sum_f \beta_{fs}}$ and measure the distance each investor's portfolio is to the market portfolio: $L_1(w_s, \bar{w})$ and $L_2(w_s, \bar{w})$. This is consistent with the literature in that the *active share* measure of Cremers and Petajisto (2009) is given by $1 - L_1(w_s, \bar{w})$.¹⁸

Our goal is to quantify how indexed each investor is on a scale of $[0, 1]$, with 1 being perfectly indexed. We compute the similarity between an investor's portfolio w_s and our constructed "market portfolio" \bar{w} . In Figure 8, we report the weighted average of these similarity measures, where we weight each investor by assets under management (AUM).¹⁹ As one might expect, at least on an asset weighted basis, investor portfolios become much more similar to

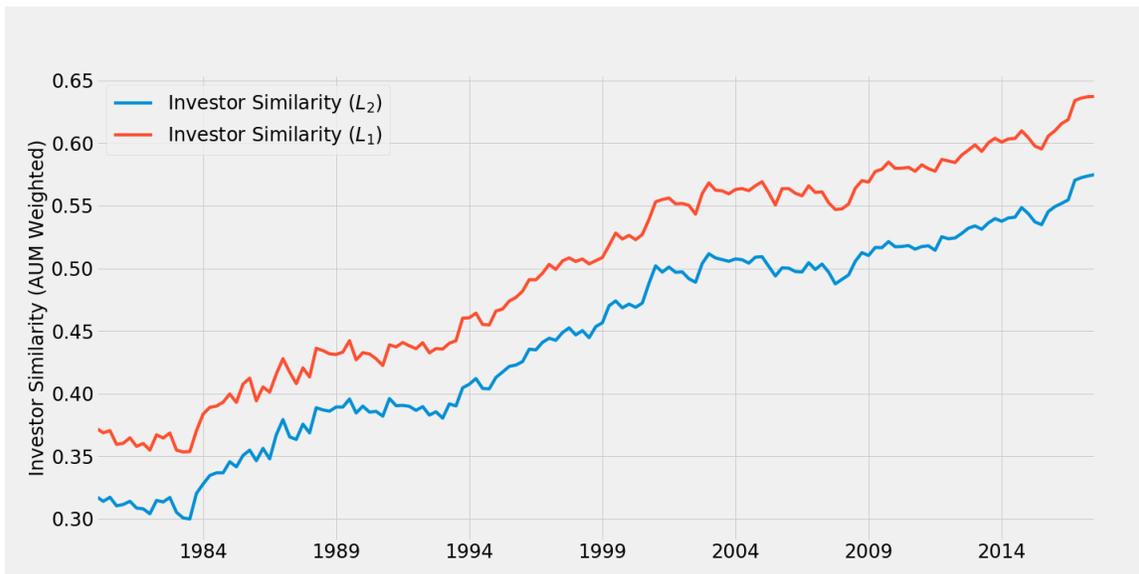
¹⁶Absent retail investors $\sum_{\forall s} \beta_{fs} = 1$. In practice, $\sum_{\forall s} \beta_{fs} < 1$, because the set of investors contains only large institutional investors who provide 13(f) filings to the SEC. We can think about $\sum_{\forall s} \beta_{fs} = 1 - r_f$ where r_f represents the retail investor share in firm f . As r_f grows, the L_1 measure declines, which may (or may not) be the desired behavior.

¹⁷Our measure of the "market portfolio" is based on cashflow shares rather than market-cap weights. But for the "retail share" of non 13(f) filers, these two measures would coincide. One interpretation of our measure is as the "market portfolio" weights among large institutional investors only. We obtained S&P weights for the most recent period and our "market portfolio" weights were highly similar. Note: we ignore all non S&P 500 securities from our calculation of portfolio weights.

¹⁸We should point out that our analysis is at the investor/manager level from 13(f) filings not at the level of an individual fund.

¹⁹Again restricted to the set of S&P 500 securities.

Figure 8: Similarity Between Investor Portfolios and S&P 500 Index



Notes: This figure depicts L1 and L2 similarity measures comparing investor portfolios weighted by investor AUM within our sample of S&P 500 assets.

the “market portfolio”.

Taken together, these facts are meant to highlight what we think are the two main trends driving long run changes in common ownership profit weights: (1) the positions of investors in firms (f, g) become more similar to each other over time and (2) the similarity is largely driven by a broad trend towards indexing among asset managers. This contrasts what appears to be the developing narrative that common ownership is largely a function of rising investor concentration particularly among the “Big Three.”

4.4 Correlates of Profit Weights

Next we ran a series of regressions of κ_{fgt} on potential covariates. In each, we include quarter and pairwise FE, where the pairs are ordered (i.e., a different FE for the time series κ_{fg} and κ_{gf}). Results are presented in Table 4.

Across all specifications we obtain a strong positive relationship between κ and the retail share. This is consistent with the theory. Recall that $IHHI_f = \sum \beta_{fs}^2$, but $\sum \beta_{fs} < 1$, where $1 - \sum \beta_{fs}$ is taken to be the retail share. Therefore the retail share is negatively

correlated with $IHHI_f$, and so mechanically positively correlated with $\kappa_{fg} \forall g$.

The log of the market cap of firm f is also consistently positively correlated with the common ownership profit weight. This reflects the inclusion of larger firms in indexes, and the corresponding increase in overlapping ownership.

In model (1) we include the sum of β_{fs} for shareholders Blackrock, Vanguard, and State Street, and find a strong positive correlation. We also find, in model (2), a strong relationship between our firm-level measure of “indexing” behavior and common ownership profit weights. For each percentage point that each of a firm’s 13(f) investors become more similar to the index, we expect an equivalent rise in the profit weights that firm f places on all other firms g . The 10th percentile of our firm-level investor L_2 similarity measure is ≈ 0.373 and the 90th percentile is ≈ 0.625 . Thus an increase from the “least indexed” firms in the sample to the “most indexed” would increase the average weight f places on other firms by ≈ 0.26 units. Much of the difference in our “indexing” measures has taken place over time, looking just at the last quarter of 2017, the range between the top and bottom deciles of “indexing” is between 0.56 and 0.70.

However, in model (3), when we include both the holdings of the Big Three as well as our indexing measure, the coefficient on the former turns negative. Likewise, in model (4), when we disaggregate the Big Three and include the individual holding, each of the correlations is negative. We take this as clear evidence that indexing, not the rise of the Big Three — or any individual institutional investor — explains the broader trend in the rise of common ownership.

5 Economic Implications

5.1 Relationship to Tunneling

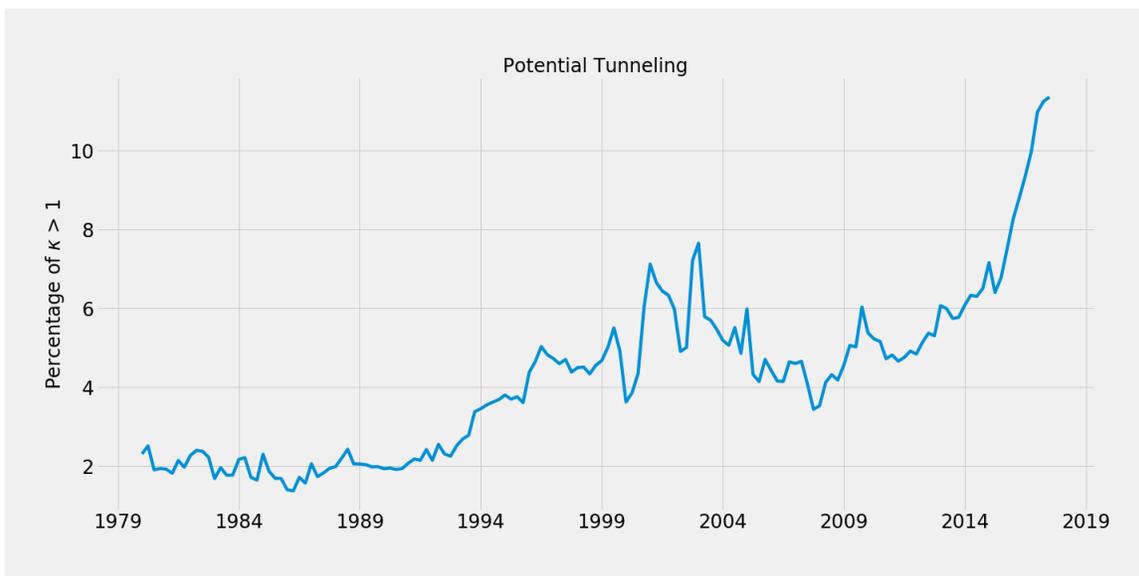
Following the language of Johnson et al. (2000), *tunneling* is the practice of transferring profits, whether via acquisition, mispriced purchase orders, or direct transfer, from one company to another in order to benefit the interests of a controlling stakeholder in both. This expropriates both creditors and minority shareholders in the former firm. The above-referenced paper offers anecdotal evidence of tunneling even in developed countries, particularly civil

Table 4: Correlations with κ

	(1)	(2)	(3)	(4)
Retail Share	0.8581*	0.7188*	0.6300*	0.6307*
	(0.0005)	(0.0005)	(0.0006)	(0.0006)
Log(Market Cap)	0.0356*	0.0335*	0.0293*	0.0290*
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Big Three Holdings	0.1473*		-0.9138*	
	(0.0019)		(0.0021)	
Investor Indexing		1.0098*	1.2850*	1.2672*
		(0.0011)	(0.0013)	(0.0013)
Blackrock Holdings				-0.3433*
				(0.0033)
Vanguard Holdings				-1.3216*
				(0.0051)
State Street Holdings				-1.3478*
				(0.0034)
R^2	0.5358	0.5598	0.5653	0.5668
Quarter FE	✓	✓	✓	✓
Pair FE	✓	✓	✓	✓
N	14852026	14852026	14852026	14852026

Notes: This table reports correlates of the common ownership profit weights. An observation is a pair of S&P 500 constituent firms in a given quarter. Robust standard errors are clustered at the firm level and reported in parentheses. * indicates significance at the 5% level.

Figure 9: Potential Tunneling Incentives $\kappa > 1$



Note: This reports the fraction of pairwise profit weights $\kappa_{f,g} > 1$ in each period under the assumption of proportional control.

law countries, and other work has found evidence in the developing world (Bertrand et al., 2002). However, tunneling is not typically believed to occur in the US for two reasons: strong investor protections that facilitate healthy financial markets (Porta et al., 1999) and the near-universal absence of a controlling interest in publicly-traded firms, as the US is the land of the “widely-held” firm (Berle and Means, 1932).

The connection between common ownership and tunneling hinges on this second point. If as the common ownership hypothesis maintains: (1) owners are sufficiently diversified and (2) firms care about the effects of their decisions the entirety of their shareholders portfolios; then firms may have an incentive to engage in tunneling even in the absence of a controlling interest. On this point we can be precise: if $\kappa_{fg} > 1$ then firm f would, if it could, transfer profits directly to firm g .

In Figure 9, we report the share of firm pairs for which $\kappa_{fg} > 1$ under the proportional control assumption. Recall that, from equation (1), since $\cos(\beta_f, \beta_g)$ is bounded above by 1, $\kappa_{fg} > 1$ implies that $\kappa_{gf} < 1$ — i.e. that tunneling is in the interest of both firms. Because tunneling is necessarily unidirectional, the maximum amount of tunneling relationships would be 50%. Therefore, twice the number described in the figure yields the fraction of pairwise relationships among S&P 500 firms in which parties have an incentive to engage in tunneling. We find a striking rise in this frequency between 1993 and 2002, and again in

the period following 2015.

There is a meaningful difference between the patterns of tunneling predicted by common ownership and the prior literature. In the latter, tunneling tended to be isolated within small groups of firms that had a common controlling interest. For example, Bertrand et al. (2002) offers econometric evidence of tunneling about documented business groups in India. Therefore, the pattern of tunneling interest is *sparse*—firms possess few tunneling “targets”. In contrast, tunneling arising from common ownership is driven by patterns of retail share via $IHHI_f$. When retail share is large, $\sqrt{\frac{IHHI_g}{IHHI_f}}$ grows for *all* potential tunneling “targets.” This suggests that the resulting patterns of tunneling will tend to be dense rather than sparse—firms which have incentives to engage in tunneling may want to tunnel funds to many partners.

