

January 2019

The Common Ownership Hypothesis: Theory and Evidence

Matthew Backus

Columbia University

Christopher Conlon

New York University

Michael Sinkinson

Yale University

This report is available online at: <https://www.brookings.edu>

B | Economic Studies
at BROOKINGS

The Brookings Economic Studies program analyzes current and emerging economic issues facing the United States and the world, focusing on ideas to achieve broad-based economic growth, a strong labor market, sound fiscal and monetary policy, and economic opportunity and social mobility. The research aims to increase understanding of how the economy works and what can be done to make it work better.

STATEMENT OF INDEPENDENCE

The authors did not receive any financial support from any firm or person for this article or from any firm or person with a financial or political interest in this article. They are currently not an officer, director, or board member of any organization with an interest in this article.

The Common Ownership Hypothesis: Theory and Evidence

Matthew Backus, Christopher Conlon, Michael Sinkinson

January 2019

1 Introduction

This paper surveys the recent literature that examines the relationship between ownership of firms in the financial space, and the strategic decisions made by firms in product markets. Of primary concern is the *common ownership hypothesis*, which suggests that when large investors own shares in many firms within the same industry, those firms have an incentive to soften competition by producing fewer units, raising prices, reducing investment, innovating less, or limiting entry into new markets. The core of the idea is quite simple: firms maximize shareholder value, but shareholders hold stakes in competitors; thus, firms may want to maximize some combination of their own profits and their competitors' profits to maximize the value of their investors' portfolios. The implications of this possibility are enormous: if firms place positive weight on rivals' profits when making strategic decisions, the entire economy of publicly-traded firms may have incentives to soften competition, resulting in significant harm to consumers.

This hypothesis is an old idea, with a history dating back to the 1980s.¹ What has renewed interest recently is increasing concentration among investment managers and new instruments for investing in diversified portfolios. The introduction of 401(k) defined-contribution retirement plans in 1978 led to a rise in diversified portfolios and mutual fund managers in the 1980s (such as Fidelity); later index funds (such as Vanguard) in the 1990s; and finally exchange-traded funds (such as those offered by Blackrock and State Street), or ETFs, in the 2000s. At the beginning of 2018, the four largest asset managers (Blackrock, Vanguard, State Street, and Fidelity) managed over \$16 trillion in assets, and for 88% of firms on the S&P 500 Index, the largest shareholder was one of those four asset managers. While index fund and ETF managers adopt “passive” investment strategies, evidence suggests that they are actively engaged in corporate governance (Fink (2018), Brav et al. (2018)).

This creates a fundamental tension for households that both consume goods and invest in diversified retirement funds. As consumers, they might be harmed if firms don't compete as fiercely as possible in markets for goods and services. As investors, they benefit from low-cost diversified investment products such as index funds or ETFs. A large number of households don't hold any investments, and so do not receive the benefit of the latter, but pay the cost of the former. This has led to an intense debate about whether large institutional investors have caused indirect harm to consumers, and what, if anything, policymakers should do about it. This also relates to the larger debate among economists about rising firm markups (De Loecker

¹See Rotemberg (1984); Bresnahan and Salop (1986).

and Eekhout, 2017) and concentration; and declining dynamism, investment (Gutiérrez and Philippon, 2016), and productivity growth. The important, unresolved question is why markups and concentration are rising while dynamism and productivity growth have been declining, and some researchers have pointed to common ownership effects as a potential explanation (Shambaugh et al., 2018).²

We formulate the common ownership hypothesis in terms of what we call *profit weights*. Put simply, under the conventional model of profit maximization, when making any strategic decision, a firm should only care about the impact on its own profits, which implies a profit weight of zero on its competitors. If the common ownership hypothesis is correct, then firms may also take into account the effect of their choices on their competitors' profit levels, due to presence of common owners with stakes in each firm. Our theory section discusses the measures researchers have proposed for these profit weights under the common ownership hypothesis. In general, these weights tend to depend on two inputs: the shareholdings of each investor in both firms, and the weight that the firm places on each investor. With some important caveats, the shareholdings are generally observable from the 13(f) financial filings of large institutional investors that are required by the Securities and Exchange Commission (SEC). The weight that a firm places on each investor is not directly observable by the researcher, nor is there a definitive theory of corporate governance which tells researchers how to choose these weights. The literature often assumes that control weights are proportional to the cash-flow rights of the investor.³

After describing the available data sources for studying the common ownership hypothesis, we will show descriptive data that are consistent with a large increase in the holdings of diversified funds over time – across the economy as a whole but also within some specific industries – as well as substantial concentration among asset managers. For a sense of scale, under the conventional assumption of profit maximization, firms would place zero weight on a dollar of another firm's profits. Alternatively, when we do merger simulations (or equivalently, in the case of perfect collusion), we posit that a firm values a dollar of profits at another firm as equal to a dollar of our own. Under the common ownership hypothesis, we show that in 1980 an average S&P 500 firm would have valued a dollar of profits to another randomly chosen S&P 500 component firm at 20 cents. By the end of 2017, this more than tripled to approximately 70 cents. This staggering rise would, if common ownership incentives translate to firm behavior, give firms a unilateral incentive to raise prices, even in the absence of collusion (which would be illegal). However, quantifying such effects would require us to also know demand and to isolate common ownership incentives among firms within a particular market. On both points we borrow from Backus et al. (2018b). To the former, we consider a simple calibration of a Bertrand pricing model with Logit demand, and show that the price effects are not dissimilar in scale to those found by Eekhout et al. (2018). On the latter point, looking at three specific product markets that have proven popular among those studying this question – banks, airlines, and ready-to-eat (RTE) cereal – we see similar or even larger effects for firms within the same industry.

Next we examine the growing empirical literature related to common ownership.⁴ A series of early empirical papers (Azar et al. (2016), Azar et al. (2018a)) drew much attention but also criticism, and spurred a large amount of new research into the topic. While the authors of these early papers deserve credit for

²The relationship between competitiveness and productivity is a robust result in industrial organization (IO), though interpretation of mechanisms remains controversial; see Syverson (2004), Backus (2014), and Asker et al. (2014) for recent treatments and discussions of the literature.

³See Backus et al. (2018b) for an extensive discussion on this point and consideration of alternative specifications. We find that these alternatives do not qualitatively change our findings concerning the rise and patterns of common ownership.

