Common Ownership in America: 1980–2017

By Matthew Backus, Christopher Conlon, and Michael Sinkinson[*]

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: L0, L21, L13, G34

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. Investors' 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 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, Raina and Schmalz, 2016; Azar, Schmalz and Tecu, 2018).[1] Contemporaneously, De Loecker, Eeckhout and Unger (2020) 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

    [1]The latter paper had over 325 citations as of the date of this draft according to Google Scholar.

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 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 prior empirical exercises. Much of that work depends on aggregate measures of common ownership based on the so–called Modified Herfindahl–Hirschman Index, which is a 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 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, both within markets and even within pairs of firms. Already building on this observation, Boller and Scott Morton (2019) show 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. Among additional implications, the presence of privately held firms would imply greater competition, and seemingly–innocuous financial events such as being de–listed from a market index may have product market effects. 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

---

[2]Exceptions include Kennedy et al. (2017) and Backus, Conlon and Sinkinson (2020a).

FIGURE 1. COMMON OWNERSHIP PROFIT WEIGHTS OVER TIME

Note: This figure depicts the mean implied profit weight across all pairs of firms in the S&P 500 index by year, denoted by κ, excluding own profit weights which are normalized to 1. The profit weights are defined as $\kappa_{fg} = \frac{\sum_{\forall s} \gamma_{fs} \beta_{gs}}{\sum_{\forall s} \gamma_{fs} \beta_{fs}}$, where β_{fs} denotes the fraction of firm f held by shareholder s, and γ_{fs} is the control weight firm f places on shareholder s. See Section I for an explicit formula for common ownership weights and the full derivation.

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 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 different weightings of the data. For comparison: a profit weight of 0 corresponds to what we expect in a world of profit–maximizing firms, and a profit weight of 1 corresponds to the weight that a merged firm places on an acquired subsidiary business (or, equivalently, full collusion). We find that when weighting observations equally, the average pairwise profit weights implied by the common ownership hypothesis more than tripled among S&P 500 firms, from just over 0.2 in 1980 to almost 0.7 in 2017. Weighting the observations by either market cap or revenue does not qualitatively change the result and we focus on the equal weight average going forward. Online Appendix Figure A-6 shows percentiles of the distribution of profit weights over time, reflecting a broad increase in these measures.

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., 2018) or proposed altogether new measures, e.g. the measure of Gilje, Gormley and Levit (2020), 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, Conlon and Sinkinson, 2019, 2020a). 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 does not 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 (Backus, Conlon and Sinkinson, 2020b). While most research to date in this area has used a commercial dataset of these holdings (Thomson Reuters, 2020), 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 II. If one were to complete our exercise using only the commercial dataset, one would reach different qualitative and quantitative conclusions, as shown in Online Appendix Figure A-3, which contrasts Figure 1 using the commercial dataset versus our novel dataset.

Our theoretical model also affords us perspective on some of the proposed policy answers to the common ownership hypothesis (Posner, Scott Morton and Weyl, 2017). We find that mergers and "break–ups" in the upstream space of insti-

---

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

tutional 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. This suggests that imports and the rise of privately held firms as a fraction of economic activity in the U.S. may dampen the most extreme predictions of the common ownership hypothesis

Work 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 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 IV.A, 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, Li and Shao, 2010; Keswani, Tran and Volpin, 2019). In some sense, the common ownership hypothesis sits at the nexus 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 I 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 II we describe our data sources as well as the advantages of our scraped dataset over the Thompson Reuters s34. In Section III we offer our main descriptive evidence on profit weights from the S&P 500. Section IV 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 V, and Section VI

concludes.

I. 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.

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). We begin with the same two assumptions as described in Backus, Conlon and Sinkinson (2020c).

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. We call an investor a common owner if $\beta_{fs} > 0$ for multiple firms. Assumption 1 is that the profit of the shareholder, v_s, is given by the sum of profits over their portfolio of investments weighted by cash flow rights,

$$(A1) \qquad v_s = \sum_{\forall g} \beta_{gs} \pi_g.$$

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:

$$(A2) \qquad Q_f(x_f, x_{-f}) = \sum_{\forall s} \gamma_{fs} \cdot v_s(x_f, x_{-f})$$

$$= \sum_{\forall s} \gamma_{fs} \cdot \left(\sum_{\forall g} \beta_{gs} \cdot \pi_g(x_f, x_{-f}) \right)$$

$$= \sum_{\forall s} \gamma_{fs}\beta_{fs}\pi_f + \sum_{\forall s} \gamma_{fs} \sum_{\forall f \neq g} \beta_{gs}\pi_g$$

$$\propto \pi_f + \sum_{g \neq f} \underbrace{\left(\frac{\sum_{\forall s} \gamma_{fs}\beta_{gs}}{\sum_{\forall s} \gamma_{fs}\beta_{fs}} \right)}_{\equiv \kappa_{fg}(\gamma_f, \beta)} \pi_g$$

$$= \pi_f + \sum_{g \neq f} \kappa_{fg}(\gamma_f, \beta) \cdot \pi_g.$$

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.[4] 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 $k_{fg} = k_{gf} = 0 \to 1$ (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.[5]

Most objections to the common ownership hypothesis can be mapped back to objections to either (A1) or (A2).[6] However, a model of common ownership must specify the Pareto weight a firm places on each of its shareholders, sometimes

---

[4]We also must assume that the inner product $\langle \beta_s, \gamma_s \rangle > 0$, so that we always divide by a positive number. This is weaker than assuming $(\beta_{fs}, \gamma_{fs}) > 0$ which would rule out short positions or punishing investors.