Taken at face value, this finding implies that in the world of the widely-held firm, i.e. in the absence of a controlling interest, the incentives for tunneling may be pervasive if firm incentives reflect common ownership concerns. It is worth emphasizing that, unlike our results in Section 4, in the later periods, the result depends heavily on our assumptions about γ . We document this in Appendix Section C.1, where putting more weight on large investors actually results in higher tunneling incentives.

The feasibility of tunneling in the world of the widely-held firm is a novel result. However, it is anticipated in Matvos and Ostrovsky (2008), who observe that institutional shareholders with cross-holdings in acquiring and target firms tend to vote in support of the merger, sometimes to the detriment of share value for the acquirer. They offer the clearest systematic documentation of tunneling arising from common ownership. Ultimately though, the implication of taking the common ownership hypothesis seriously, that in 2017 more than 10% of the S&P 500 is engaging some form of tunneling behavior, is implausibly strong. It is possible that this is held in check by strong minority shareholder protections, but perhaps more likely, these incentives may be incompletely transmitted from owner to institutional manager, and from institutional manager down the corporate chain to actors in the firm.

5.2 Quantifying the Common Ownership Channel

Eekhout et al. (2018) document that average markups rise from 21% in 1980 to 61% in 2017 across a broad range of publicly traded firms. We conduct a simple calibration exercise in

order to compare both the magnitude and the timing of the price effects implied by the common ownership hypothesis.

We start with J symmetric firms, with marginal costs c , selling differentiated products and competing in Nash-in-prices. We assume that each firm faces a logit demand such that its market share is given by:

$$s_j(p_j, p_{-j}) = \frac{e^{a-bp_j}}{1 + \sum_{k=1}^J e^{a-bp_k}}.$$

Each firm chooses its p_j simultaneously in order to maximize:

$$\tilde{\pi}(p_j, p_{-j}, \kappa) = (p_j - c)s_j(p_j, p_{-j}) + \sum_{k=1}^J \kappa_{jk} \cdot (p_k - c)s_{k'}(p_j, p_{-j}).$$

Given the parameters of the problem (a, b, c, J, κ) it is possible to solve the $J \times J$ system of equations for the equilibrium prices $\hat{p}(\kappa)$. Our goal is to hold fixed the (a, b, c, J) aspect of the problem, and to re-solve the problem with all κ_{fg} set equal to the average value reported in Figure 1 period by period. We then plot $\mu = p/c$ as Eekhout et al. (2018) does over time from 1980 to 2017.

We calibrate parameters as follows. First we set $c = 1$ without loss of generality. This means that prices and markups are one in the same: $\hat{p}(\kappa) = \mu$. Next we choose the number of firms $J = 8$ so that our $HHI \approx 1250$ to match Grullon et al. (2018).²⁰ Finally, we calibrate a and b for 1980. We construct a markup of $\mu = 1.21$ to match Eekhout et al. (2018) and an average own-elasticity of -7.21 in line with the range of elasticities reported in Eaton and Kortum (2002).²¹ This all but eliminates the outside good share.²²

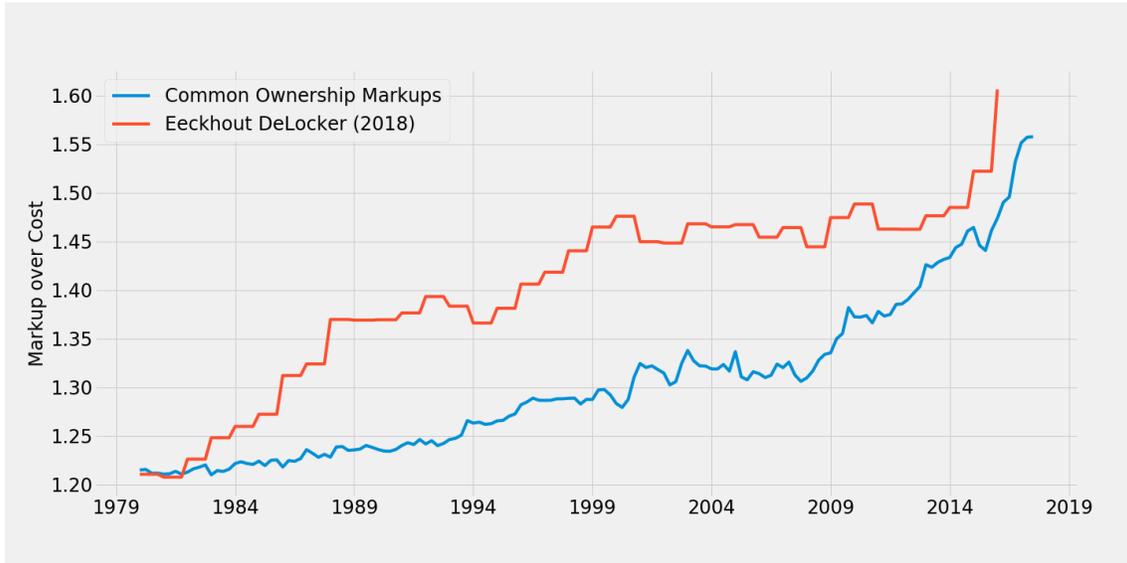
Results for this calibration exercise are presented in Figure 10. The scale of the increase in markups predicted by the rise in common ownership is substantial: from 1.21 to 1.56. This is very similar in magnitude to the rise in markups found by Eekhout et al. (2018) for the

²⁰We can obtain nearly identical results varying the number of firms from 5-15.

²¹Simonovska and Waugh (2014) obtain elasticities about half as large ≈ -4.0 which suggests that demand is too inelastic to get markups as small as $\mu = 1.21$ in 1980.

²²Alternatively, one could eliminate the parameter a as well as the outside good, but the existence of even a very small outside good option substantially improves convergence of the simulated prices when computing equilibrium. This computation is done with the freely available `pyblp` python package (?).

Figure 10: Simulated Markups: 1980–2017



Notes: This figure presents predicted markups for the calibration exercise. See the text for exact specification.

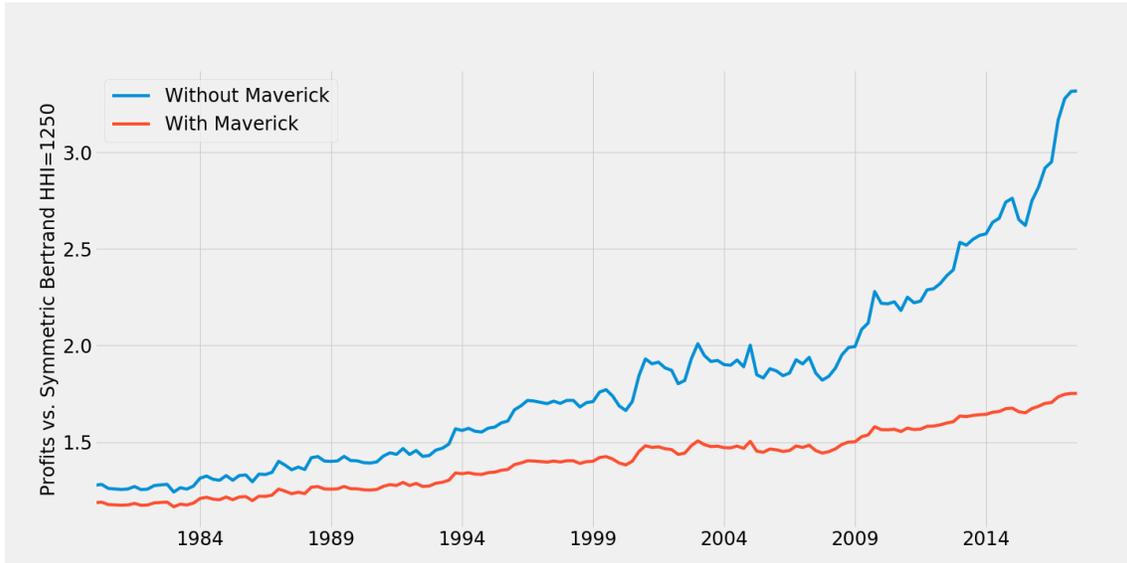
same period.²³ Moreover, a Granger test rejects the null that the lagged simulated common ownership markups are not predictive of the Eekhout et al. (2018) markups.²⁴ Despite this, we take the time series in Figure 10 as clear evidence that the bulk of the rise in markups described in Eekhout et al. (2018) is inconsistent with the predicted price effects of the common ownership hypothesis in our toy example. In particular, we note the timing, which is largely insensitive to the specification of the example: the former observe a substantial increase in markups in the 1980s, whereas the increase in markups predicted by common ownership follows largely after 2010. However, we do observe some coincidence in the two time series, particularly in 2009 and following 2015.

Whether or not one believes that common ownership price effects are manifested, we are also interested to know by how much profits would be greater if they were. This speaks directly to the incentives of institutional managers, who tax portfolio value uniformly, or the incentives of the ultimate owners to delegate control to institutional managers who will exercise their

²³An important aspect of the Eekhout et al. (2018) results that this exercise misses is the reallocation from low- to high-markup firms, since in our exercise all firms are symmetric. Alternatively one could match κ_{fg} to firm-level markups, but making sense of that relationship would require a pricing game (which firms compete with which and how, a particularly difficult question at this bird’s-eye level), a problem we elude with our logit pricing example.

²⁴The test (Granger, 1969) is based on a VAR in first differences with two lags and a time trend, and rejects with a critical p value of 0.05.

Figure 11: Simulated Markups: 1980–2017



Notes: This figure presents predicted profits for the calibration exercise. The blue line represents the baseline model. The red line, i.e. With Maverick, depicts a modification of the model in which the eighth firm is held by an entirely undiversified owner.

corporate governance rights in a fashion consistent with the common ownership hypothesis. Therefore we depict profits associated with the pricing equilibrium in the blue line in Figure 11.

We find a dramatic, more than threefold increase in profits associated with the rise in simulated markups for our calibrated example. However, this result is sensitive to the symmetry of the profit weights. If even one firm in the market prices aggressively, then the resulting markups (and profits) are much lower. To demonstrate this, the red line in Figure 11 depicts profits under a modification of the example in which the eighth firm is a “Maverick,” i.e. they price as if held by an entirely undiversified owner. The change in profits is now substantially lower, an approximately 70% increase instead of a more than threefold increase. Nonetheless, the magnitude of these numbers emphasizes the fact that if owners could successfully incentivize institutional managers and firms to behave in a manner consistent with common ownership pricing incentives, they may stand to gain substantially.

We learn two more things from the Maverick exercise, however, which we believe are important for the literature on testing common ownership moving forward. First, that within-market dispersion in common ownership profit weights can generate dramatic variation in the predictions of the model. That dispersion can even appear between pairs of firms: from

(1), $\kappa_{fg} \simeq \kappa_{gf}$ only when $IHHI_f$ and $IHHI_g$ are close. These latter expressions, however, are sensitive to variation in retail share. Within-market dispersion and between-firm asymmetries are obscured when the econometrician aggregates common ownership profit weights up to the market level using measures such as MHHI. In doing so, we believe they throw away some of the most interesting variation and its testable implications. This is a discussion we continue in Backus et al. (2018). Second, however, it suggests that the presence of privately-held firms is not merely a data nuisance, but has testable and useful implications for the manifestation of common ownership price effects.