⁴Schmalz (2018) also reviews this growing literature.

shining a spotlight on this issue, the methods used in the early papers – regressions of prices on opaque and theoretically problematic measures of ownership concentration – are unreliable in many contexts. In particular, for differentiated products markets, firms that are pricing competitively could spuriously generate correlations between prices and measures of ownership concentration. Regressions therefore may simply be measuring correlations as opposed to true causal effects. This is a familiar concern for the empirical industrial organization literature, which for decades reported cross-industry regressions with equilibrium outcomes as dependent variables and measures of concentration, also an equilibrium outcome, as “independent” regressors. Criticism of these methods grew in the 1980s as researchers pointed out both theoretical problems interpreting concentration measures, as well as the difficulty of finding statistically clean variation (see [Schmalensee \(1989\)](#) for a thorough discussion). Today, such methods are viewed with much skepticism in the industrial organization literature. So, while we view the initial evidence as provocative and important, we do not believe that it offers a framework for the growing literature, neither in terms of rigorous testing of conduct nor policy analysis.

Subsequently, a number of articles have presented empirical results which relate common ownership incentives to various other market outcomes, such as investment and entry. We briefly discuss these papers as well.

Beyond basic facts, another question raised by the literature is precisely how increases in common ownership might lead to softer competition in the product market. A number of mechanisms have been proposed, with much attention being paid to executive compensation (see [Antón et al. \(2018a\)](#)). Here, we might imagine that as an industry has more common ownership, owners want to see weaker managerial incentives; in industries with less common ownership, owners would favor steeper incentive structures. Other potential mechanisms involve direct communication between investors and managers on capacity ([Aryal et al. \(2018\)](#)) or investment decisions. Another potential mechanism is through corporate governance actions. Common owners may be more likely to side with management against activist investors, or to oppose management strategies which lead to stronger competition within the industry. Large common owners may also seek to shape the agenda for voting at annual meetings.

To summarize, we find that while the existing correlations explored in this growing literature are provocative and important, the methods and measures used to date make it difficult to draw clear conclusions. We believe that these early contributions are the beginning of a literature rather than the end, and that before policy measures can be debated we need first to test these hypotheses in more settings, using modern methods; and second, to better understand the mechanisms by which such effects are generated. Our criticisms of this literature mirror those made by empirical IO economists of the literature on prices and concentration indices. However, this also guides our thinking for new directions – just as merger analysis moved from concentration indices to careful modeling of demand and pricing, these same tools can be applied to the analysis of common ownership. In [Backus et al. \(2018a\)](#) we take a step in this direction and offer new tools for testing conduct, however we believe that there is room for more work in this area.

The paper is organized as follows: Section 2 discusses the formal theory underlying the common ownership hypothesis. Section 3 discusses data sources, data challenges, and shows trends in common ownership as well as implied profit weights. Section 4 discusses the empirical findings in this literature as well as many responses and criticisms. Section 5 concludes.

2 Theory

The theory of common ownership takes seriously an argument that economists often make informally. Why do firms maximize profits? Because, it goes, they answer to their investors. Therefore, continues the common ownership argument, we need a model of investors' interests. The theory of common ownership posits that firms seek to maximize the value of investors' portfolios. This raises two immediate problems. First, investors may own shares not only in my firm but also in my competitors. Second, each investor's portfolio may be vastly different from the next, and investors may disagree about firm objectives. The theoretical literature on common ownership presents a simple model where firms maximize a weighted average of investor *portfolio* profits.⁵

Suppose there are many investors. We index them by s . Their portfolios consist of shares of many firms, which we'll index by f . For notation, let β_{sf} be investor s 's ownership share of firm f . That ownership share entitles them to a fraction of the profits of firm f , henceforth π_f , which means we can write the value of their portfolio as:

$$\sum_f \beta_{sf} \pi_f.$$

In order to resolve disputes when portfolios disagree, we assume that the firm places weight $\gamma_{sf} > 0$ on the cash flow received by investor s . We interpret these as Pareto weights in the firm's choice problem, i.e. the weight that the firm places on a marginal dollar of each investor. These weights stand in for a model of corporate governance and control rights, and while it is tempting to interpret them as the outcome of an explicitly specified voting game, to our knowledge there is no rigorous microfoundation in the literature.⁶

With all of this in hand, we can now write the objective function of the firm as maximizing:

$$\sum_s \gamma_{sf} \sum_g \beta_{sg} \pi_g.$$

After some simple algebra, we can re-arrange things so that the firm maximizes:

$$\pi_f + \sum_g \kappa_{fg} \pi_g \quad \text{with } \kappa_{fg} \equiv \frac{\sum_s \gamma_{sf} \beta_{sg}}{\sum_s \gamma_{sf} \beta_{sf}}. \quad (1)$$

This objective function is rather different than the classical model, in which firm f maximizes π_f alone. A useful and intuitive way to think about the difference is in terms of the common ownership *profit weights*, denoted by κ_{fg} , which represent the value to firm f of a dollar of profit generated for a competing firm g . These profit weight objects are the channel through which common ownership of stock by large investors affects firm behavior. To anchor interpretation in familiar terms, if $\kappa_{fg} = 0$ whenever $f \neq g$, then we have standard own-firm profit maximization. If $\kappa_{fg} = 1$ for all f and g in the market, then the model is equivalent to market monopoly pricing (or, observationally, perfect collusion). In general, the theory of common ownership yields intermediate values. However, as we will show in a moment, there is no reason that these objects need be constrained between zero and one. Let us consider a few examples borrowed from

⁵This model of firm behavior traces to [Rotemberg \(1984\)](#), and appears implicitly in the cross-ownership profit weights of [Bresnahan and Salop \(1986\)](#). In what follows we adopt the (γ, β) notation of [O'Brien and Salop \(2000\)](#), which studied the implied profit weights in the context of Cournot and Bertrand competition; see their Appendix C.

⁶This is perhaps not surprising: in order to model the voting game we would need to know the payoffs to the players in the downstream product market, a point on which we have been so far agnostic.

Table 1: Example 1 Ownership Structure

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

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

Backus et al. (2018b).