[5]It 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.

[6]Assumption 2 in particular paves over a number of questions that have been raised anew by the controversy over common ownership: by what mechanisms and in what settings are the interests of ownership represented by management? And what is the role of fiduciary duties, both those of management to owners as well as those of institutional investment managers to clients. See Hemphill and Kahan (2020) for a recent overview, and Anton et al. (2018) which addresses these questions directly in the context of common ownership.

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 U.S. economy. 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 $f(0) = 0$.[7] A convenient choice is $f(\beta_{fs}) \propto (\beta_{fs})^{\alpha}$, which satisfies both.[8] 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. 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 V.

A.   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}:

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

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

---

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

[8]We 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 $\kappa_{fg} = \frac{\langle \gamma_f, \beta_g \rangle}{\langle \gamma_f, \beta_f \rangle} = \frac{\langle a\gamma_f, \beta_g \rangle}{\langle a\gamma_f, \beta_f \rangle}$.

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 that investors hold in f and those that 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 $\cos(\beta_f, \beta_g) \to 1$. This suggests a link between indexing strategies, e.g. investing in the "market portfolio," and common ownership profit weights, which we explore further in our empirical exercise.

Overlapping ownership is what, in general, the literature construes to be "common ownership." 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, Lopez-De-Silanes and Schleifer, 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}.

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

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

B.   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 an undiversified investor. Firms 2 and 3 have the following identical ownership structure: 60 percent of each is held by small, undiversified retail investors. Another 20 percent of each is 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:

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

The calculation follows in the same manner as Example 1. This example high-

Table 2—Example 2 Ownership Structure

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

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

lights 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 IV.A.

II. 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 the 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."[9] 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.[10] We augment this ownership data by scraping the data ourselves from the SEC filings. These data are available from 1999 onward (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.[11] We also gather data on prices and shares outstanding from The Center for Research in Securities Prices (2020) (CRSP).

---

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

[10]Recently, WRDS and some researchers (Ben-David et al. (2018)) noticed data quality issues regarding 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, Conlon and Sinkinson (2019) and in Appendix A, where we document problems that remain after the 2018 update.

[11]A 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 U.S. Securities and Exchange Commission Office of Inspector General (2010).

We use our scraped data on 13(f) holdings from 2000 onward, and the s34 database for filings from 1980–1999. We provide additional details on dataset construction and comparisons of the two databases in Appendix A and the Online Appendix. 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.[12] We also exclude firms where the U.S. 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 in our scraped data.[13]

We document a number of additional discrepancies between our scraped dataset and the Thomson Reuters s34 dataset in Appendix A. Online Appendix Figure A-1 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, Online Appendix Figure A-2 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, Online Appendix Figure A-3 shows that if one were to create Figure 1

---

[12]Occasionally, 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. Similarly, there is additional information on firm cross–holdings in 13(d) and 13(g) reports, which are more difficult to incorporate because they are not filed on a quarterly basis. 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, Schmalz and Tecu (2018) do when they compute the profit weights for airlines and Backus, Conlon and Sinkinson (2020a) do when they compute the profit weights for cereal.

[13]This is one of the documented issues with the s34 database; see Ben-David et al. (2018).
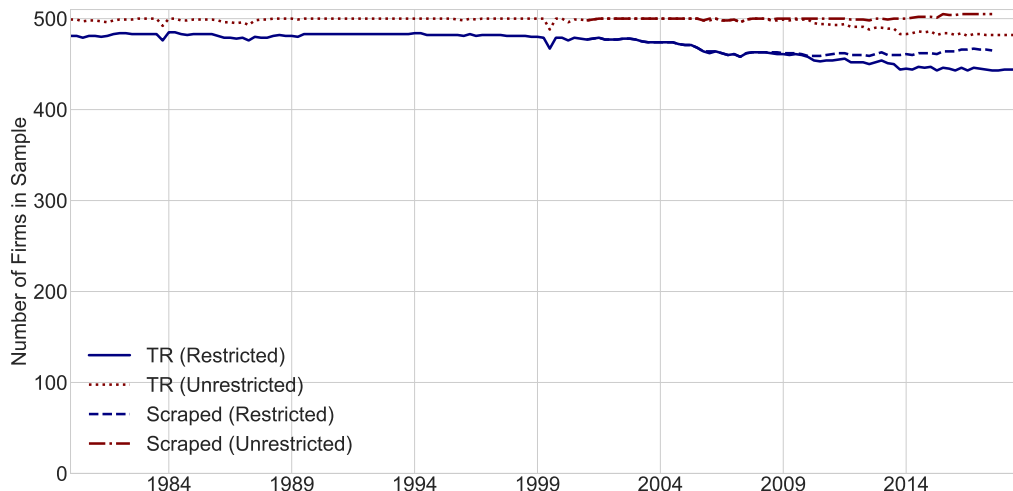
FIGURE 2. NUMBER OF FIRMS IN THE S&P 500 SAMPLE

Note: 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).