5.3 The “Big Three”: Mergers and Breakups

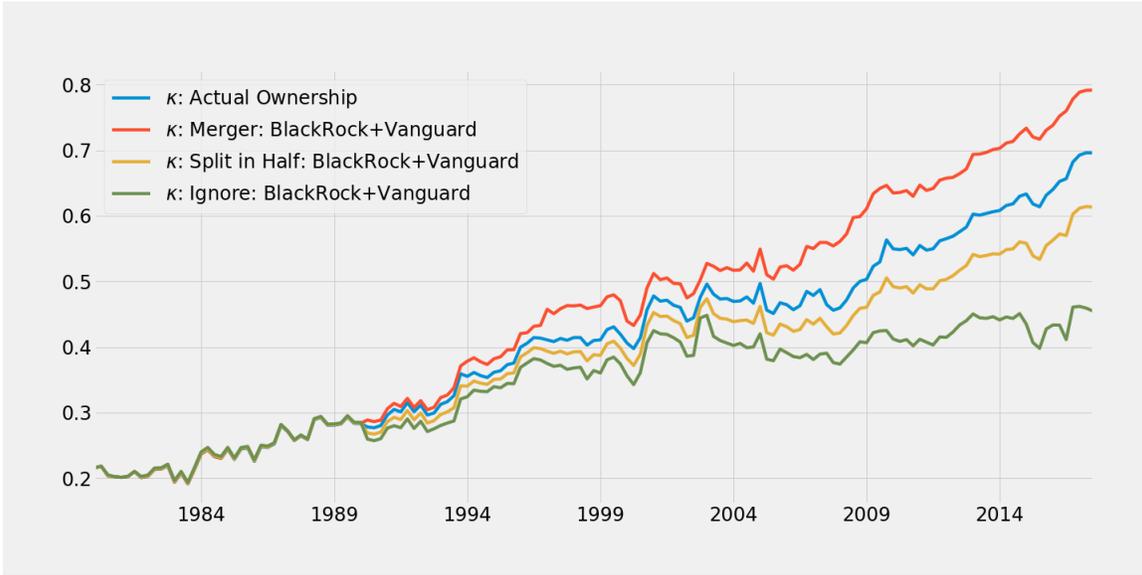
There has been much discussion of the role played by the “Big Three” investment management firms (BlackRock, Vanguard, and State Street) with respect to common ownership incentives, including various proposals to restrict the size of large institutional investors in different ways Posner et al. (2017). Here we consider a simple exercise where we either: (a) allow BlackRock and Vanguard to merge; (b) we take BlackRock and Vanguard and split them each into two firms (BlackRock A/B, Vanguard A/B) with identical holdings that are half as large as the current firm;²⁵ or (c) we tell firms to “ignore” BlackRock and Vanguard by setting $\gamma_{f,s} = 0$, which implicitly treats them as “retail” investors.

We report our findings in Figure 12. Up through 2004 there are limited effects on κ values of either allowing BlackRock and Vanguard to merge, or breaking them up into identically sized smaller firms. By the end of the sample, there begin to be more substantial differences. Under our baseline scenario of proportional control and the observed ownership structure $\bar{\kappa} \approx 0.7$, the merger would increase this to $\bar{\kappa} \approx 0.8$, while breaking them up would decrease this to $\bar{\kappa} \approx 0.62$. Qualitatively, the trend over time is similar to our baseline case. The most drastic difference comes when we “ignore” BlackRock and Vanguard by setting $\gamma_{f,s} = 0$. This gives $\bar{\kappa} \approx 0.46$ in 2017, and it implies that average profit weights are essentially unchanged since 2000.

First, we change the ownership structure of the two largest firms without changing the degree to which investors are “indexed” because we either merge them or split them into smaller

²⁵We do not split holdings based on overlapping industries (one of the suggestions in Posner et al. (2017)) but rather simply increase or decrease the overall size of BlackRock and Vanguard. This shouldn’t matter because we are reporting the average profit weight $\bar{\kappa}$ for the entire S&P 500 index.

Figure 12: Alternative Ownership Structures for BlackRock and Vanguard

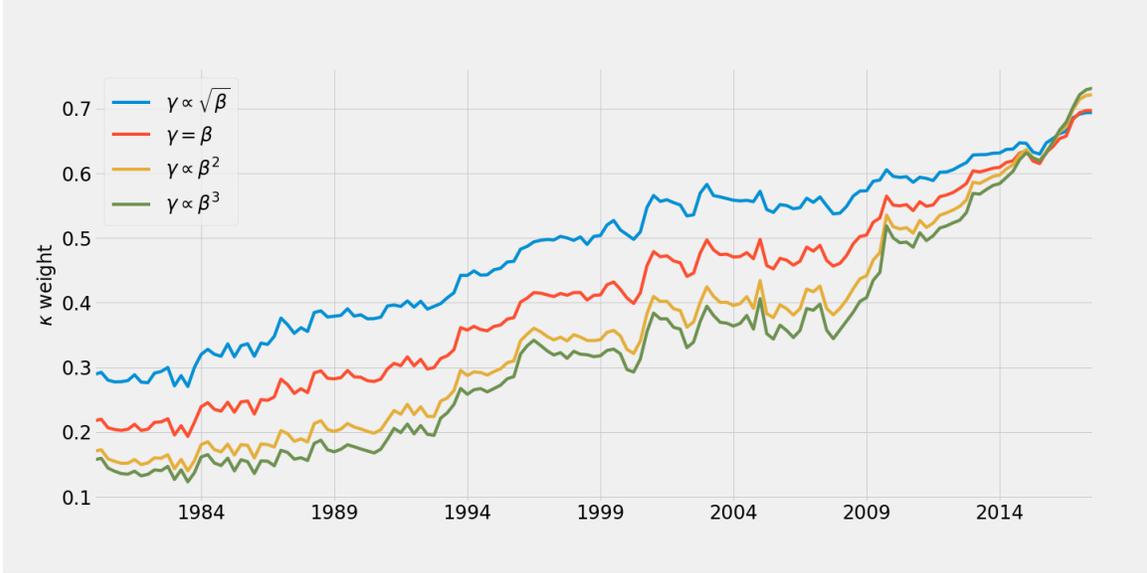


Note: Actual ownership uses proportional control assumption from Figure 1. Splitting firms in half preserves portfolio weights but constructs two identical firms each half as large as the original. Merger combines BlackRock and Vanguard into single firm with combined holdings. Ignoring BlackRock and Vanguard sets $\gamma_{f,s} = 0$ for those investors and implicitly includes them in retail share.

firms with identical holdings. This tells us two things. While large firms like BlackRock and Vanguard play a role in the rise in common ownership incentives, they play a smaller role (controlling for “indexing”) than one might think, because splitting them in half reduces $\bar{\kappa}$ by only ≈ 0.08 units. Likewise the combined BlackRock and Vanguard firm would be enormous (owning more than 15% of most S&P constituents). Under proportional control this increases the average profit weights, albeit not dramatically. Taken together, this highlights that indexing behavior, rather than the growth of the largest investment managers, seems to be driving the long-run trends in profit weights.

When we “ignore” BlackRock and Vanguard by setting $\gamma_{f,s} = 0$ for those two investors, we are implicitly treating them as if they are retail investors. This both drastically reduces the degree of “indexing” in the market by concentrating control in the remaining institutional investors who tend to be less “indexed” than BlackRock and Vanguard. We explore this in Appendix C.2, where Appendix Figure 24 shows the impact of removing those two firms from our measures of indexing developed in Section 4.3. More disagreement among the remaining investors tends to lead to lower profit weights overall. We can think of this scenario as similar to the “put the shares in a drawer” proposal of Posner et al. (2017), where institutional investors above a certain size would agree not to participate in corporate governance activities. As several have pointed out, while this remedy may be effective at

Figure 13: Profit Weights Under Different Control Assumptions



Notes: This figure reports $\bar{\kappa} = \frac{1}{F(F-1)} \sum_f \sum_{g \neq f} \kappa_{f,g}$ under different maintained assumptions of control weights, with $\gamma \propto \beta^\alpha$.

curbing common ownership incentives, this proposal might have unintended consequences in reducing the effectiveness of other corporate governance actions.

6 Robustness

6.1 Profit Weights and Control

In Figure 1, we saw that under the assumption of proportional control, $\gamma = \beta$, there is a stark positive trend in common ownership incentives (κ) among S&P 500 firms, growing from an average of 0.2 to 0.7 between 1980 and 2018. Figure 13 plots the average κ for every pair of S&P 500 firms by quarter for different control assumptions. We set $\gamma_{fs} \propto \beta^\alpha$ and vary the α parameter. As we increase the exponent α , we concentrate more control among the largest investors in firm f . We see that the increasing trend is relatively robust to assumptions about corporate control, and that toward the end of the sample (2012-2017), the average κ profit weight does not appear to depend on our choice of γ .

Perhaps contrary to expectations, as we increase α , the average weight κ that a firm places on its competitors' profits *decreases*. Toward the very end of the sample this relationship

inverts, though differences among average profit weights become negligible.

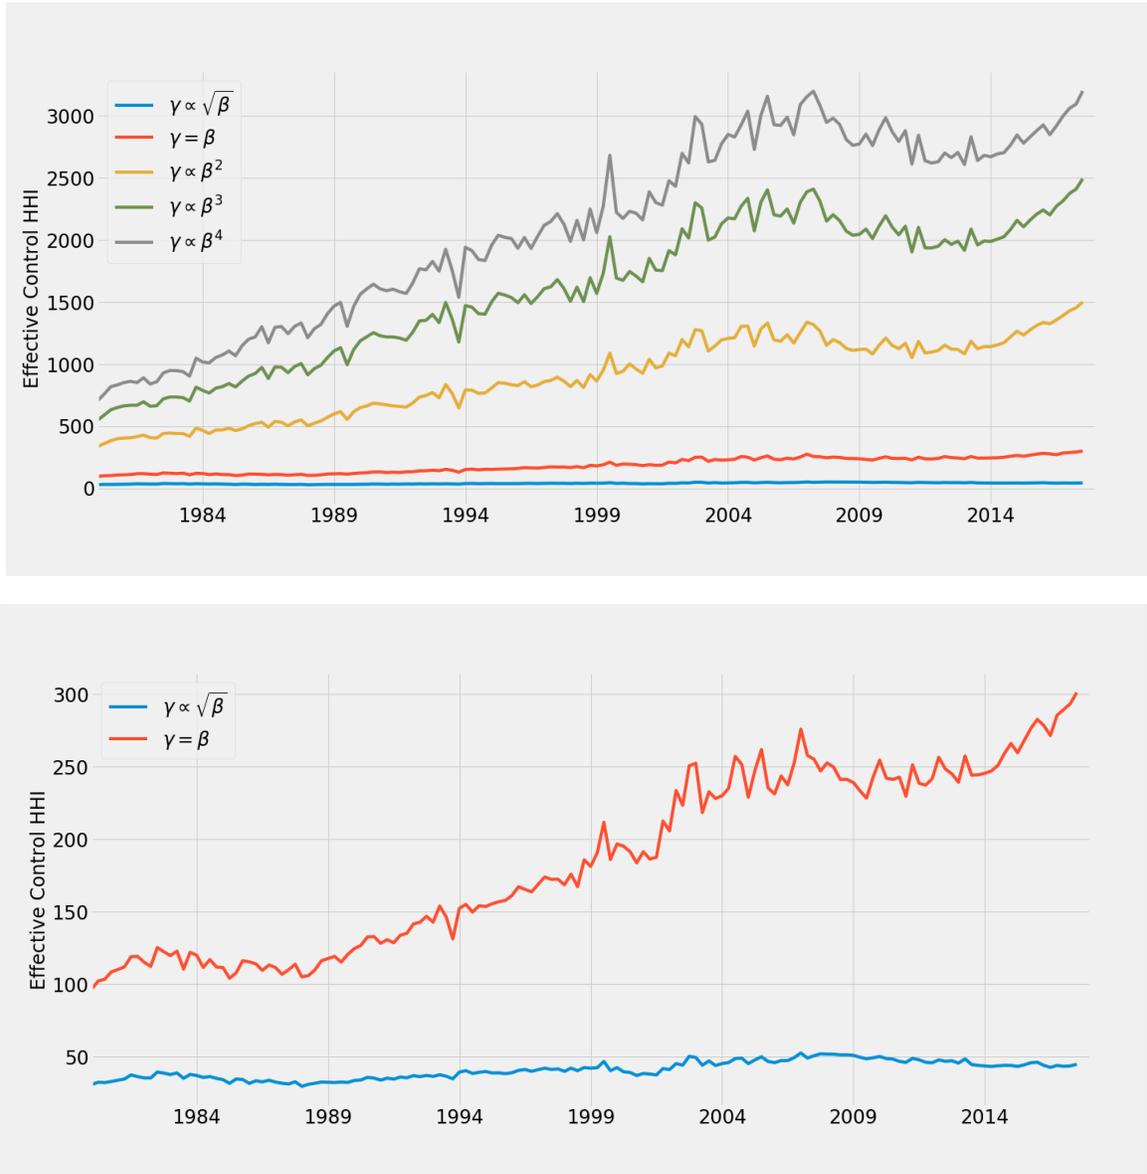
These results challenge some previously held assumptions regarding common ownership. If common ownership effects were driven entirely by the rise of the largest institutional investors, we would expect the profit weights to be more sensitive to different assumptions about effective control γ . Instead, we find that for most of the sample, more weight on large investors acts to reduce rather than increase κ . The second is that, while we know very little about how ownership translates into control, in recent years average profit weights are relatively insensitive to a wide range of control assumptions.