Example 1: Suppose there are three firms and four institutional investors. Firm 1 is controlled by an undiversified owner. Firms 2 and 3 have the following, identical structure: 60% held by a continuum of retail investors; 20% held by two undiversified institutional investors; and 20% held by a single common, diversified investor. This ownership structure is depicted in Table 1.

If we assume $\gamma_{sf} = \beta_{sf}$, i.e. proportional control, then profit weights can be computed directly from ownership data. Applying equation (1), we obtain the following profit weights: Firm 1, with independent ownership, maximizes its own profit alone, with weight $\kappa_{12} = \kappa_{13} = 0$. Similarly, firms 2 and 3, sharing no common owners with firm 1, place zero weight on its profits, so $\kappa_{21} = \kappa_{31} = 0$. However, because firms 2 and 3 have a large common owner, they place substantial weight on each other’s profits: $\kappa_{23} = \kappa_{32} = 1/2$. This implies that these two firms value a dollar of profit at the other at 50 cents to a dollar of their own. It does not imply conspiracy or collusion – but rather, that the firms will naturally internalize the effect of their decisions on each other, in a way that emulates collusive outcomes. We have not specified the strategic interactions between the three firms, so there are no welfare implications at this point in the analysis, but the implications for competition are grim.

A remarkable feature of the example is how large the profit weight on a competitor is given the 20% stake of the common owner. Part of this is coming from the 60% retail share, which is the portion of shares held by retail investors. Individual retail investors are assumed to be infinitesimally small in the model, and so when we multiply γ and β in equation (1), their contribution drops out, effectively magnifying the influence of the institutional investors.⁷ This mathematical artifice is perhaps not without intuition, given the low participation rates of retail investors in corporate governance, but it is essential to the conclusions. Indeed, it can have strong implications, as we illustrate in the next, more stylized example.

Example 2: Let there be two firms and ten investors. Each of the ten investors holds a 1% stake in firm 1, and so the retail share of firm 1 is 90%. Each of the ten investors also holds a $x\%$ stake in firm 2 (so $x \leq 10\%$), and so the retail share of firm 2 is $(100 - 10x)\%$. This ownership pattern is summarized in Table 2.

⁷That retail investors drop out may not be an unrealistic feature of the model. For example, a recent report by ProxyPulse examined 3,379 proxy meetings in the first half of 2017 and found that while institutional investors own 70% of shares and vote them 91% of the time, retail investors holding the remaining 30% voted those shares only 29% of the time. ”2017 Proxy Season Review”, ProxyPulse. However, as we discuss later, it is important to consider large individual investors, e.g. insiders and founding families, when identifying the retail share, since it would be incorrect to treat them as atomistic.

Table 2: Example 2 Ownership Structure

	Firm 1	Firm 2
Investor 1	1%	$x\%$
Investor 2	1%	$x\%$
\vdots	\vdots	\vdots
Investor 10	1%	$x\%$
Retail Share	90%	$(100-10x)\%$

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

In such a world, by applying equation (1), we obtain that firm 1 places a profit weight of x on firm 2, and firm 2 places a profit weight of $1/x$ on firm 1. So, if $x = 3$, where the ten investors also each own 3% of firm 2, then firm 1 *prefers* firm 2 profits to its own by a factor of three. What would this mean in practice? It is difficult to fathom the market outcome where one competing firm makes decisions placing a significantly higher weight on rival profits than on own profits. If a competitor were as efficient as myself, then in a pricing game I should set my prices so high as to divert almost all sales to that firm whenever our products are substitutes.

Though the example is designed to be striking, it is not without precedent. The literature on “tunneling” has shown how the divorce of control rights and cash-flow rights create perverse behavior, in which an owner will wish to transfer assets and cash flows from one firm (the one in which the owner has low cash-flow rights) to another (the one in which the owner has high cash-flow rights), defrauding minority investors in the former firm (Porta et al., 1999; Johnson et al., 2000; Bertrand et al., 2002). However, there are major differences: first, in the tunneling literature, the divorce of control rights and cash-flow rights comes from institutional structures (e.g., “golden shares”). Instead, in the theory of common ownership, this divorce is coming through investor concentration. If control rights (γ) are an increasing function of ownership share (β) as they are in our specification, which assumes $\gamma = \beta$, then an investor with a 10% ownership share has more influence than two investors with 5% each. And, if $\gamma \rightarrow 0$ as $\beta \rightarrow 0$, then retail investors have no influence. This dilution of control rights for small and fragmented owners affords out-sized influence for the largest investors.⁸ Second, however, though tunneling is possible for extremal values of $\kappa > 1$, under common ownership it is not generally the case that minority investors are harmed. Indeed, if profit weights are roughly symmetric and approaching 1, then the theory, as we outline it next, predicts that prices will approach monopoly levels, in which case minority investors, like common owners, will benefit at the cost of consumer welfare.

Implications for Pricing: In order to translate the common ownership hypothesis into economic outcomes, both for understanding its economic significance and for drawing out testable empirical predictions, we need to be more specific about the strategic interaction of firms. We focus on the canonical case of differentiated Bertrand price competition with single product firms where firms compete by setting prices,

⁸It might be tempting to argue that two investors with symmetric ownership shares ought to have the same influence as one with their combined ownership. For this to be true everywhere, it would require that γ be invariant to β , e.g. $\gamma = 1$. Such a model has particularly perverse consequences. Revisiting Example 1, suppose that investor 1, who owns firm 1 in its entirety, bought a single share of firm 2’s stock (an epsilon-small stake). If $\gamma = 1$, then this seemingly innocuous purchase would discontinuously raise κ_{21} from 0 to 2.5. Even though the purchase is small, $\gamma = 1$ implies that the welfare of firm 1 would then be entirely on firm 2’s books. In order to avoid such discontinuities, we need $\gamma \rightarrow 0$ as $\beta \rightarrow 0$, which $\gamma = \beta$ trivially satisfies (but so, of course, would many other specifications).

as in O'Brien and Salop (2000). They each set a price p_f , and demand for product f is given by a function D_f that maps all prices into quantity sold, i.e. $q_f = D_f(p_1, \dots, p_f)$. Then writing marginal costs c_f , we have:

$$\pi_f = (p_f - c_f)D_f(p_1, \dots, p_f). \quad (2)$$

In a world without common ownership, firms choose p_f to maximize π_f . In the presence of common ownership, firms solve:

$$\text{Choose } p_f \text{ to max } \underbrace{\pi_f}_{\text{own profit}} + \underbrace{\sum_g \kappa_{fg} \pi_g}_{\text{common ownership concerns}}.$$