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.

III. 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 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.[14] 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,

---

[14]Fichtner, Heemskerk and Garcia-Bernardo (2017) maps the historic rise of the Big Three and raises concerns for their role in corporate governance.

FIGURE 3. NUMBER OF 13(F) MANAGERS HOLDING S&P 500 CONSTITUENTS

Note: This figure depicts the number of managers filing 13(f) reports each quarter over time. For the scraped dataset, a manager is a Central Index Key (CIK). In the Thomson Reuters data, a manager is identified by a "mgrno."


FIGURE 4. SHARE OF S&P 500 OWNED BY 13(F) MANAGERS

Note: 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.

FIGURE 5. SHARE OF TYPICAL FIRM OWNED BY BIG THREE INSTITUTIONAL OWNERS

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

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 III.D 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.

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.[15] We use the S&P 500 as it is designed to reflect the broader U.S. economy; it consists of widely held firms, and many investment funds offer products tied to the constituent firms in one way or another.

A. *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

---

[15]We 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
Raw	68.67%	31.33%
Cross-Section	54.89%	45.11%
Time-Series	67.96%	32.04%
Panel	61.69%	38.31%

Note: This table describes the attribution of variance according to the decomposition in equation (2). The raw sample is unmodified; the cross–section is residualized on quarter fixed effects; the time series is residualized on ordered firm pair fixed effects, and finally the panel case is residualized on both quarter and pair fixed effects.

investor concentration. Taking logs, these sources of variation are additively separable, and so we can attribute the variance due to each component:

$$Var(\log \kappa_{fg}) = Var(\log \cos(\beta_f, \beta_g)) + Var\left(\log \sqrt{\frac{IHHI_g}{IHHI_f}}\right)$$

$$(2) \qquad + 2 \cdot Cov\left(\log \cos(\beta_f, \beta_g), \log \sqrt{\frac{IHHI_g}{IHHI_f}}\right).$$

These are observable objects, and so the decomposition helps us to understand the sources of variation in the common ownership profit weights. The covariance term can be shown to be mechanically zero in our data since for each log relative investor concentration (say, for κ_{fg}), we also observe its inverse (for κ_{gf}). 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 ordered 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 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, representing over 45% of the variation. This is easily reconciled;
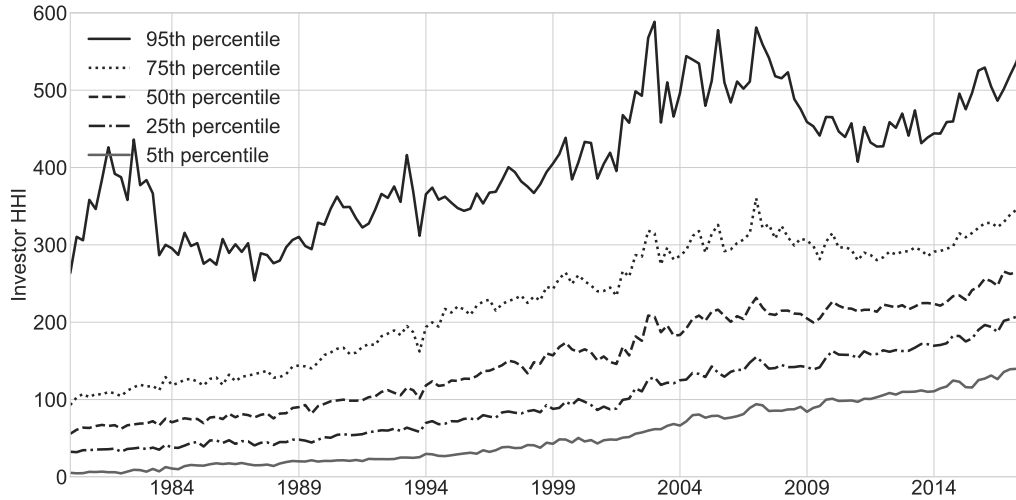
FIGURE 6. INVESTOR β CONCENTRATION (IHHI)

Note: 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 $10{,}000 \times IHHI_f$.

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, that explains the lion's share of the rise of common ownership in the time series.

B.   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.[16]

---

[16]Note 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 guidelines for product markets are appropriate to apply to investors.

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.