While our $\gamma_{fs} \propto \beta_{fs}^\alpha$ parameterization is convenient, our choice of $\alpha \in \{\frac{1}{2}, 1, 2, 3\}$ is not obviously interpretable, other than that larger values of α place more weight on the largest shareholders. In order to quantify the effects of α on effective control, we calculate a concentration measure for effective control for a particular firm f . We define $CHHI_f = \sum_s \gamma_{fs}^2$ and plot average $CHHI_f$ under different choices of α where $\gamma_{fs} \propto \beta_{fs}^\alpha$. Because this measure resembles an HHI, we can compute the equivalent number of symmetric controllers as $\frac{1}{CHHI_f}$.²⁶

In Figure 14, we report our concentration measures for effective control which we multiply by 10,000 as is common in the antitrust literature. Under proportional control, $\alpha = 1$, $CHHI = IHHI$, so that a typical firm had the equivalent of 65 symmetric “controllers” ($CHHI \approx 150$) in 1980 and around 33 symmetric “controllers” ($CHHI \approx 300$) by 2018. As we increase α , we place more weight on a small number of larger investors. For example, when $\alpha = 3$, in 2018 we find that $CHHI \approx 2500$, or that firms effectively pay attention to the four largest investors. We can also see that this measure has grown substantially over time, as it was only $CHHI \approx 600$ in 1980 (or around 17 symmetric “controllers”). This suggests we have considered the range of relevant values for α .

²⁶Unlike in our calculation of κ where we can multiply γ_{fs} by a scalar a without loss of generality, because $CHHI_f = \sum_s \gamma_{fs}^2$ the normalization of $a_f \cdot \beta_{fs}^\alpha$ matters. We choose our normalization $a_f = \left(\frac{\sum_s \beta_{fs}}{\sum_s \gamma_{fs}}\right)^2$ so that $\sum_s \beta_{fs} = \sum_s \gamma_{fs}$. This keeps the overall institutional investor share the same as we change the convexity α .

Figure 14: Control Weights γ Concentration (CHHI)

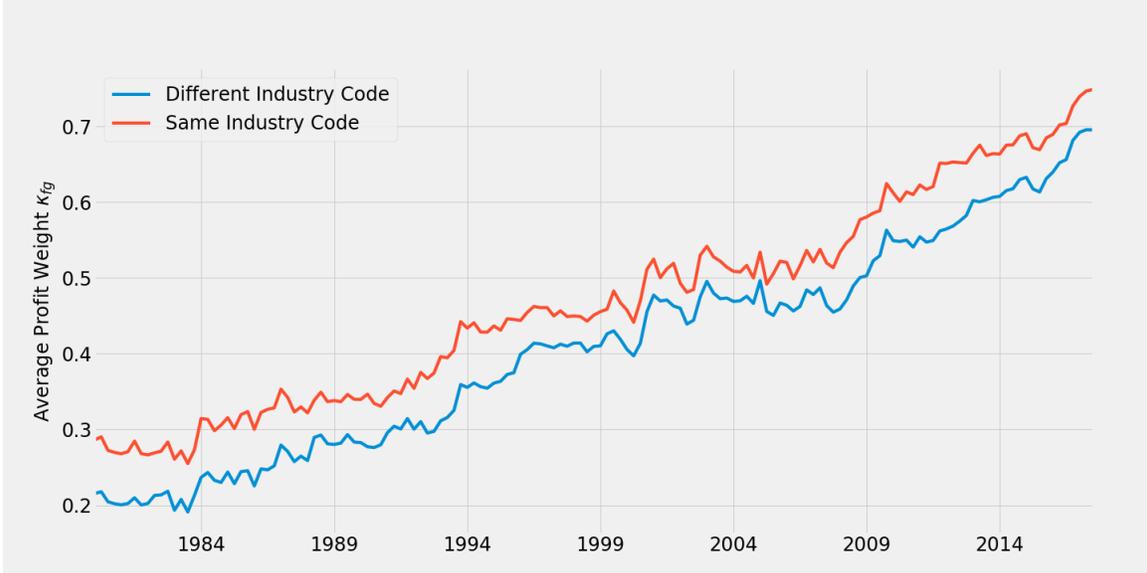


Notes: These figures average CHHI under different maintained assumptions of control weights, with $\gamma \propto \beta^\alpha$. The second zooms in on $\gamma \propto \sqrt{\beta}$ and $\gamma = \beta$.

6.2 Within-Industry and Case Studies

An obvious criticism of the above economy-wide analysis is that a pharmaceutical firm's decisions hardly affect the profits of an airline, so why do these profit weights tell us anything? What are profit weights within relevant product markets? Answering this question requires us to make assumptions about market definition, which we have eschewed so far.

Figure 15: Within & Between Industry Profit Weights

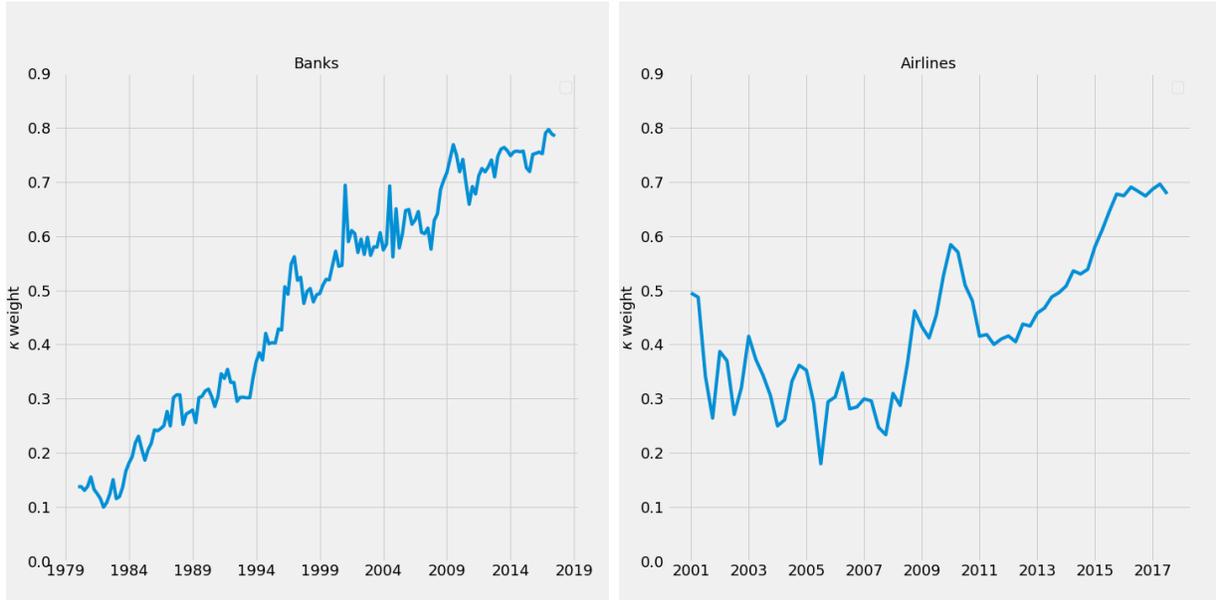


Note: This figure presents average pairwise profit weights separately between and within industry codes. Industry codes are from Compustat data as reported at the 4 digit SIC code level. Profit weight κ_{fg} computed under proportional control $\gamma_{fs} = \beta_{fs}$.

Here we follow the literature and adopt, perhaps unsatisfyingly, four-digit SIC codes as “markets.” We show average profit weights κ_{fg} over time where both firms f and g are in the same four-digit SIC code according to Compustat. While these industry classifications are often criticized, it would be problematic if the overall trends we document did not hold under this restriction. Figure 15 shows the results: the overall trend is the same, and the level is, if anything, slightly higher within SIC code.

Next, we present the average profit weight for a pair of specific industries: commercial banks, as defined by SIC code 6021 (National Commercial Banks) in Compustat that are also S&P 500 constituents, and airlines, using a hand-collected sample of 27 nationwide airline securities. The airline sample required extensive data cleaning due to the many bankruptcies and mergers over the timeframe. Details are in Appendix B.2.5. Results are depicted in Figure 16. We see that the qualitative and quantitative patterns are similar to those in the S&P 500 as a whole: a large increase in profit weights for competing firms over the past few decades.

Figure 16: Profit Weights Among Commercial Banks and Airlines



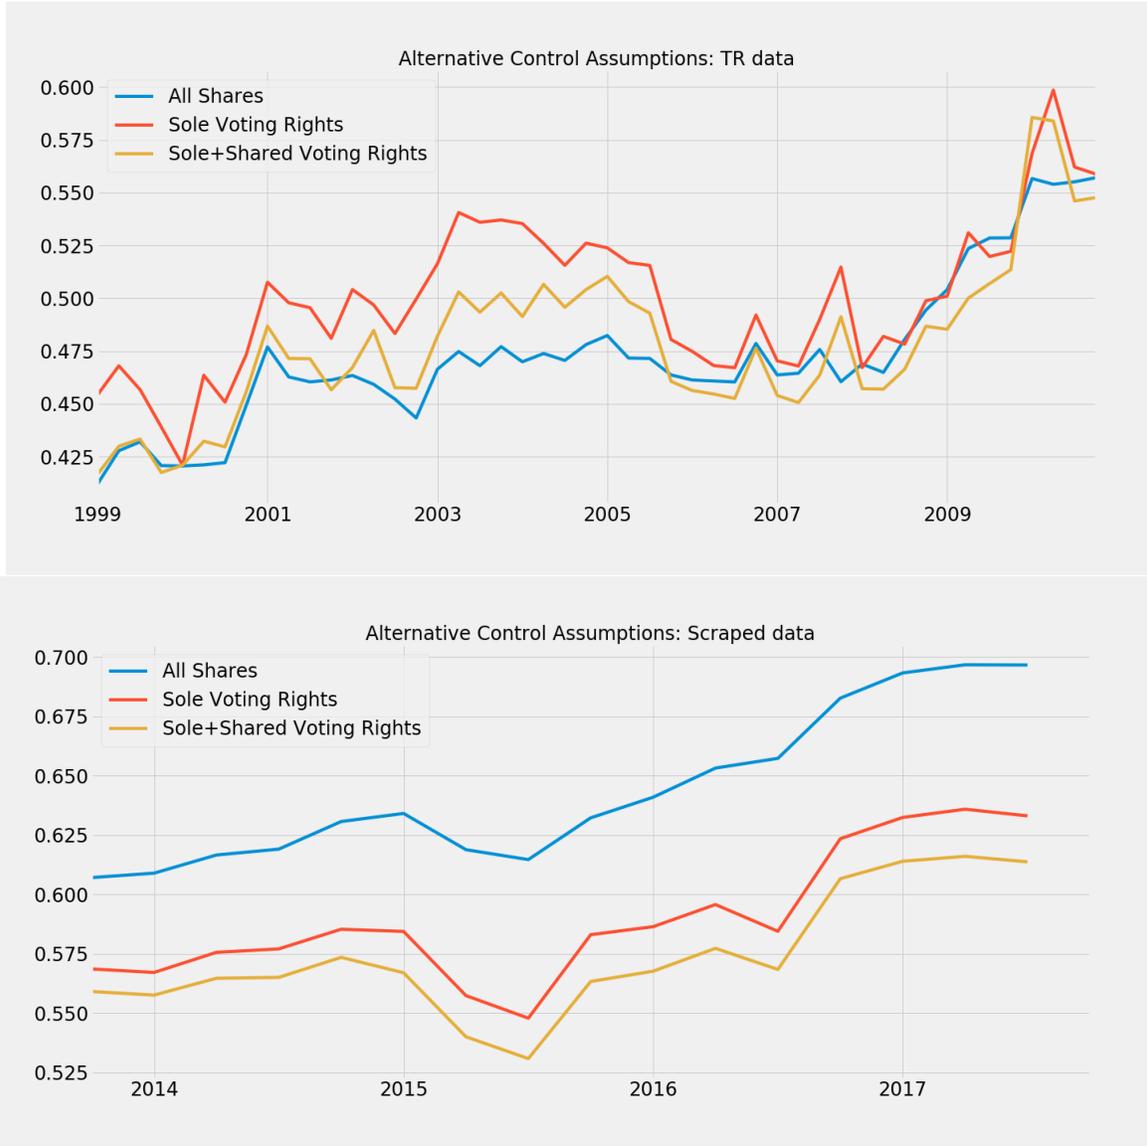
Note: This figure presents average pairwise profit weights separately for banks and airlines. Banks are for SIC code 6121 from S&P 500 sample only. Airlines are separately scraped from 13(f) filings and are available only for 2000-2017. Profit weight κ_{fg} computed under proportional control $\gamma_{fs} = \beta_{fs}$.