To characterize the prices firm f will set in equilibrium, we use the first order condition of firm f 's optimization problem. Plugging in (2) and taking the derivative with respect to p_f yields

$$0 = \underbrace{D_f(p_1, \dots, p_f)}_{\text{inframarginal}} + \underbrace{(p_f - c_f) \frac{\partial D_f(p_1, \dots, p_f)}{\partial p_f}}_{\text{marginal}} + \underbrace{\sum_g \kappa_{fg} (p_g - c_g) \frac{\partial D_g(p_1, \dots, p_f)}{\partial p_f}}_{\text{diverted}}. \quad (3)$$

The first two terms are the usual countervailing inframarginal and marginal effects of raising prices. Inframarginal customers do not switch, and so the price increase raises profits. Marginal customers do, and so those lost sales diminish profits. The third term, however, is new. It captures the profits generated by sales that are diverted to firm f 's competitors when they raise their price. According to the theory of common ownership, firm f values the profits of competitor g according to weights κ_{fg} . If products are net substitutes, then this third term in the first-order condition is always positive, and therefore it can be shown that prices will be strictly higher than they would be in a world in which firms maximize only own-firm revenue.^{9,10}

This result summarizes the heart of the controversy over the common ownership hypothesis, as well as its central testable implication: that consumers face higher prices because of the presence of common owners. Though their size is an empirical question, there is reason to believe they may be large. For comparison, consider a scenario we routinely regulate: the merger of two large competitors selling substitute products. In practice, the computation of the price effects of such mergers is equivalent to setting the profit weights of these firms on each other to 1, following the work of Nevo (2000). These effects are often large, but can be dampened by the presence of competitive third parties in a market. In contrast, common ownership affects *all* publicly traded firms. In the absence of competitive pressure from privately (or narrowly) held firms, this can lead to dramatic predictions for pricing and consumer welfare. However, we note several caveats to this theoretical analysis.

Remark: corporate governance, owners, and managers. The model of common ownership, that firms maximize a γ -weighted sum of investors portfolios, given by β , is implicitly a model of corporate governance. Why exactly do firms care more about larger shareholders? If they have more voting power, how does that translate into effects on firm decision-making?

⁹By dividing the last term in (3) by $\partial D_f / \partial p_f$ we obtain a diversion ratio, which has many applications in antitrust analysis (Werden, 1996; Farrell and Shapiro, 2010). For a recent treatment the estimation of diversion ratios, see Conlon and Mortimer (2019).

¹⁰It is also worth noting that if two firms do not compete in a product market, there should be no effect of κ_{fg} on prices. If two firms are part of the same vertical supply chain or produce complements, a positive κ_{fg} may even result in an increase in efficiency.

As we discussed above, the literature has followed Rotemberg (1984) in assuming $\gamma_{sf} = \beta_{sf}$, for two reasons: first, for empirical tractability. As we discuss below, data on large institutional investors' shareholdings are publicly available. But second, there is little guidance from the corporate governance literature on how to best model γ . As an alternative, one could use Banzhaf power, a model of voting power borrowed from the political science literature.¹¹ Alternatively (and in a different context), Crawford et al. (2018) propose that firms with higher $\sum_f \beta_{sf}$ be given less weight γ_{sf} . In our own computations we follow Rotemberg (1984) in setting $\gamma_{fg} = \beta_{fg}$. However, this is an area where new research in the corporate finance literature could yield great progress, and we hope to see more work in this area going forward. Beyond intuitive qualitative features – e.g., γ_{sf} ought to be monotone in β_{sf} – any parametric form seems arbitrary. In the presence of common ownership effects, though, we note that these forms may be empirically testable.¹²

Whatever the functional form, an oft-raised objection to the common ownership hypothesis concerns the conflation of ownership and management of assets. Large asset managers such as Blackrock and Vanguard work on behalf of their clients, who are the ultimate owners. Do they have a financial incentive to affect the coordination implied by common ownership? The common ownership profit weights κ assume that they do. While some work has endeavored to directly quantify their financial incentives to maximize the value of the assets they manage, a more direct reply is to perhaps take them at their word when they claim to have a fiduciary duty to their clients. Moreover, there is extensive empirical evidence (see Brav et al. (2018), Gilje et al. (2018b)) suggesting that they are, in fact, active in corporate governance. Whether their incentives are strong enough – or whether they take that fiduciary duty seriously enough – to want to soften competition in product markets is a difficult question with little direct empirical traction. Still less is known about the feasibility and costs of conveying such intentions to management. For this reason we believe that the debate ought to focus on the directly testable empirical predictions for economic outcomes – e.g. prices – in specific markets.

Remark: alternative strategic interactions. The model above is consistent with the most controversial claims of prior work, that common ownership means that firms set higher prices. However, firms do much more than price-setting. They make R&D choices, they engage in vertical contracting, and participate in a plethora of other activities for which it is not obvious that common ownership is a threat. Indeed, it may be efficiency-enhancing. For example, if common ownership arises between two vertically-oriented firms such as a retailer and a supplier, then common ownership may help to mitigate double marginalization (although it is worth noting that if such concerns were large, firms in such situations have better tools to align incentives, such as contracts). Alternatively, López and Vives (2018) show that the welfare implications of common ownership are ambiguous in the setting of R&D spillovers. The point being that while the literature has been focused on price-setting (and perhaps reasonably so, given the potential antitrust implications), there are other games in which the welfare effects of common ownership are not decisively negative.

This observation raises an important empirical point for future work. It may be that as the literature

¹¹A voter's Banzhaf power is proportional to the number of winning coalitions in which they are a swing voter Banzhaf (1965). It is computed by first determining the set of all winning coalitions of voters. For each winning coalition, pivotal voters are identified. The Banzhaf power for a voter is the number of cases where that voter is found to be pivotal, divided by the total number of pivotal votes.

¹²For example, the common ownership weights κ_{fg} described above are very sensitive to increasing concentration among asset managers and increasing scale. In contrast, the weights proposed by Crawford et al. (2018) normalize out investor scale, and emphasize portfolio composition.

unfolds, the clearest evidence on common ownership will come not from pricing games, but from models of entry, product positioning, R&D, or other strategic interactions. Modeling these other implications will generate a bevy of additional testable empirical implications for the theory, and is an open area for research.