C.   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 equation (1). Cosine similarity is an L_2 measure, and it measures how similar the investors' positions in firm f are to those in 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.[17] To be explicit we can write:

$$(3) \qquad L_2(\beta_f, \beta_g) = \cos(\beta_f, \beta_g) = \frac{\sum_s \beta_{fs}\beta_{gs}}{\|\beta_f\| \, \|\beta_g\|}.$$

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 address 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 when investors hold either a position in firm f or in firm g, and thus are not common owners. We construct

---

[17]As 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 $\cos(\beta_f, \beta_g) = 1$

FIGURE 7. COSINE SIMILARITY AMONG INVESTORS

Note: 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.

an L_1 measure of similarity that varies from $[0, 1]$:[18]

$$(4) \qquad L_1(\beta_f, \beta_g) = \frac{1}{2} \sum_s \left(\beta_{fs} + \beta_{gs} - |\beta_{fs} - \beta_{gs}| \right).$$

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 does not 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 $\overline{w}_f = \frac{\sum_s \beta_{fs}}{\sum_{f,s} \beta_{f,s}}$ that represent the market portfolio.[19] We can then compare the normalized portfolio

[18]Absent 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.

[19]Our measure of the "market portfolio" is based on cash flow shares rather than market–cap weights.

FIGURE 8. SIMILARITY BETWEEN INVESTOR PORTFOLIOS AND S&P 500 INDEX

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

weights $w_{fs} = \frac{\beta_{fs}}{\sum_f \beta_{fs}}$ and measure the similarity of each investor's portfolio to the market portfolio: $L_1(w_s, \overline{w})$ and $L_2(w_s, \overline{w})$. This is consistent with the literature in that the *active share* measure of Cremers and Petajisto (2009) is given by $1 - L_1(w_s, \overline{w})$.[20]

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" $\overline{w}$ among S&P 500 securities in our dataset. In Figure 8, we report the weighted average of these similarity measures, where we weight each investor by assets under management (AUM). As one might expect, at least on an asset weighted basis, investor portfolios become much more similar to 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

---

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 that we ignore all non S&P 500 securities from our calculation of portfolio weights.

[20]However our analysis is at the investor/manager level from 13(f) filings not at the level of an individual fund.

particularly among the Big Three.

D.   Correlates of Profit Weights

Next we ran a series of regressions of κ_{fgt} on potential covariates. In each, we include quarter and pairwise fixed effects, where the pairs are ordered (i.e., a different fixed effect 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 indices and the corresponding increase in overlapping ownership. We also observe a robust negative correlation between quarterly operating margins at firm f and κ_{fg}.[21] We caution that this is essentially uninterpretable, as there are reasons to believe that causality might go both ways. Common ownership effects might permit firms to obtain higher margins; alternatively, institutional investors might be attracted to firms that for other reasons obtain high margins.

In model (1) we include the sum of β_{fs} for shareholders BlackRock, Vanguard, and State Street, and we find a strong positive correlation. Next, in models (2) – (4) we add our measure of investor indexing. This measure aggregates the investor–level L_2 indexing measure of Section III.C above to the firm level according to weights $\beta_{fs} / \sum_f \beta_{fs}$. We find, in model (2), a strong relationship between our firm–level measure of indexing behavior and common ownership profit weights. 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, two of three of the correlations are 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.

IV.   Economic Implications

### A.   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

---

[21]We measure quarterly operating margins using data from Compustat as the ratio (Sales - Cost of Goods Sold)/Sales.

TABLE 4—CORRELATIONS WITH κ

	(1)	(2)	(3)	(4)
Retail Share	0.8555*	0.6927*	0.6699*	0.6764*
	(0.0006)	(0.0006)	(0.0006)	(0.0006)
Log(Market Cap)	0.0797*	0.0719*	0.0705*	0.0721*
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Operating Margin	-0.0044*	-0.0047*	-0.0045*	-0.0046*
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Big Three Holdings	0.6498*		-0.2357*	
	(0.0022)		(0.0024)	
Investor Indexing		1.0427*	1.1094*	1.0993*
		(0.0013)	(0.0014)	(0.0014)
BlackRock Holdings				-0.3246*
				(0.0032)
Vanguard Holdings				0.4469*
				(0.0061)
State Street Holdings				-0.4570*
				(0.0051)
R^2	0.7079	0.7204	0.7206	0.7209
N	13,230,003	13,230,003	13,230,003	13,230,003

Note: This table reports correlates of the common ownership profit weights. An observation is a directional pair of S&P 500 constituent firms in a given quarter. All specifications include quarter and directional firm-pair fixed effects. Robust standard errors are reported in parentheses. * indicates significance at the 1% 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.

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 law countries, and other work has found evidence in the developing world (Bertrand, Mehta and Mullainathan, 2002). However, tunneling is not typically believed to occur in the U.S. for two reasons: strong investor protections that facilitate healthy financial markets (Porta, Lopez-De-Silanes and Schleifer, 1999) and the near–universal absence of a controlling interest in publicly–traded firms, as the U.S. 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 on 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 number 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, Mehta and Mullainathan (2002) offer econometric evidence of tunneling in 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 that 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 III, in the later periods, the result depends heavily on our assumptions about γ.[22]

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 in some form of tunneling behavior, is implausibly strong. It is possible that this is held in check by strong minority shareholder protections. Perhaps more likely, and in contrast to mergers, where shareholder activity is direct and measurable, these incentives may be incompletely transmitted from owner to institutional manager, and from institutional manager down the corporate chain to actors in the firm.