6.3 Voting Authority

An objection that has been raised to the literature on common ownership is that many large institutional owners do not have full discretion in voting the shares that they control. To the extent that the Pareto weights γ_f represent control rights that derive from a voting game, this would cause us to potentially over-represent common ownership concerns.

Fortunately, the 13(f) filings require investors to report not only total share holdings, but to divide these among “sole,” “shared,” and “no” voting authority shares. Therefore, to show the sensitivity of our results to alternative assumptions, we next recompute profit weights under the assumption of proportional control ($\gamma_{fs} = \beta_{fs}$) where we limit attention to either “sole” and “shared” voting authority shares, or only “sole” voting authority shares. We restrict our attention to the period beginning in 2013 when we can reliably scrape this information from XML 13(f) filings. We display the results in Figure 17 where we observe that on average κ profit weights appear to be slightly higher when we exclude nonvoting shares or shares with shared voting rights. In general, the differences between the average κ measures appear to be miniscule.

Figure 17: Alternative κ by Voting Authority



Note: This reports robustness checks where we compare the measure we report in our main results *All Shares* (blue-line) to cases where we exclude shares marked as *No Voting Rights* or *Shared Voting Rights* from the investment manager’s portfolio. These data are available in our Scraped data only for the period where we have XML filing (post 2013) and for the TR data only after 1999.

7 Conclusion

This paper has taken the common ownership hypothesis seriously to work through the economic implications at an aggregate level, examining the universe of firms in the S&P 500 from 1980 to 2017. This began with a data challenge, and so in addition to the sources already exploited by the literature, we manually recompiled investor holdings from 13(f) reports

downloaded from the SEC. We are making the source code and output of this compilation available for future researchers. From the exercise can draw a number of conclusions.

First, the implied common ownership incentives have risen substantially over the period, more than tripling from an average of 0.2 in 1980 to almost 0.7 in 2017. This rise is economically significant. A simple calibration exercise suggests that much of the rise in markups observed in Eekhout et al. (2018) is similar in magnitude to that predicted by the common ownership hypothesis over the period in our stylized example, however a closer look at the timing (which is less sensitive to the specification of the example) suggests that this relationship cannot explain much of the purported rise in markups.

Moreover, merger analysis would look substantially different in a world where firms placed a weight of 0.7 on one another's profits before a merger. This would suggest (perhaps contrary to evidence (Kwoka et al., 2015)) that we should systematically overpredict the price effects of mergers among public firms. Likewise, it suggests that purely financial transactions which eliminate or reduce common owners such as "taking firms private" or de-listing from indices might promote competition).

Second, even though the big three index funds have dominated the public debate on common ownership, much of the historic rise in common ownership incentives predates them, and is driven not by concentration in asset management but rather by a broader increase in diversification of investor portfolios. Indeed, the growth of these firms has an ambiguous relationship to common ownership incentives, as investor concentration appears both in the numerator and the denominator of the profit weight.

Third, we find a strong relationship between common ownership and retail share. We see this both in the theory, by decomposing the common ownership profit weights, as well as in the cross-sectional variation of common ownership weights between firms. Taken at face value, this implies that large firms popular with individual investors (e.g., PepsiCo, which owns Quaker Oats) should be among the least aggressive competitors towards other public firms (e.g., Kellogg's, General Mills, and Post).

Under the common ownership theory, a large retail share tends to inflate common ownership incentives by giving outsized control rights to a small set of large, diversified institutional investors. In extreme cases, which are becoming more common, this can even yield profit weights that exceed one. This is a necessary condition for "tunneling," and overturns the

traditional defense of the “widely held firm,” that in the absence of a controlling interest, investors are safe from expropriation. Again, taken at face value, our calculations imply that 10% of S&P 500 firms would have incentives to tunnel assets from or to another firm. However, unlike our main results, this finding is sensitive to the specification of control weights. Also, unlike price effects of common ownership, tunneling benefits common owners at the expense of undiversified shareholders, and so legal protections for minority shareholders may bind the expression of these incentives.

It is important to emphasize that the goal here has not been to explicitly test the common ownership hypothesis, but rather to articulate its implications in order to better form the policy debate and research efforts that are already underway. There is much more work to be done and we believe that there are two important areas for future research in particular. The first is a forensic question of understanding the mechanisms of corporate governance and the means by which common ownership incentives are, or are not, manifested. The second is to develop tests to detect effects of common ownership on market outcomes. The literature so far, including our companion piece (Backus et al., 2018), has focused on pricing. We hope that we have contributed to this effort in part by highlighting the theoretically-motivated and empirically salient variation and asymmetries in common ownership profit weights driven by, e.g., retail share, market capitalization, and the growth of indexing. This variation is entirely lost when researchers use dated, market-level indices such as MHHI. Above and beyond pricing, however, we hope that this will be useful as researches go on to examine other strategic interactions, from entry and location decisions to advertising and product development, as well as mergers and tunneling, to test the implications of common ownership more fully.

References

- Anton, M., Ederer, F., Gine, M., and Schmalz, M. C. (2016). Common ownership, competition, and top management incentives. Ross School of Business Paper No. 1328.
- Azar, J., Raina, S., and Schmalz, M. C. (2016). Ultimate ownership and bank competition. Working Paper.
- Azar, J., Schmalz, M. C., and Tecu, I. (2018). Anti-competitive effects of common ownership. *Journal of Finance*, 73(4).
- Backus, M., Conlon, C., and Sinkinson, M. (2018). Common ownership and competition in the ready-to-eat cereal industry. Working Paper.
- Backus, M., Conlon, C., and Sinkinson, M. (2019). The common ownership hypothesis: Theory and evidence. Brookings Working Paper.
- Ben-David, I., Franzoni, F., Moussawi, R., and Sedunov, J. (2018). The granular nature of large institutional investors. Swiss Finance Institute Research Paper Series.
- Berle, A. and Means, G. (1932). *The Modern Corporation and Private Property*. MacMillan.
- Bertrand, M., Mehta, P., and Mullainathan, S. (2002). Ferreting out tunneling: An application to indian business groups. *Quarterly Journal of Economics*, 117(1)].
- Boller, L. and Scott Morton, F. (2019). Testing the theory of common stock ownership. Working Paper.
- Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, 35(4):457–482.
- Bresnahan, T. F. and Salop, S. C. (1986). Quantifying the competitive effects of production joint ventures. *International Journal of Industrial Organizatoin*, 4:155–175.
- Crawford, G. S., Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*.
- Cremers, K. J. M. and Petajisto, A. (2009). How active is your fund manager? a new measure that predicts performance. *The Review of Financial Studies*, 22(9):3329–3365.
- Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

- Eekhout, J., De Loecker, J., and Unger, G. (2018). The rise of markups. Working Paper.
- Fichtner, J., Heemskerk, E. M., and Garcia-Bernardo, J. (2017). Hidden power of the big three? passive index funds, re-concentration of corporate ownership, and new financial risk. *Business and Politics*, 19(2):298–326.
- Friedman, M. (1953). *Essays in Positive Economics*, chapter The Methodology of Positive Economics. University of Chicago Press. (unfinished citation).
- Gilje, E., Gormley, T. A., and Levit, D. (2019). Who’s paying attention? measuring common ownership and its impact on managerial incentives. European Corporate Governance Institute (ECGI) - Finance Working Paper No. 568.
- Gramlich, J. and Grundl, S. (2017). Testing for competitive effects of common ownership. *Finance and Economics Discussion Series*, 29.
- Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37.
- Grullon, G., Larkin, Y., and Michaely, R. (2018). Are u.s. industries becoming more concentrated? forthcoming, *Review of Finance*.
- Gutiérrez, G. and Philippon, T. (2016). Investment-less growth: An empirical investigation. Technical report, National Bureau of Economic Research.
- Heim, S., Hüschelrath, K., Laitenberger, U., and Spiegel, Y. (2019). Anticompetitive effect of minority share acquisitions: Evidence from the introduction of national leniency programs. Working Paper.
- Jiang, W., Li, K., and Shao, P. (2010). When shareholders are creditors: Effects of simultaneous holding of equity and debt by non-commercial banking institutions. *Review of Financial Studies*, 23:3595–3637.
- Johnson, S., Porta, R. L., Lopez-De-Silanes, F., and Schleifer, A. (2000). Tunneling. *American Economic Review: Papers and Proceedings*, 90(2):22–27.
- Kahan, M. and Rock, E. (2008). The hanging chads of corporate voting. *The Georgetown Law Journal*, 96.
- Kennedy, P., O’Brien, D. P., Song, M., and Waehrer, K. (2017). The competitive effects of common ownership: Economic foundations and empirical evidence. Working Paper.

- Keswani, A., Tran, A., and Volpin, P. (2019). Institutional debt holdings and governance. Working Paper.
- Kwoka, J., Greenfield, D., and Gu, C. (2015). *Mergers, Merger Control, and Remedies: A Retrospective Analysis of U.S. Policy*. Mit Press.
- Matvos, G. and Ostrovsky, M. (2008). Cross-ownership, returns, and voting in mergers. *Journal of Financial Economics*, 89:381–403.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.
- O’Brien, D. and Salop, S. (2000). Competitive effects of partial ownership: Financial interest and corporate control. *Antitrust Law Journal*, 67(3):559–614.
- Porta, R. L., Lopez-De-Silanes, F., and Schleifer, A. (1999). Corporate ownership around the world. *The Journal of Finance*, 54(2):471–517.
- Posner, E. A., Scott Morton, F. M., and Weyl, E. G. (2017). A proposal to limit the anti-competitive power of institutional investors. *Antitrust Law Journal*, 81(3).
- Reynolds, R. J. and Snapp, B. R. (1986). The competitive effects of partial equity interests and joint ventures. *International Journal of Industrial Organization*, 4(2):141–153.
- Rotemberg, J. (1984). Financial transactions costs and industrial performance. Working Paper.
- Securities, U. and of Inspector General, E. C. O. (2010). Review of the sec’s section 13(f) reporting requirements. Report No. 480.
- Shambaugh, J., Nunn, R., Breitwieser, A., and Liu, P. (2018). The state of competition and dynamism: Facts about concentration, start-ups, and related policies. Technical report, The Hamilton Project: Economics Facts.
- Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50.
- Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information acquisition and under-diversification. *The Review of Economic Studies*, 77(2):779–805.

Appendices

A Main Appendix

A.1 Common Ownership and Oligopoly Models

Relationship to Cournot

Much attention in the common ownership literature has been paid to the Modified Herfindahl–Hirschman Index (MHHI) concentration measure, which is derived from a Cournot oligopoly model of competition in O’Brien and Salop (2000).²⁷ MHHI extends the traditional concept of HHI to incorporate common ownership, and is defined from the following firm objective function:

$$\max_{q_f} \pi_f(q_f, q_{-f}) + \sum_g \kappa_{fg} \pi_g(q_f, q_{-f}).$$

After taking the FOC (where η represents the elasticity of demand) we get:

$$\frac{P_f - MC_f}{P_f} = \frac{1}{\eta} \sum_g \kappa_{fg} s_g.$$

Which gives the share weighted average markup of:

$$\sum_f s_f \frac{P_f - MC_f}{P_f} = \frac{1}{\eta} \underbrace{\sum_f \sum_g \kappa_{fg} s_g s_f}_{MHHI}$$

— where $MHHI = \underbrace{\sum_f s_f^2}_{HHI} + \underbrace{\sum_f \sum_{g \neq f} \kappa_{fg} s_f s_g}_{\Delta MHHI}$.