Remark: indices of common ownership. Finally, we want to address the theoretical foundations (and limitations) of what are becoming commonly-used indices of market power, adjusted for common ownership. They go by names such as MHHI, MHHI-D, GHHI, etc. These indices are derived from a theoretical model of common ownership under symmetric Cournot competition – i.e., quantity-setting competition among a set of identical firms selling a *homogenous* output. These measures face the same problems that plague measures of market concentration such as HHI: They are theoretically uninterpretable as measures of market power and often misleading in differentiated products markets. Moreover, in contrast with direct measurement of the common ownership weights κ , concentration indices such as MHHI are extremely sensitive to market definition, as they are functions of market shares. For an explicit derivation of the MHHI measure from investor holdings, see Appendix C of [O’Brien and Salop \(2000\)](#). As we will discuss later, the use of these measures can raise identification problems in empirical applications, and has clouded the debate on the common ownership hypothesis.

3 Data

This section describes the different data sources available to study the common ownership hypothesis. The first part would be of use to anyone interested in conducting applied research in this area, but other readers should consider skipping ahead to the descriptives in Section 3.2.

Data on holdings of large investors come from multiple original sources and are also aggregated and distributed by third-party data firms. In particular, institutional investment managers that manage over \$100 million in “Section 13(f)” securities, which are defined by the SEC, must file a form 13(f) with the SEC on a quarterly basis to report their holdings. These reports include the CUSIP (Committee on Uniform Security Identification Procedures) number for each security held, which can be used to interface these data with other datasets. This is the primary source of data used in assessing common ownership. In addition, mutual funds must report their holdings on a semi-annual basis at the individual fund level using form N-30D. All of these filings are available on the SEC’s EDGAR platform going back to 1999, when electronic filing became mandatory; Thomson Reuters makes these data available going back to 1980 in different commercial datasets.

Ancillary datasets are often used in conjunction with the Thomson Reuters dataset. A first is from CRSP, the Center for Research in Securities Prices, which contains data on share prices as well as the number of shares outstanding, and is indexed by what is called a PERMNO for each security. CRSP also maintains lists of stock market index composition for a large set of indices. Wharton Research Data Services (WRDS) created a crosswalk from CUSIP codes to PERMNO codes. A second dataset often linked to the Thomson Reuters data is Compustat, which can be searched by CUSIP or their own proprietary GVKEY variable, and contains firm filings accounting data.

3.1 Data quality concerns

There have been a number of data quality concerns relating to the Thomson Reuters 13(f) dataset (referred to as the “S34” dataset) dating back to 2010. The dataset is commonly accessed through the WRDS platform, and WRDS has been active in trying to correct for a number of the data quality issues. [Ben-David et al. \(2018\)](#) identified a number of additional data quality issues with the Thomson Reuters dataset. As of July 2018, WRDS has addressed a number of the concerns raised and produced an updated version of the S34 dataset.¹³ However, some issues with the data relate to how they should be interpreted in light of the theory above, while other issues are strictly errata that must be corrected for. Below, we discuss various data issues and whether or not they have been addressed in the WRDS July 2018 update.

Short positions One oddity in the data is that on occasion, more than 100% of a firm’s outstanding shares as reported by CRSP are reported as owned in the 13(f) filings for that quarter. One explanation is double-counting of shares due to short-selling: if investment firm A lends out shares to investor B, who then sells those shares to investor C to create a short position, both A and C may report owning the shares. Data on short positions is available via Computstat. From a theoretical standpoint, only one agent per share should be able to exercise any control over the target firm at any point in time, and so this potential double-counting introduces some error into common ownership analyses. [Christoffersen et al. \(2007\)](#) investigate “vote trading” as an activist investor strategy, where shares are borrowed to exercise control at low cost, and find that “double voting” of lent shares is a pervasive phenomenon. The WRDS dataset presents holdings as reported and so makes no correction for possible double-counting. [Lewellen \(2011\)](#) finds that shares reported exceeding shares outstanding to be a rare occurrence and of a small magnitude. [Ben-David et al. \(2018\)](#) recommends updating the shares outstanding variable in the S34 dataset to be the values found in CRSP whenever the two disagree.

Dual-class Shares Many publicly traded firms have multiple classes of shares with different levels of voting control. From a data reporting standpoint, this is not a concern, as the securities have different CUSIP designations. However, the theory models above effectively assume one vote per share. To the extent that this is not the case, then any implied profit weights are incorrect. In principle, when studying a particular setting, one could explicitly adjust for dual-class shares using the γ term. In practice, in the descriptive exercise below, dual-class firms are removed from the sample. It should be noted that as with the retail investor share discussed below, different approaches to dual-class shares make very different predictions about profit weights under different models of common ownership.

Thomson Reuters S34 dataset WRDS and Thomson Reuters began to notice data irregularities in the S34 dataset in 2010. In addition, [Ben-David et al. \(2018\)](#) noted additional data issues, such as Blackrock disappearing from the dataset in 2014 and re-appearing with far too low a level of assets in 2015. A collaboration between those authors, WRDS, and Thomson Reuters resulted in an updated dataset as of July 2018 that is believed to address a large number of the concerns raised. WRDS has released a document

¹³In particular, 13(f) filings were mandated to be in an XML format starting in the third quarter of 2013, and WRDS has parsed these filings to guarantee accuracy of the Thomson Reuters dataset.

detailing the fixes.¹⁴

Retail and other small investors Given the reporting requirements for institutional investors features a specific cutoff for assets under management, the holdings of firms managing less than \$100 million are not reported. Individual investors need not necessarily report holdings, although in many cases individual investors employ large investment management firms to manage their portfolio and so those holdings are reported by those investment managers. To the extent that an individual has large holdings at multiple investment management firms, the data will not necessarily reflect the “common” element of those holdings. From a theoretical point of view, this creates an additional problem: suppose we observe the ownership of 60% of a firm’s common shares in the S34 dataset. The remaining 40% is held by diffuse investors about whom we know nothing. One assumption might be that these shares are held by atomistic investors, and so would each have $\beta_s = 1/N$, where N is shares outstanding. Another assumption might be that these investors are all undiversified and act in lock step, and so should be considered as one representative agent with $\beta_s = 0.4$. Both of these approaches could be justified, but would result in very different profit weights.