B. Quantifying the Common Ownership Channel

De Loecker, Eeckhout and Unger (2020) 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

---

[22]We document this in Online Appendix Section 2, where putting more weight on large investors actually results in higher tunneling incentives.

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:

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

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 $\widehat{p}(\kappa)$. 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 De Loecker, Eeckhout and Unger (2020) do over the time period 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 and the same: $\hat{p}(\kappa) = \mu$. Next we choose the number of firms $J = 8$ so that our $HHI \approx 1250$ to match Grullon, Larkin and Michaely (2018).[23] Finally, we calibrate a and b for 1980. We construct a markup of $\mu = 1.21$ to match De Loecker, Eeckhout and Unger (2020) and an average own–elasticity of -7.21 in line with the range of elasticities reported in Eaton and Kortum (2002).[24] This all but eliminates the outside good share.[25] We also consider an alternative scenario where one of the firms instead prices as if it were held by an entirely undiversified owner, or was privately held. We denote this firm as a "maverick."

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 De Loecker, Eeckhout and Unger (2020) for the same period.[26] Moreover, a Granger test rejects the null that the lagged simulated common ownership markups are not predictive of the De Loecker, Eeckhout and Unger (2020) markups.[27] Despite this, we take the time series in Figure 10 as clear evidence

---

[23]We can obtain nearly identical results varying the number of firms from 5–15.

[24]Simonovska 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.

[25]Alternatively, 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 (Conlon and Gortmaker, Forthcoming).

[26]An important aspect of the results in De Loecker, Eeckhout and Unger (2020) 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.

[27]The test (Granger, 1969) is based on a VAR in first differences with two lags and a time trend and

Figure 10. Simulated Markups: 1980–2017

Note: This figure presents predicted markups, defined as p/c to align with De Loecker, Eeckhout and Unger (2020), for the 8–firm calibration exercise described in the text. "Maverick" refers to a scenario where one firm is privately held, and therefore has $\kappa = 0$. See the text for exact specification.

that the bulk of the rise in markups described in De Loecker, Eeckhout and Unger (2020) 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. It is notable that turning one single firm into a "maverick" greatly disciplines the pricing effect of common ownership in this simplified setting. This may provide a testable implication for those studying common ownership effects on prices.

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 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 mar-

rejects with a critical p value of 0.05.

FIGURE 11. SIMULATED PROFITS: 1980–2017

Note: 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.

ket prices aggressively, then the resulting markups (and profits) are much lower. The dashed red line in Figure 11 depicts profits when one firm is a "maverick." 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 equation (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, Conlon and Sinkinson (2020a). 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.

C.   *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, Scott Morton and Weyl, 2017). Here we consider a simple exercise where we either: (a) allow BlackRock and Vanguard to merge; (b) 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;[28] or (c) 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 begins to be more substantial differences. Under our baseline scenario of proportional control and the observed ownership structure $\overline{\kappa} \approx 0.7$, the merger would increase this to $\overline{\kappa} \approx 0.8$, while breaking them up would decrease this to $\overline{\kappa} \approx 0.62$. 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 $\overline{\kappa} \approx 0.46$ in 2017, and it implies that average profit weights are essentially unchanged since 2000.

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 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 $\overline{\kappa}$ 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 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 Online Appendix 3, where Online Appendix Figure A-8 shows the impact of removing those two firms from our measures of indexing developed in Section III.C. More disagreement among the remaining investors tends to lead to lower profit weights overall. We can think of this scenario

---

[28]We do not split holdings based on overlapping industries (one of the suggestions in Posner, Scott Morton and Weyl (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 $\overline{\kappa}$ 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 a single firm with combined holdings. Ignoring BlackRock and Vanguard sets $\gamma_{f,s} = 0$ for those investors and implicitly includes them in retail share.

as similar to the "put the shares in a drawer" proposal of Posner, Scott Morton and Weyl (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 curbing common ownership incentives, this proposal might have unintended consequences in reducing the effectiveness of other corporate governance actions.

V. Robustness

A. 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

FIGURE 13. PROFIT WEIGHTS UNDER DIFFERENT CONTROL ASSUMPTIONS

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

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 $\frac{1}{CHHI_f}$.[29]

In Figure 14, we report our concentration measures for effective control which we