Note that many of the papers that regress price on measures of ownership separately include HHI and $\Delta MHHI$ as independent variables. It is important to point out that both measures vary only at the *across markets* while the incentive terms κ_{fg} vary *across firms* within a market.

²⁷Originally the MHHI was derived by Bresnahan and Salop (1986) in the context of a joint-venture.

Relationship to Bertrand

The Price Pressure Index (PPI) is similarly defined for differentiated Bertrand competition. We consider the objective function for firm f when setting the price p_j holding fixed the prices of all other products p_{-j} . As firm f raises the price p_j some consumers substitute to other brands owned by f : $k \in \mathcal{J}_f$ on which it receives full revenue, and substitute brands owned by competing firms g : $k' \in \mathcal{J}_g$ for which it acts as if it receives a fraction of the variable profit κ_{fg} :

$$(p_j - mc_j)q_j(p_j, p_{-j}) + \sum_{k \in \mathcal{J}_f} (p_k - mc_k)q_k(p_j, p_{-j}) + \sum_g \kappa_{fg} \cdot \left(\sum_{k' \in \mathcal{J}_g} (p_{k'} - mc_{k'})q_{k'}(p_j, p_{-j}) \right).$$

When solving the FOC it is helpful to do two things: (1) divide through by $-\frac{\partial q_j}{\partial p_j}$; (2) define the diversion ratio $D_{jk} = -\frac{\frac{\partial q_k}{\partial p_j}}{\frac{\partial q_j}{\partial p_j}}$, this gives:

$$p_j - mc_j = -q_j / \frac{\partial q_j}{\partial p_j} + \sum_{k \in \mathcal{J}_f} (p_k - mc_k)D_{jk} + \sum_g \kappa_{fg} \cdot \left(\sum_{k' \in \mathcal{J}_g} (p_{k'} - mc_{k'})D_{jk'} \right). \quad (4)$$

This clarifies what common ownership does under differentiated Bertrand competition. It raises the effective opportunity cost of selling product j . Now as p_j rises, some customers are *recaptured* by other products controlled by the same firm $k \in \mathcal{J}_f$ (this is the usual multiproduct oligopoly effect), also by products controlled by competing (but commonly owned firms) $k' \in \mathcal{J}_g$ with $\kappa_{fg} > 0$.

A.2 Alternative Similarity Measures

Our primary interest is how overlapping ownership relates to *profit weights* or *cooperation incentives* among firms in the product market. The measure in Rotemberg (1984) or O'Brien and Salop (2000) is shown in (1) to be an L_2 measure. We could construct alternative measures of investor overlap, such as an L_1 measure.

For example:

$$L_1(\beta_f, \beta_g) = \frac{1}{2} \left(\sum_s \beta_{fs} + \sum_s \beta_{gs} - \sum_s |\beta_{fs} - \beta_{gs}| \right).$$

It is important to point out that $\sum_s \beta_{fs} < 1$. This is because the set of investors $s \in S$ contains only large institutional investors who provide a 13(f) form to the SEC. We can think about $\sum_s \beta_{fs} = 1 - r_f$ where r_f represents the retail investor share in firm f . The L_1 measure varies from $[0, 1]$. It is highest if we don't have any retail investors, yet all investors hold the same portfolio so that $\beta_{fs} = \beta_{gs}$. Likewise the L_1 measure declines as portfolios become more dissimilar $|\beta_{fs} - \beta_{gs}|$ becomes large.

A.3 Alternative κ Weights

An alternative specification is offered in Crawford et al. (2018), who develop profit weights in the different context of vertical incentives.²⁸ There each investor constructs their ideal weight $\tilde{\beta}_{fs} = \frac{\beta_{fs}}{\sum_g \beta_{gs}}$, and firms form a γ weighted average of their investors' desired weights. In the construction of κ_{fg} we find it more transparent to normalize $\tilde{\gamma}_{fs} = \frac{\gamma_{fs}}{\sum_g \beta_{gs}}$ so that:²⁹

$$\kappa_{fg}^{CLWY} = \frac{\langle \gamma_f, \tilde{\beta}_g \rangle}{\langle \gamma_f, \tilde{\beta}_f \rangle} = \frac{\langle \tilde{\gamma}_f, \beta_g \rangle}{\langle \tilde{\gamma}_f, \beta_f \rangle} = \frac{\beta'_f \cdot g_f(\beta) \cdot \beta_g}{\beta'_f \cdot g_f(\beta) \cdot \beta_f}. \quad (5)$$

The CLWY weights are just the common ownership weights, but with a different assumption on control (γ). Investors with large diversified portfolios (such as index funds), have larger values for $\sum_g \beta_{gs}$ and receive a smaller weight $\tilde{\gamma}_{fs}$.³⁰ One justification for this re-weighting might be that investors with large portfolios may become *inattentive* (Van Nieuwerburgh and Veldkamp (2010), Gilje et al. (2019)).

An alternative measure which more explicitly addresses investor inattention is proposed by Gilje et al. (2019). Inattention is related to the *portfolio share* of firm f rather than the *normalized cash flow*. With a bit of work, one can show that their measure:

$$GGL_{fg} = \beta'_f \cdot g_f(\beta_s) \cdot \beta_g. \quad (6)$$

where $g_f(\beta_s)$ is a $S \times S$ diagonal matrix with entries which are a function of the *portfolio*

²⁸It is important to note that Crawford et al. (2018) are not considering the common ownership hypothesis directly but rather examining incentives for vertical integration and bargaining among MVPDs and content providers where the former often have a partial ownership stake in the later.

²⁹ $g_f(\beta)$ represents a diagonal $S \times S$ weighting matrix with entries $\frac{1}{\sum_g \beta_{gs}}$

³⁰The Crawford et al. (2018) paper only considers firms within the same industry in $\sum_g \beta_{gs}$ (albeit in a very different context). It is hard to understand what an equivalent assumption would be for the entire S&P 500 Index.

share which f comprises in s 's portfolio: $\alpha_{fs}(\beta_s) = \frac{\beta_{fs}v_f}{Y + \sum_g \beta_{gs}v_g}$, where v_f represents the market capitalization of firm f . The authors consider a number of functional form choices for $g_f(\alpha_{fs})$,³¹ which has the interpretation as the probability that investor s pays attention to the actions of firm f . It is important to note that while this overlap measure is a quadratic form like (1) it is not normalized by $\gamma'_f \cdot \beta_f$ and thus does not directly represent a change in the *profit weights* but rather some other difference in f 's objective function.³²

Because both the CLWY and GGL measures have the effect of putting less weight γ_{fs} on investors who are more diversified, we don't apply them to our study of the entire S&P 500 index. These measures seem more appropriate when examining a single industry at a time (such as the cable television industry).

B Data Appendix

B.1 Data Sources

Our main data source is the universe of 13(f) filings from 1980–2017. The 13(f) form is a mandatory SEC filing for institutional investors with over 100M USD in assets. We compile 13(f) filings from two sources: for the period 1980–1999, we use the Thompson Reuters s34 database. For 2000–2017, we use our own proprietary dataset, for which we are making the code publicly available, based on scraped and parsed source documents from the SEC. The latter dataset is discussed below in Appendix B.2.

For many filings there are multiple filing dates (fdate) for the same report date (rdate). This happens when filings are amended, often because of an error in the original submission or in the case of a stock split. For an ordinary revision, e.g. in case of error, we would like to take the last fdate for each rdate. However, revisions following a stock split are often retroactively applied to report dates prior to the split event itself, and in these cases we want to use the first filing date. This is a frequent issue in the data.

³¹For example: linear, convex, etc.

³²We also caution against normalizing the GGL_{fg} measure and interpreting it as a *profit weight* under any circumstances. Consider the case of the breakfast cereal industry: Kellogg's has a market cap of \$21 Billion and derives most of its revenue from its breakfast cereal business. Pepsi has a market cap of \$165 Billion, and derives only a small fraction of its revenue from its breakfast cereal business (Quaker Oats). An investor holding the S&P 500 index would then value \$1 of profits 8 times as much for Pepsi as it would for Kellogg's even though they own the same fraction of each firm.

In order to resolve the problem, we identify the universe of stock splits for all S&P 500 firms in our sample using the CRSP data CFACSHR multiplier, and from that identify a set of quarter-firm pairs at which we use the *first*, rather than the last fdate for duplicate rdate reports.

In addition, there is a notable exception: in several instances BlackRock holdings appear to conflate the two dates, and so for BlackRock we use the filing date exclusively. This resolves the otherwise inexplicable disappearance of BlackRock Inc. from the s34 in 2010q2 and 2010q3.

13(f) filings use investor-reported values and tallies of shares outstanding and these frequently contain errors, so we use the CRSP monthly database, merged on contemporaneous CUSIP codes (nCUSIP), to compute these figures.

From CRSP we also obtain historical data on membership in the S&P 500.

From Compustat we obtain additional fields: Aggregate short interest for each member firm by quarter, and the number of business segments, as reported in the Compustat (North America) Database. There are two limitations of this data. First, coverage is imperfect. Of the 1,587 firms that ever appear in the S&P 500 between 1980 and 2017, we lack data on business segments for 209 of them. Second, the data are self-reported. What constitutes a “business segment” is an ill-defined notion, and may vary from firm to firm. Moreover, as suggested by [where’s that citation again], there may be incentives for strategic misreporting here.

B.2 Alternative Dataset

Given our concerns with the Thomson Reuters dataset, as well as the concerns voiced by others such as WRDS and Ben-David et al. (2018), we also recreated a dataset of 13(f) holdings directly from the source filings. This involved gathering approximately 25GB of 13(f) filings from the SEC, for the time period 1999-2017. Mandatory electronic filing of 13(f) forms began in 1999; for earlier years, coverage is poor. These files are then parsed to extract holdings of S&P 500 firms. The parsing is handled slightly differently for filings made before the third quarter of 2013, as starting then, the SEC mandated an XML filing format. The code is written in Perl and uses regular expressions to match text patterns

corresponding to holdings. The code is freely available from the authors. Note that we do not claim that every single one of the nearly 19M observations in our scraped and parsed sample are correct; we have a number of examples of filings that are so irregular as to be un-parsable. However, we believe this alternative dataset does capture many filings missing from Thomson Reuters, and is more consistent over time in a number of measures.

B.2.1 Pre-XML Parsing

In these filings (covering January 31, 1999 through June 30, 2013), most reports are fixed-width tables of holding name, holding CUSIP, value, number of shares, and then a possible breakout of shares by voting rights. For each file, our code first extracts the reporting date, filing date, CIK of the filing firm, and form type from the filing header. The code then looks for any line of text that contains an S&P 500 CUSIP for that form’s reporting period. As firms on occasion report derivative holdings for a CUSIP, we drop any records that match any of the following words (case insensitive, with word boundaries on both sides): put, call, conv bd, conv bond, opt. The code then attempts to match a pattern that is consistent with most filings: a CUSIP, followed by a value, followed by a number of shares. As filings are far from uniform, the code also attempts to correct a number of common problems: for example, in some cases there is no space in between the value and the number of shares; the code attempts to discern the correct breakdown based on the price and shares outstanding for that holding in that quarter, as reported by CRSP. The code then outputs a list of share holdings at the CIK-CUSIP-reporting date level.

B.2.2 XML Parsing

For filings beginning in the third quarter of 2013, our code exploits the XML structure when parsing for filings. As before, we first extract the reporting date, filing date, CIK of the filing firm, and form type from the filing header. We then separate the file into “infotable” XML objects. We keep all such objects that have a CUSIP element that contains an S&P 500 CUSIP for that form’s reporting date. We further drop any records that have a “put” or “call” element, or a “principal amount” element. We finally drop any where the title of class contains “put” or “call” surrounded by word boundaries, or that begins with “opt” or “var” (all case insensitive). The code also extracts the reported value from the value element of the information table, and compares that to the extracted number of shares times

the CRSP-reported price at the reporting date. If the two values differ by less than 10%, we also include a flag in the output that the data appear valid (we use this when there are multiple filings per reporting date for a CIK-CUSIP).