Insider Holdings The SEC requires “insiders” to make additional disclosures regarding holdings of a firm (forms 3, 4 and 5). An insider is defined as an officer, a director, or anyone owning 10% or more of a firm’s shares. Researchers investigating a particular set of firms should investigate any insider holdings information to see if the holdings are being reported by a larger custodial institution. For example, [Backus et al. \(2018a\)](#) show that the Kellogg Foundation held a major stake in Kellogg’s, but that the foundation’s shares were reported in Bank of New York’s 13(f) filings in many quarters. The foundation itself is not diversified and so this reporting mixed the holdings of an undiversified owner (the Kellogg Foundation) with the holdings of their diversified bank (Bank of New York).

Aggregation One final point raised by [Ben-David et al. \(2018\)](#) and not yet addressed by WRDS is that there may be multiple entities reporting in the S34 dataset that are actually subsidiaries of one large entity. To make a simple example, in the final quarter of 2016, the following entities report holdings separately: “BLACKROCK INC,” “BLACKROCK ADVISORS, LLC,” “BLACKROCK ASSET MGMT IRELAND,” “BLACKROCK INVESTMENT MGMT, LLC,” “BLACKROCK JAPAN CO., LTD.,” and “BLACKROCK INVT MGMT (UK) LTD.” (each under a unique “mgrno” in the S34 dataset). If the control rights of these entities are coordinated, then one would want to consolidate them into one entity before computing profit weights. While this has not been addressed by WRDS, it is not difficult for an applied researcher to consolidate these entities.

Source Documents Finally, some researchers have noted that the source documents themselves – Form 13(f) filings made by institutional investors – contain errors. In particular, [Anderson and Brockman \(2016\)](#) document irregularities and caution the use of 13(f) filings in research or in investing. The SEC’s Inspector General in 2010 noted several issues with how 13(f) filings are handled, see [Securities and of Inspector General \(2010\)](#). The authors of this current paper noticed when examining ownership of airlines that many filings

¹⁴S12/S34 Regenerated Data, accessed at https://wrds-www.wharton.upenn.edu/documents/952/S12_and_S34_Regenerated_Data_2010-2016.pdf

contained errors around bankruptcy events. This emphasizes the importance of manual data-cleaning when studying common ownership in particular industries.

Alternative Data Sources Due to the irregularities in the Thomson Reuters dataset, particularly around 2011-2013, the authors in (Backus et al., 2018b) gather all original 13(f) filings from the SEC starting in 1999 (the era of mandatory electronic filing) and parse them to obtain holdings for S&P 500 firms.¹⁵ The authors show examples where major firms are only reported to have a few institutional holders in the Thomson Reuters dataset in some quarters but in fact have many holders in the raw 13(f) filings.

3.2 Descriptive data on common ownership

While there have been a number of analyses and rebuttals regarding the common ownership hypothesis, it is undeniable that a few large investment management firms have amassed large, diversified portfolios of all major publicly-traded firms over the past few decades through their mutual funds or exchange-traded funds. Below, we document some basic facts contained in the S34 dataset and our own novel dataset of 13(f) holdings about ownership of firms listed in the S&P 500 Index. It should be noted that descriptives based on 13(f) data require no market definition. While one firm may have an implied profit weight for another, if the two firms do not compete in a product market, it is unlikely that there would be any effect from common ownership. After examining the economy as a whole, we will pick some specific product markets to examine. We have made some data corrections to the dataset to, for example, consolidate all “BLACKROCK” entities, as well as consolidate all “STATE STR” entities.

As a starting point, Figure 1 shows the number of 13(f) managers reporting over time.¹⁶ The number of such managers has greatly increased, and at least part of the increase is due to the nominal reporting barrier. One dollar in 1980 is worth approximately \$3.24 in 2018, and yet the number of 13(f) investment managers has more than quintupled in that time. This suggests both substantial entry into the investment management space, and more diffuse management of investments over time.

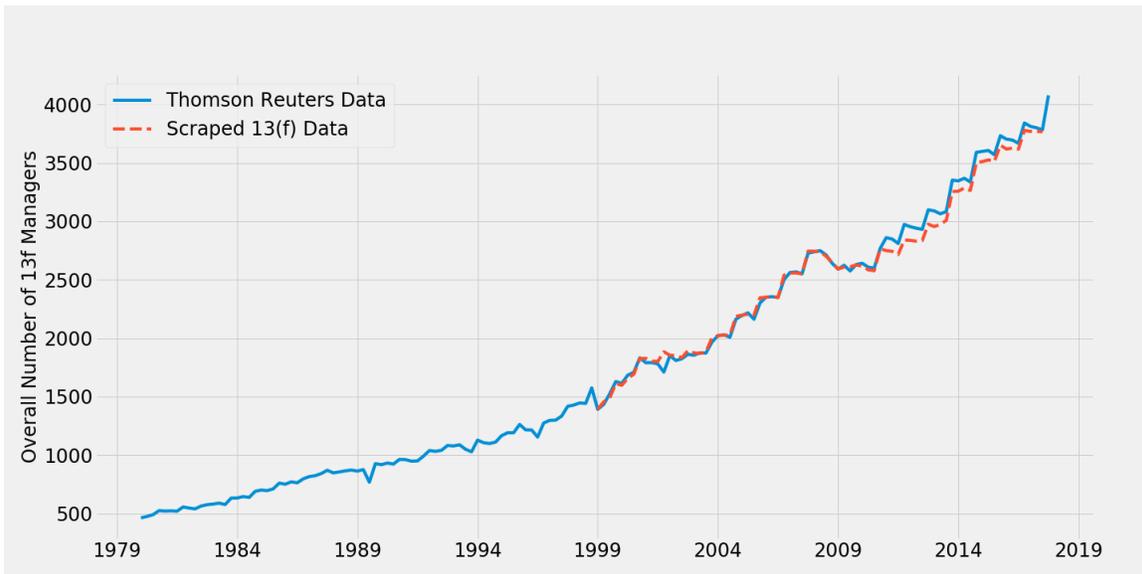
The red line in Figure 1 indicates that most of the new managers are classified as *undiversified managers*, which we define as having holdings in fewer than 50 of the S&P 500 firms. However, Figure 3 shows that the share of the average S&P 500 firm owned by all 13(f) managers has greatly increase over time, from under 40% in 1980, to approximately 80% in 2017. While it is true that publicly traded firms represent a decreasing share of the economy over time, the S&P 500 is specifically chosen to reflect the overall United States economy, and so this change is reflective of a change in investment concentration.