---

[29]Unlike in our calculation of κ where we can multiply $\gamma_{\cdot s}$ 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.
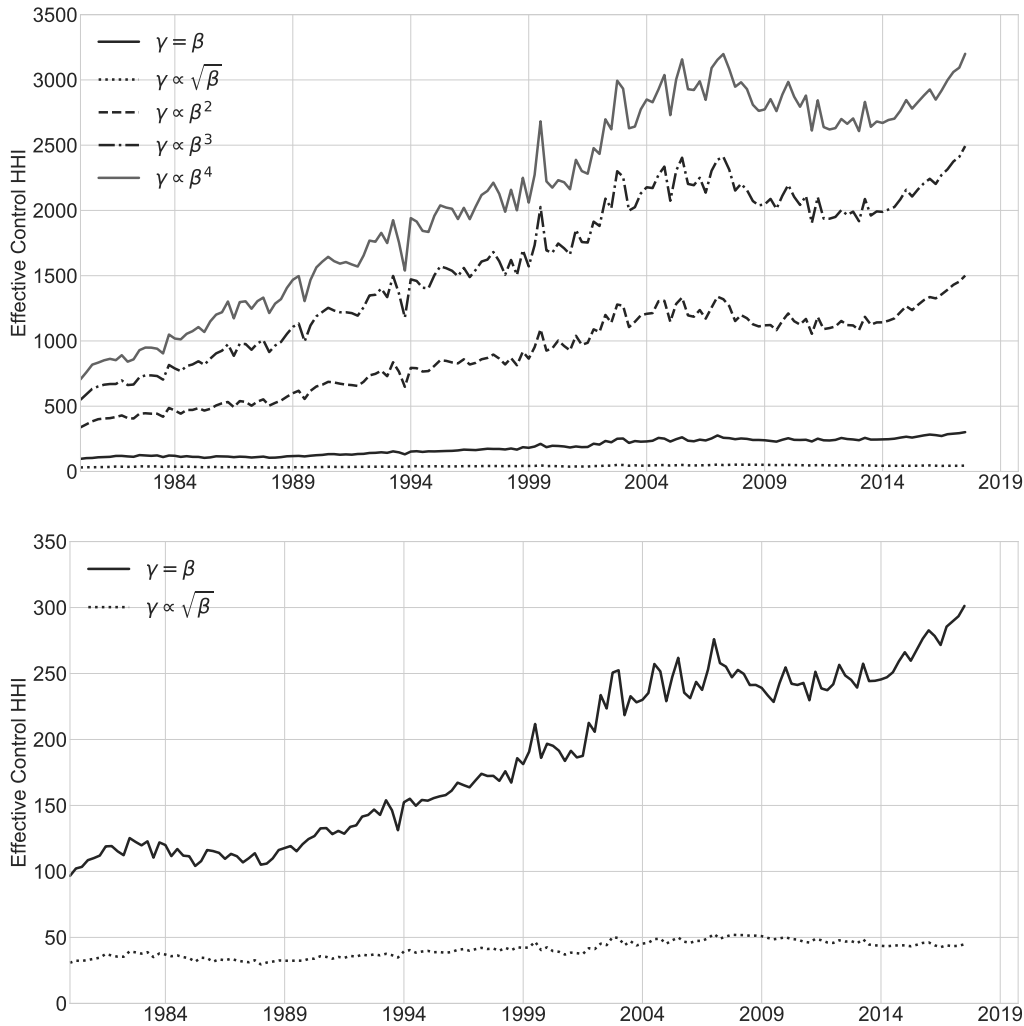
FIGURE 14. CONTROL WEIGHTS γ CONCENTRATION (CHHI)

Note: 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 $\gamma = \beta$.

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

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 $\gamma_{fs} = \beta_{fs}$.

suggests we have considered the range of relevant values for α.

B.   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.

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 Standard & Poor's (2020). 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 set of specific industries: commercial banks, as defined by SIC code 6021 (National Commercial Banks) in Compustat that are also S&P 500 constituents; airlines, using a hand–collected sample of 27 nationwide airline securities; and ready–to–eat (RTE) cereal. The airline sample required extensive data cleaning due to the many bankruptcies and mergers over the time frame. Details are in Online Appendix 2. Results are

FIGURE 16. PROFIT WEIGHTS AMONG AIRLINES, BANKS, AND RTE CEREAL

Note: This figure presents average pairwise profit weights for banks and airlines. Banks are for SIC code 6021 from S&P 500 sample only. Airlines are separately scraped from 13(f) filings and are available only for 2000–2017. Cereal is Kellogg's, General Mills, Quaker Oats (a unit of PepsiCo) and Post (or its controlling entity). Profit weight κ_{fg} computed under proportional control $\gamma_{fs} = \beta_{fs}$.

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.