B.2.3 Final Cleaning

We take the output of the parsing steps above and obtain a dataset of institutional holdings. In the case of restated filings, we keep the initial filing unless the reported value and number of shares appears impossible, in which case we keep the first rational report filed within 90 days of the mandatory reporting date. We consolidate all BlackRock entities into the same entity and collapse their holdings (while the argument could be made for collapsing other investment management firms' sub-entities, we solely do this for BlackRock given the practice in the literature). Finally, we drop 331 observations where the reported shareholdings are greater than 50% of shares outstanding. Some of these observations are correct: for example, Loews Corporation, an S&P 500 component, controlled more than 50% of common stock of Diamond Offshore Drilling, another S&P 500 component, from 2009-2016. Other records among these 331 observations appear to be either parsing errors or raw data errors. For example, in 2014, Guardian Life (CIK: 901849) reported holdings in Noble Corp (CUSIP: G6543110) of over 144 billion shares valued at \$144 billion dollars, while Noble Corp had a just over 250M shares outstanding and a market capitalization of \$5.6B.³³ The result is a dataset of 18,968,596 observations of unique CIK-CUSIP-record date holdings across 75 reporting quarters.

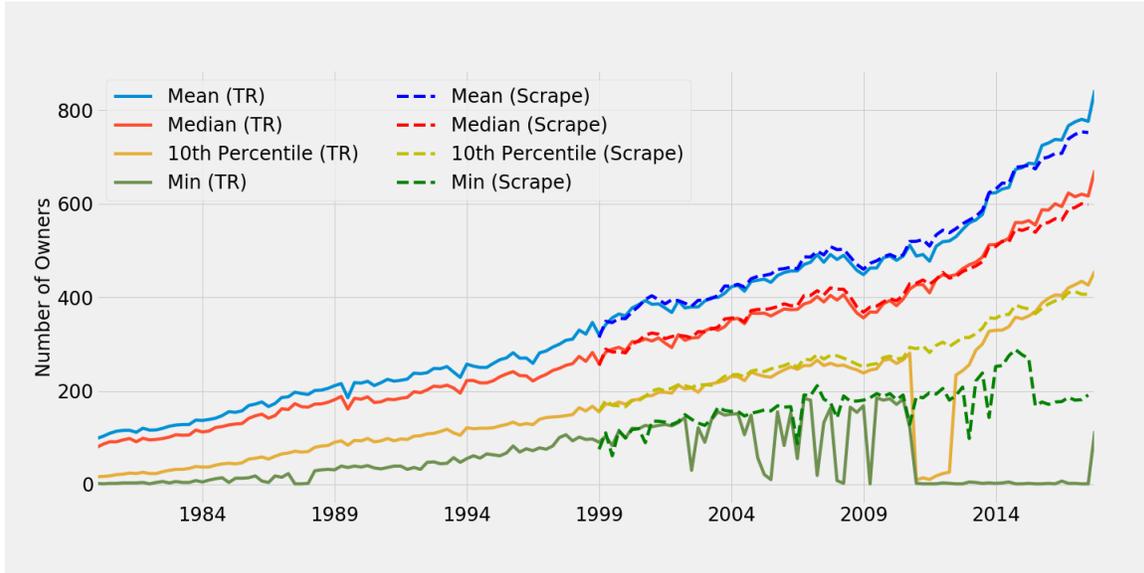
B.2.4 Comparison to Thomson Reuters

We do two primary comparisons against the Thomson Reuters (TR) dataset, followed by a deep-dive on some particular holdings where the TR dataset seems deficient. First, we consider the number of 13(f) owners per S&P 500 firm. Second, we consider the number of S&P 500 single-class of share firm that has over 100 owners in the dataset. In both cases we indicate the TR data with solid lines and our scraped data with dashed lines.

Appendix Figure 18 plots the mean, median, 10th percentile, and minimum of the number

³³Guardian's XML filing is available at:
https://www.sec.gov/Archives/edgar/data/901849/000072857214000014/xslForm13F_X01/SepGLIC.xml

Figure 18: Owners Per Firm



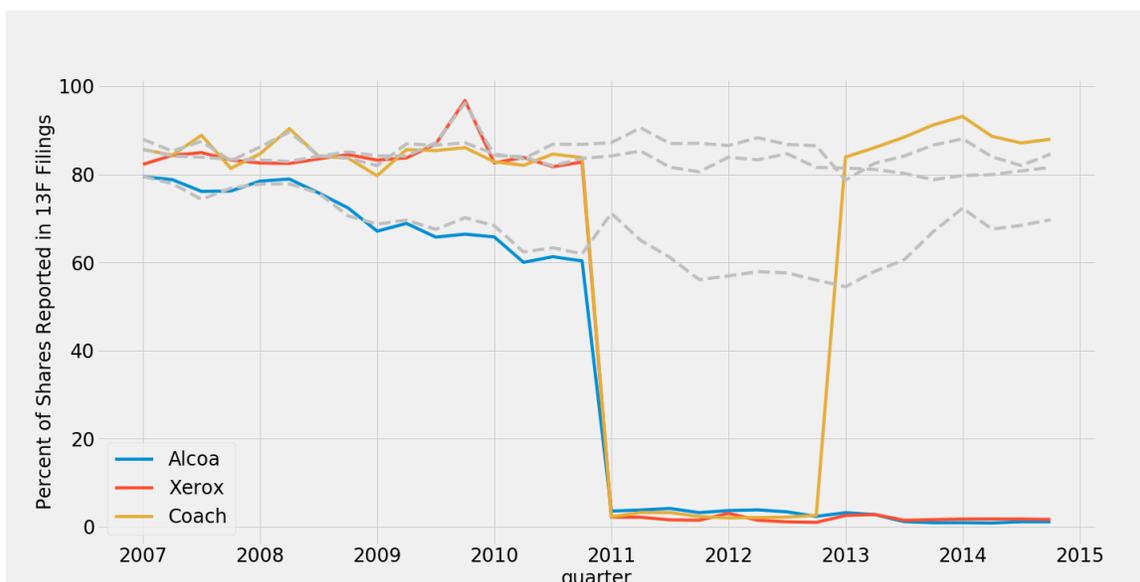
Notes: This figure depicts statistics of the number of investment managers per issue in the S&P 500 over time. The TR data uses “mgno,” manager number, as a manager while the scraped data uses the SEC’s CIK number for a manager.

of owners of S&P 500 firms over time. Solid lines are the TR data, dashed are the scraped data. As is clear, there appears to be an issue in the TR data where some firms show few owners, as evidenced by the “min” line. In addition, the “10th percentile” line shows that there is a series of quarters beginning in 2011 where over 10% of S&P 500 firms have very few reported owners. In contrast, the dashed lines show more consistent patterns in the scraped data.

Appendix Figure 2 presents these data in a different way: for each quarter, it plots the number of single-class of share firms held by 13(f) managers in the respective datasets, limited to issuances held by at least 100 investment managers. Note that this should be below 500 as we omit firms with multiple classes of shares. As is immediately clear, there is an issue with the TR dataset beginning in 2011. If a firm appears to have very few owners, this directly impacts κ through the $IHHI$, as shown in equation 1.

Finally, Appendix Figure 19 does a “deep dive” for three S&P 500 securities around the 2011 window where the TR dataset appears to have deficiencies. The plot shows, in solid color lines, the percent of shares outstanding reported to be held by 13(f) managers for three major firms: Alcoa, Xerox, and Coach in the TR dataset. The solid lines show that prior 2011, 13(f) investment firms held between 60% and 90% of these firms. However, in 2011, that falls dramatically to under 10%, before reverting back in 2013 for one of the three

Figure 19: Examples of TR Coverage Issues



Notes: This figure sums the holdings of all 13(f) managers for three firms: Alcoa (permno: 24643, CUSIP: 03965L10), Xerox (permno: 27983, CUSIP: 98412110), and Coach (permno: 88661, CUSIP: 87603010). TR data series are plotted solid lines, the authors' scraped and parsed data in dashed lines.

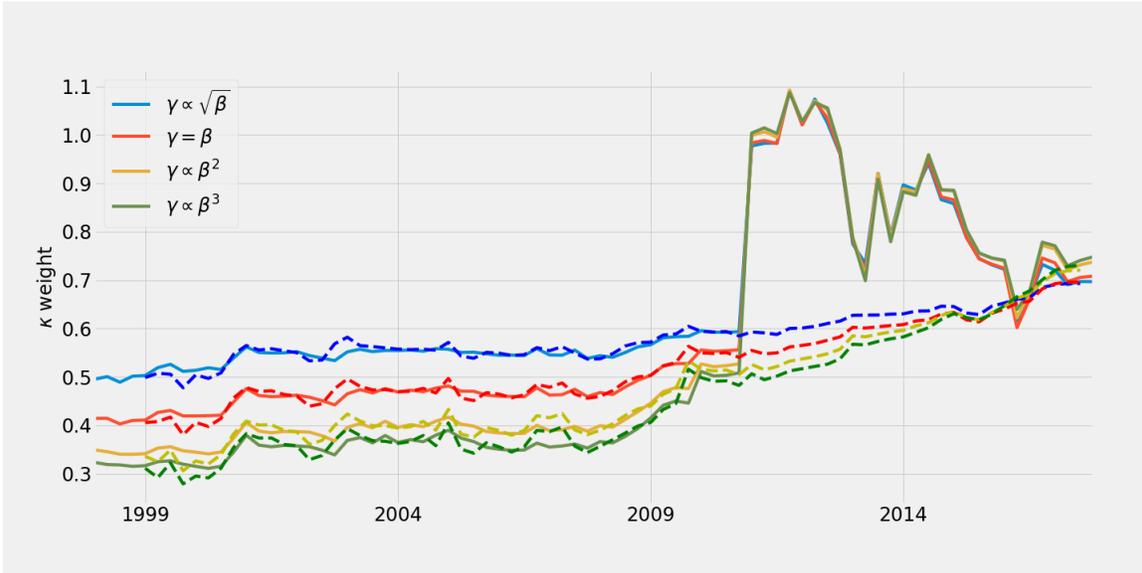
firms. In dashed lines are the percent of shares outstanding found in our scraped and parsed dataset. The TR data seem unreliable while the scraped data present a reasonable time series for institutional ownership.

To summarize the issue with the Thomson Reuters dataset, Appendix Figure 20 shows what the average computed profit weights (the κ values) would be using the raw Thomson Reuters data in solid lines, and our new dataset in dashed lines. As is clear, the Thomson Reuters dataset has coverage deficiencies in several years that result in large swings of the average κ , even reaching improbably high values starting in 2010.

B.2.5 Airline Sample

Most airlines are not S&P 500 constituents during this time period (one notable exception is Southwest Airlines). Therefore, we began by assembling a set of CUSIPs for airlines from CRSP and arrived at a set that consisted of major airlines (excluding foreign and regional). We were careful to drop any reported holdings after any bankruptcy declaration: there are many cases of institutional investors continuing to report holdings of non-existent securities. We also gather CUSIPs for entities that emerge from bankruptcy, or from mergers. The final

Figure 20: Profit Weight Comparison, TR vs Scraped Data



Notes: This figure presents average profit weights for both the TR and scraped datasets by control assumption. Solid lines denote Thomson-Reuters dataset and dashed line denotes scraped and parsed 13(f) dataset.