In contrast, Figure 3 specifically plots, over time, the average percent of S&P 500 firms held by Blackrock, Vanguard, and State Street over time. Blackrock here includes holdings by Barclays, whose iShares product was acquired in 2009. These three major firms went from each holding under 1% of a typical S&P 500 firm to holding roughly 5% (State Street) to over 7% (Vanguard). These three investment managers were chosen given their emphasis in the literature and large holdings at the end of the sample, but we cannot rule out

¹⁵The authors make the profit weights available for the firms they study and are happy to provide code to researchers interested in studying different firms.

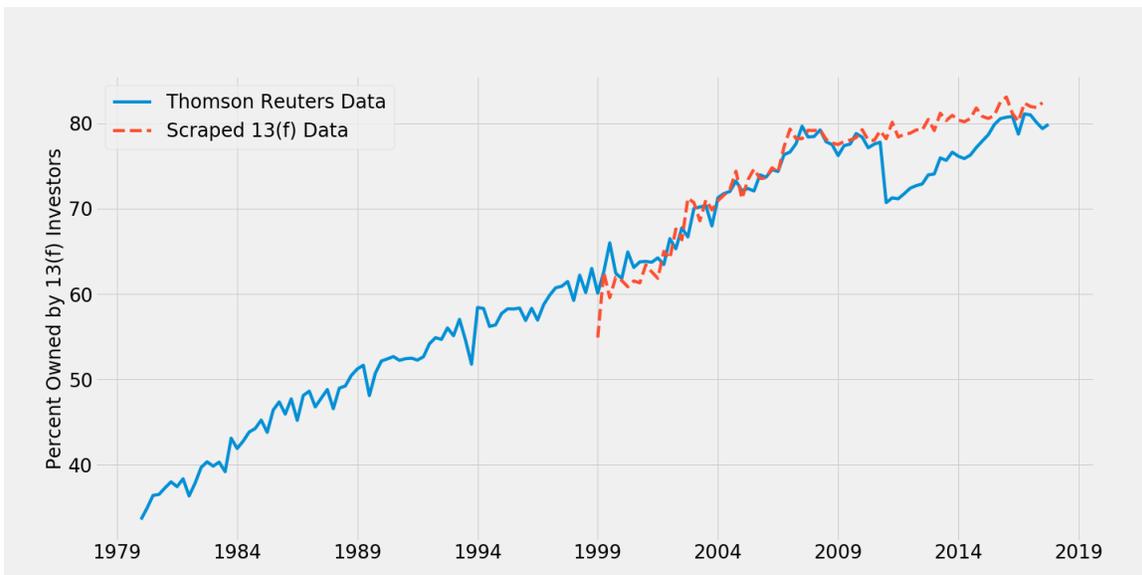
¹⁶As a reminder, a 13(f) manager is a legal definition by the SEC. The relevant part of the definition is that the firm manages over \$100M of securities.

Figure 1: Number of 13(f) Managers



Notes: This figure depicts the number of 13(f) managers in our dataset. The red line depicts those that are “undiversified,” i.e. have holdings in fewer than 50 S&P 500 firms, whereas the blue line includes all investors.

Figure 2: Share of S&P 500 Ownership Over Time

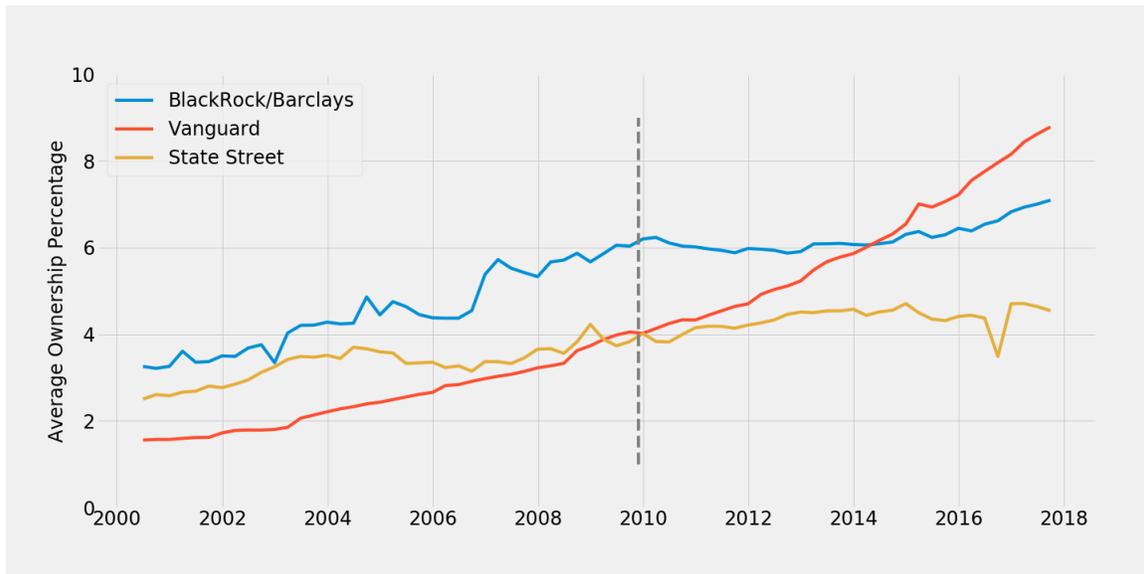


Notes: This figure depicts the sum of the holdings of investors in our dataset in a representative firm or, equivalently, what we are calling the institutional share of ownership.

that other investment managers could also be plotted here and would show qualitatively similar patterns.

The implication of these two figures is that there has been a massive increase in assets under management at investment firms over time; this increase has coincided with the entry of many new, large investment managers; but there has also been a large amount of consolidation in the “diversified” investment manager

Figure 3: Ownership of S&P 500 Firms



Notes: This figure depicts the assets under management (AUM) over time of Blackrock/Barclays, Vanguard, and State Street. Data are from our scraped and parsed dataset of raw 13(f) filings.

space, particularly with firms like Blackrock, Vanguard, State Street, and others.