C. 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 to "sole" voting authority shares. We use Thomson Reuters data prior to 2010, and then our scraped sample 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 measures appear to be
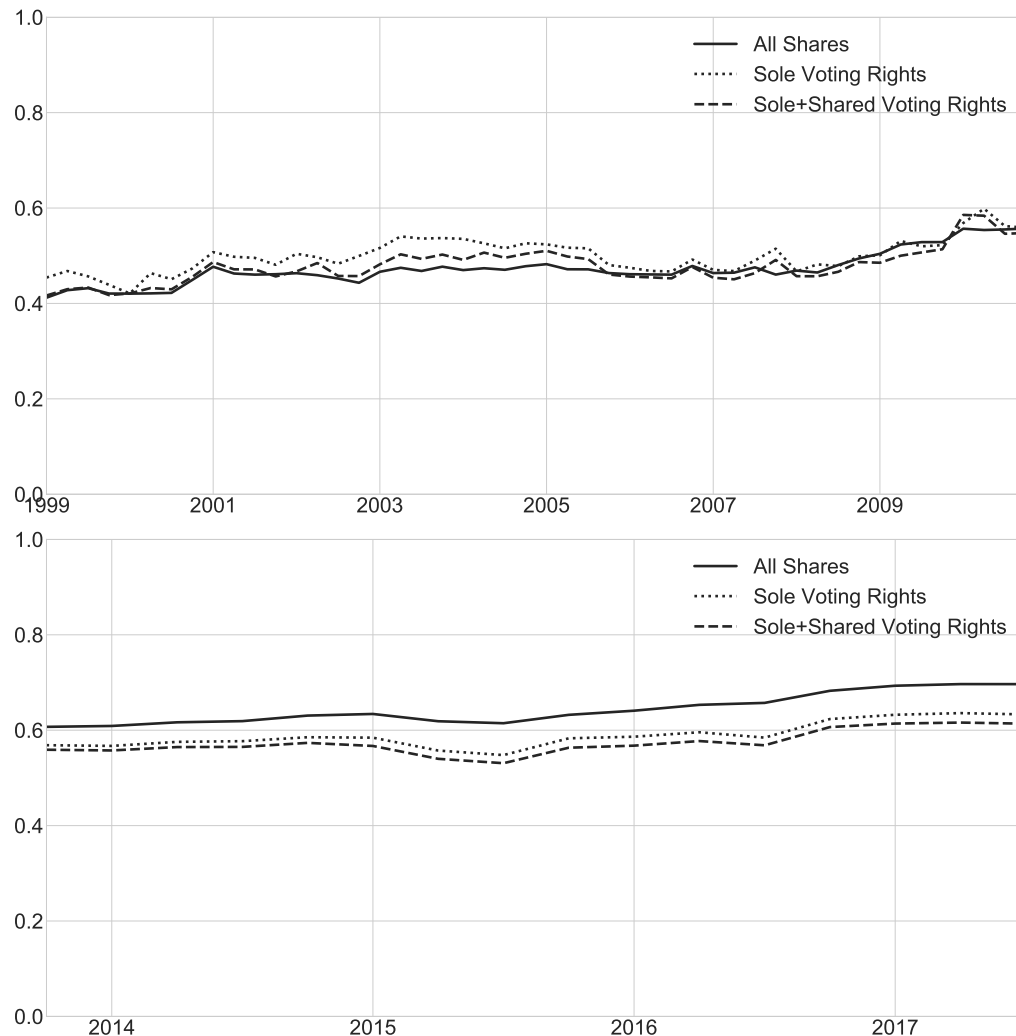
FIGURE 17. ALTERNATIVE κ BY VOTING AUTHORITY

Note: This reports robustness checks where we compare the measure we report in our main results *All Shares* (solid–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.

miniscule.

VI.   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, one 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 De Loecker, Eeckhout and Unger (2020) 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, contrary to evidence from Kwoka, Greenfield and Gu (2015), that the standard analyses are systematically over–predicting the price effects of mergers. Likewise, it suggests that purely financial transactions that 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 the effects of investor concentration appear 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, and 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, Conlon and Sinkinson, 2020a), 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 researchers 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, Miguel, Florian Ederer, Mireia Gine, and Martin C. Schmalz. 2018. "Common Ownership, Competition, and Top Management Incentives." Ross School of Business Paper No. 1328.

Azar, Jose, Martin C. Schmalz, and Isabel Tecu. 2018. "Anti-Competitive Effects of Common Ownership." *Journal of Finance* 73(4).

Azar, Jose, Sahil Raina, and Martin C. Schmalz. 2016. "Ultimate Ownership and Bank Competition." https://ssrn.com/abstract=2710252. Accessed on 2020-01-08.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2019. "The Common Ownership Hypothesis: Theory and Evidence." Brookings Working Paper.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2020*a*. "Common Ownership and Competition in the Ready-To-Eat Cereal Industry." Unpublished.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2020*b*. "Scraped Common Ownership Data." *Available online*, https://sites.google.com/view/msinkinson/research/common-ownership-data.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2020*c*. "Theory and Measurement of Common Ownership." *AEA: Papers and Proceedings* 110.

Ben-David, Itzhak, Francesco Franzoni, Rabih Moussawi, and John Sedunov. 2018. "The Granular Nature of Large Institutional Investors." Swiss Finance Institute Research Paper Series.

Berle, Adolf, and Gardiner Means. 1932. *The Modern Corporation and Private Property.* MacMillan.

Bertrand, Marianne, Paras Mehta, and Sendhil Mullainathan. 2002. "Ferreting out Tunneling: An Application to Indian Business Groups." *Quarterly Journal of Economics* 117(1]).

Boller, Lysle, and Fiona Scott Morton. 2019. "Testing the Theory of Common Stock Ownership." Working Paper.

Bresnahan, Timothy F. 1987. "Competition and Collusion in the American Automobile Industry: The 1955 Price War." *The Journal of Industrial Economics* 35(4): 457–482.