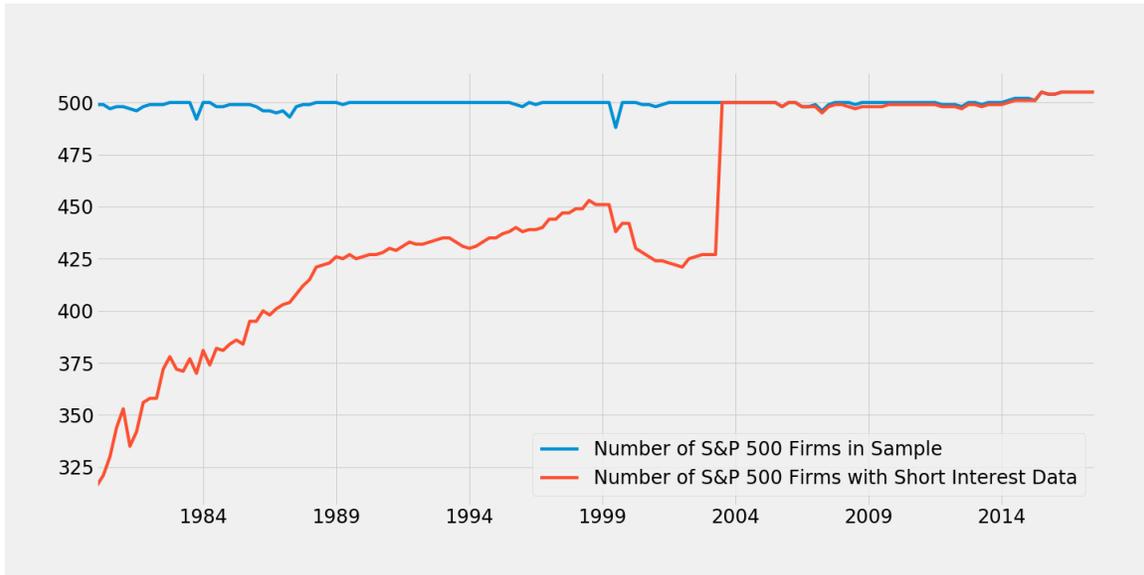
set of airlines consists of: AirTran, Alaska, American, Continental, Delta, Eastern, Hawaiian, JetBlue, Northwest, Pan Am, Southwest, Spirit, Trans World, United, US Airways, and Virgin America. Several of these have multiple CUSIPs over this time period. We do not adjust for insider holdings in this exercise, although in practice this may be a good thing to do if insider holdings are significant.

B.3 Short Interest

A known limitation of the 13(f) data for calculating institutional ownership is that short interest is double-counted. When an investor takes a short position they borrow shares from another investor and sell them, with a promise to repay the shares at a later date. These shares are then double-counted, reported on form 13(f) by both the initial investor as well as the investor to whom they are sold. It is for this reason that one can often observe “institutional ownership,” as reported in online sources, in excess of 100%.

Data on short interest are obtained from the Compustat short interest supplemental dataset. These data are available at the firm level, not the investor level. Moreover, evidence suggests that even if we had data at the investor level, it is not clear how we should think about control

Figure 21: Coverage of Short Interest Data

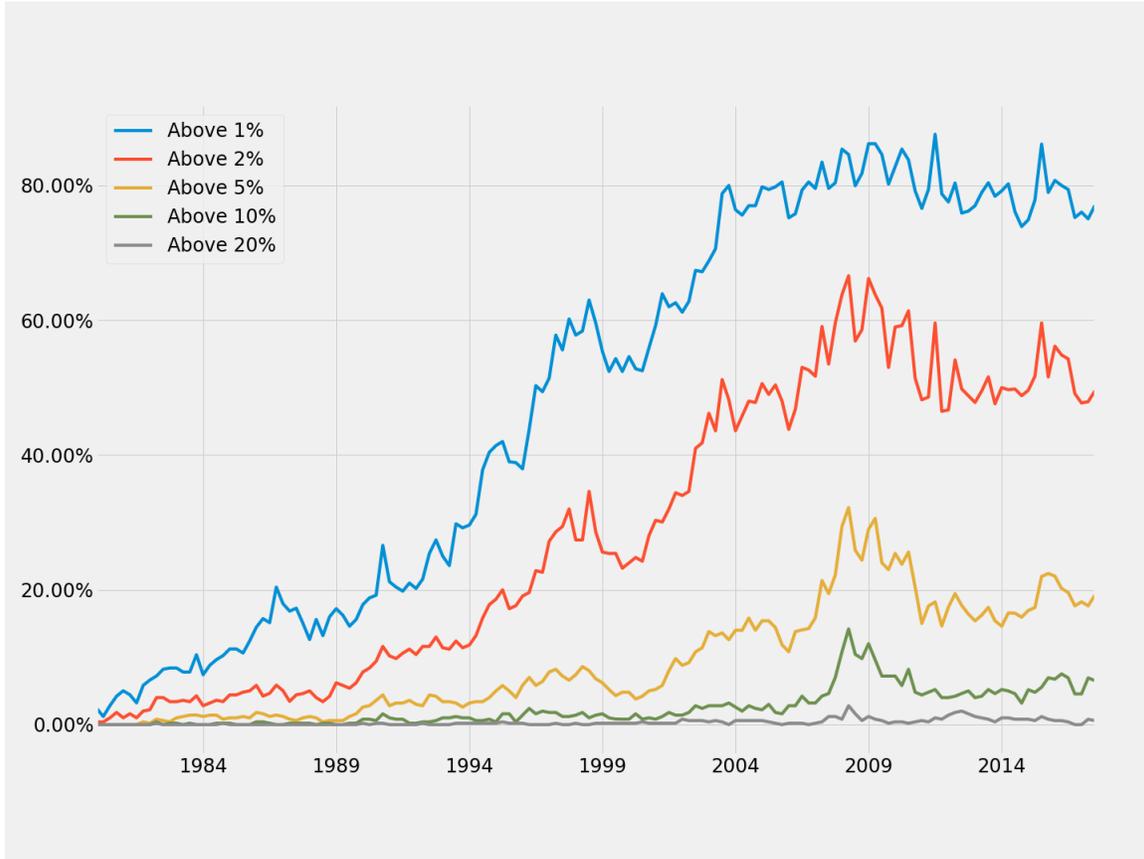


Notes: This figure compares the number of firms in the sample against the number of firms for which we observe the level of short interest in Compustat.

rights. While it seems intuitive that only the actual holder of the stock should cast votes in corporate governance activities, in practice it seems that *both* the initial investor as well as the current holder may end up voting the same shares, see Kahan and Rock (2008).

Appendix Figure 21 characterizes the coverage of the short interest data in our sample, which improves dramatically after 2004. Appendix Figure 22 documents the degree of short interest. We see that while short interest in excess of 2% is quite common, short interest in excess of 20% is quite rare.

Figure 22: Short Interest Levels



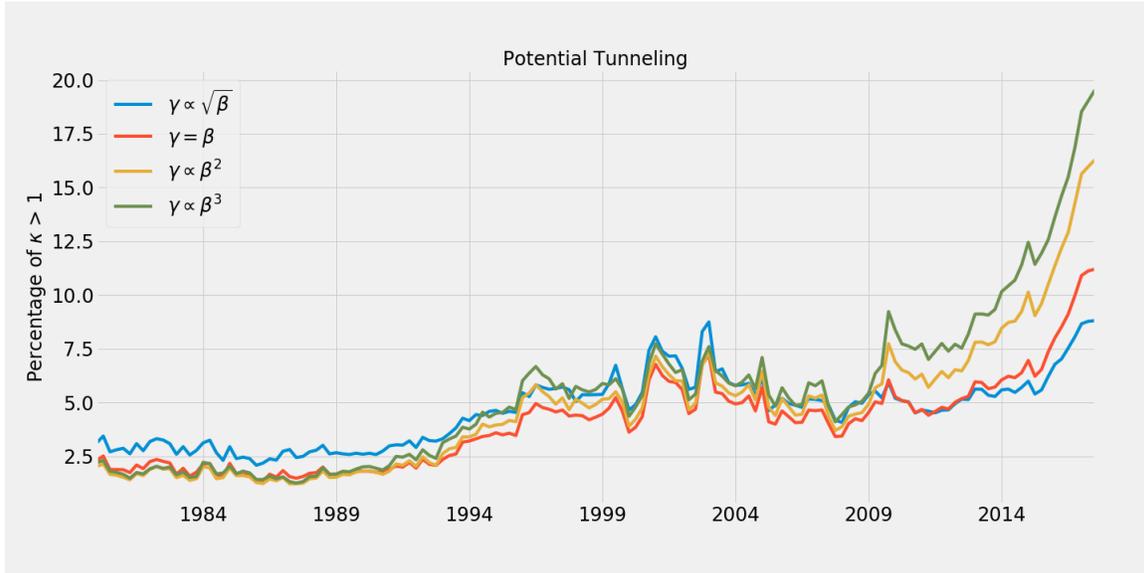
Notes: This figure shows the distribution of short interest levels over time reported for our sample firms in Compustat.

C Additional Tables and Figures

C.1 Tunneling and Specifications of Control

Here we re-create figure 9 under alternative specifications of γ . The results are depicted in Figure 23. While the proportion of pairwise profit weights greater than one is insensitive to specification from 1980 to the late aughts, it becomes very sensitive in the period following. This coincides with the rise of the “Big Three” from Figure 5. If we place more weight on the holdings of these large firms in constructing control rights, we find substantially greater incentives for firms to engage in tunneling.

Figure 23: Potential Tunneling Incentives $\kappa > 1$, Alternative Control Specifications

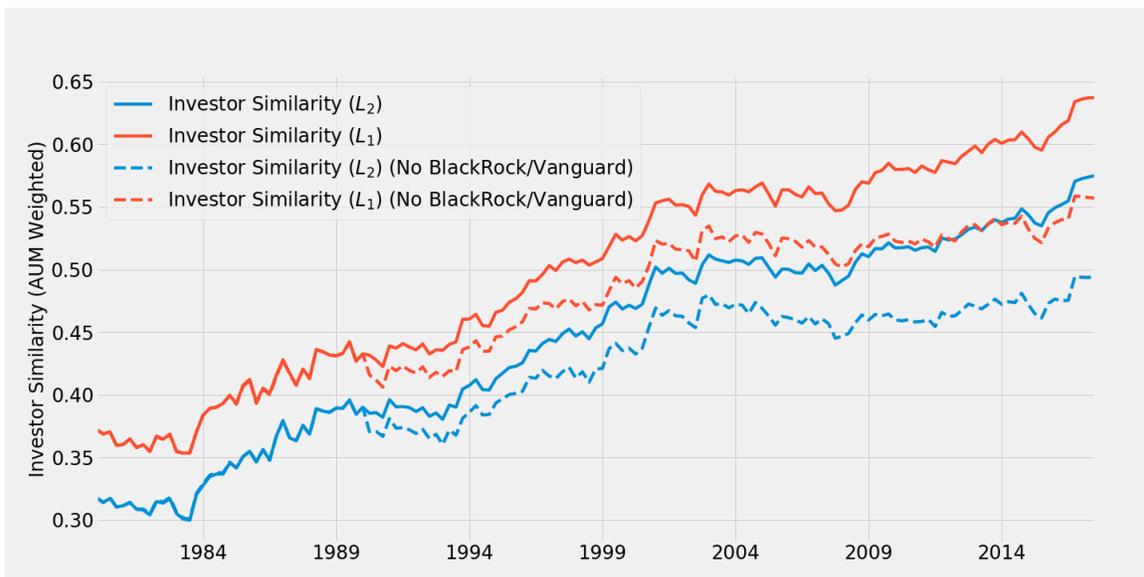


Note: This figure reports the fraction of pairwise profit weights $\kappa_{f,g} > 1$ in each period under different control assumptions.

C.2 Rise of Indexing

To understand the role that large institutional investment firms BlackRock and Vanguard have had on the rise of indexing, we now revisit Figure 8, which plotted the similarity of investor portfolios to the “market” portfolio of the S&P 500. The figure computes the average in each time period weighted by assets under management. We now re-compute these figures removing BlackRock and Vanguard entirely. This allows us to see the contribution to the rise in indexing attributable to those firms. Appendix Figure 24 shows the results. From this, we see that while those particular firms are indeed large, they are large and particularly indexed, and as a result they have had a sizable effect on the increase in indexing.

Figure 24: Similarity Between Investor Portfolios and S&P 500 Index



Notes: This figure depicts L1 and L2 similarity measures comparing investor portfolios weighted by investor AUM within our sample of S&P 500 assets. Dashed lines show the result if we exclude BlackRock and Vanguard.