Bresnahan, Timothy F., and Steven C. Salop. 1986. "Quantifying the Competitive Effects of Production Joint Ventures." *International Journal of Industrial Organizatoin* 4: 155–175.

Conlon, Christopher, and Jeff Gortmaker. Forthcoming. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *Rand Journal of Economics.*

Cremers, K. J. Martijn, and Antti Petajisto. 2009. "How Active Is Your Fund Manager? A New Measure That Predicts Performance." *The Review of Financial Studies* 22(9): 3329–3365.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics* 135.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70(5): 1741–1779.

Fichtner, Jan, Eelke M. Heemskerk, and Javier Garcia-Bernardo. 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, Milton. 1953. "Essays in Positive Economics." Chapter The Methodology of Positive Economics. University of Chicago Press. (unfinished citation).

Gilje, Erik P., Todd A. Gormley, and Doron Levit. 2020. "Who's paying attention? Measuring Common Ownership and its Impact on Managerial Incentives." *Journal of Financial Economics.*

Gramlich, Jacob, and Serafin Grundl. 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, Gustavo, Yelena Larkin, and Roni Michaely. 2018. "Are U.S. Industries Becoming More Concentrated?" forthcoming, *Review of Finance.*

Gutiérrez, Germán, and Thomas Philippon. 2016. "Investment-less Growth: An Empirical Investigation." *Brookings Papers on Economic Activity.*

Heim, Sven, Kai Hüschelrath, Ulrich Laitenberger, and Yossi Spiegel. 2019. "Anticompetitive Effect of Minority Share Acquisitions: Evidence from the Introduction of National Leniency Programs." Working Paper.

Hemphill, Scott, and Marcel Kahan. 2020. "The Strategies of Anticompetitive Common Ownership." *Yale Law Journal* 129.

Jiang, Wei, Kai Li, and Pei Shao. 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, Simon, Rafael La Porta, Florencio Lopez-De-Silanes, and Andrei Schleifer. 2000. "Tunneling." *American Economic Review: Papers and Proceedings* 90(2): 22–27.

Kennedy, Pauline, Daniel P. O'Brien, Minjae Song, and Keith Waehrer. 2017. "The Competitive Effects of Common Ownership: Economic Foundations and Empirical Evidence." Working Paper.

Keswani, Aneel, Ahn Tran, and Paolo Volpin. 2019. "Institutional Debt Holdings and Governance." Working Paper.

Kwoka, John, Daniel Greenfield, and Chengyan Gu. 2015. *Mergers, Merger Control, and Remedies: A Retrospective Analysis of U.S. Policy.* Mit Press.

Matvos, Gregor, and Michael Ostrovsky. 2008. "Cross-Ownership, Returns, and Voting in Mergers." *Journal of Financial Economics* 89: 381–403.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69(2): 307–342.

O'Brien, Daniel, and Steven Salop. 2000. "Competitive Effects of Partial Ownership: Financial Interest and Corporate Control." *Antitrust Law Journal* 67(3): 559–614.

Porta, Rafael La, Florencio Lopez-De-Silanes, and Andrei Schleifer. 1999. "Corporate Ownership Around the World." *The Journal of Finance* 54(2): 471–517.

Posner, Eric A., Fiona M. Scott Morton, and E. Glen Weyl. 2017. "A Proposal to Limit the Anti-Competitive Power of Institutional Investors." *Antitrust Law Journal* 81(3).

Reynolds, Robert J., and Bruce R. Snapp. 1986. "The competitive effects of partial equity interests and joint ventures." *International Journal of Industrial Organization* 4(2): 141–153.

Rotemberg, Julio. 1984. "Financial Transactions Costs and Industrial Performance." http://hdl.handle.net/1721.1/47993. Accessed on 2020-01-08.

Shambaugh, Jay, Ryan Nunn, Audrey Breitwieser, and Patrick Liu. 2018. "The State of Competition and Dynamism: Facts about Concentration, Start-Ups, and Related Policies." The Hamilton Project: Economics Facts.

Simonovska, Ina, and Michael E. Waugh. 2014. "The elasticity of trade: Estimates and evidence." *Journal of International Economics* 92(1): 34–50.

Standard & Poor's. 2020. "Compustat." *Available through Wharton Research Data Services.*

The Center for Research in Securities Prices. 2020. "CRSP Database." *University of Chicago*, Available through Wharton Research Data Services.

Thomson Reuters. 2020. "s34 Database." *Available through Wharton Research Data Services.*

U.S. Securities and Exchange Commission Office of Inspector General. 2010. "Review of the SEC's Section 13(f) Reporting Requirements." Report No. 480.

DATA APPENDIX

A1. 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 in Appendix A.A2.

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.

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 we 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 these 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.

A2. 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 is 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. The data and code are available at the following website as of the date of publication: `https://sites.google.com/view/msinkinson/research/common-ownership-data`.

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.

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 that have 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 "war" (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).

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, while Noble Corp had a just over 250M shares outstanding and a market capitalization of $5.6B.[30] The result is a dataset of 18,968,596 observations of unique CIK–CUSIP–record date holdings across 75 reporting quarters.

---

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