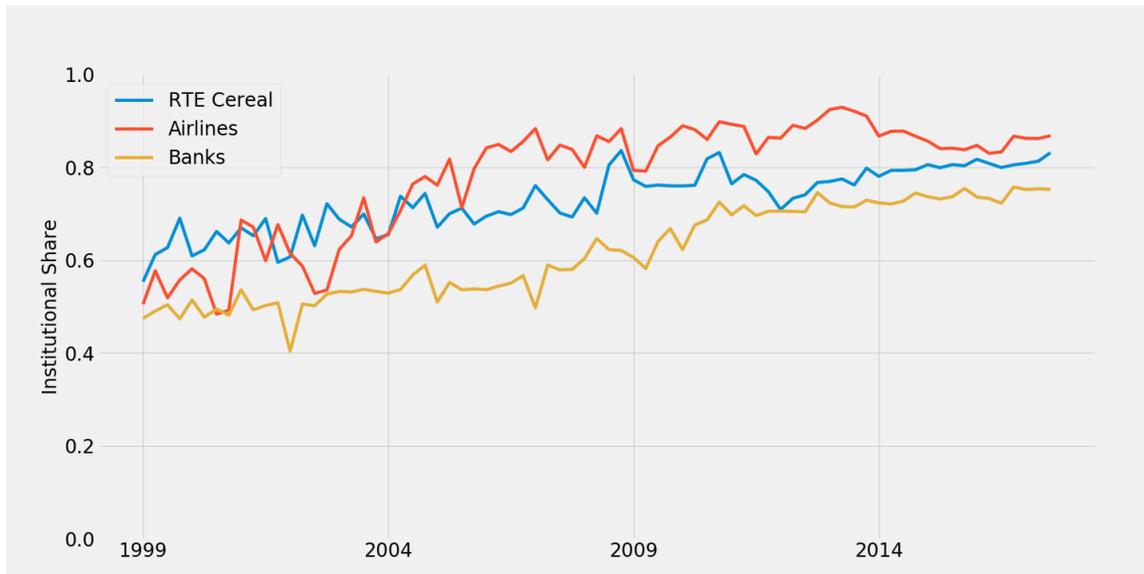
One concern about looking at the S&P 500 as a whole is that many of the firms are operating in separate markets, and so common ownership is unlikely to affect strategic decisions. For example, even if a pharmaceutical company and an airline share a large investor, since their decisions are unlikely to affect each other's profit, there is no change in incentives from common ownership. To address this, Figure 4 shows institutional ownership in three specific sectors: ready-to-eat cereal, airlines, and commercial banks. We hand-collect data on firms for the first two product markets and use the Compustat SIC code to classify commercial banks in the S&P 500 by their SIC code of 6021.¹⁷ As can be seen, the overall trends in these specific product markets are similar, with large shares being accumulated by investment management firms over time. When looking at specific industries, larger swings in holdings are visible due to mergers, bankruptcies, entry, etc.

3.3 Implied profit weights

Given this data, we can also compute implied profit weights that firms in the S&P 500 Index have for other firms in the index over time. The figures below are from Backus et al. (2018b) and highlight the long-term trends in implied profit weights. The trend is steadily increasing over time. This reflects the ownership patterns described in 3.2 – both increasing diversification of investor portfolios, as well as increasing concentration in the market for asset managers. Figure 5 shows that for a typical firm in the S&P 500, the profit weight for any other firm in the S&P 500 implied by the common ownership hypothesis has increased

¹⁷This calculation therefore excludes privately held banks and credit unions that are likely to be narrowly held and thus overstate the overall common ownership effect for the banking sector.

Figure 4: Institutional Ownership of Firms Within Product Markets



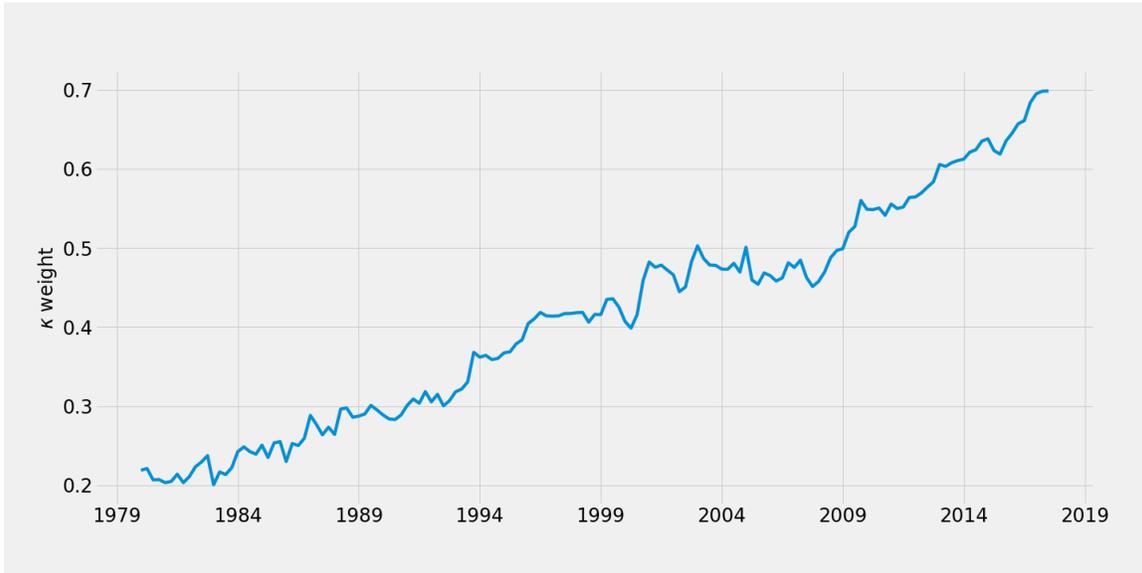
Notes: This figure depicts the sum of the holdings of investors in our dataset in a representative firm or, equivalently, what we are calling the institutional share of ownership for airlines (left), and banks (right).

from approximately 0.2 in the 1980s to approximately 0.7 in the late 2010s.

In terms of heterogeneity, one would expect profit weights to differ across firms in terms of the retail share of their ownership. As discussed in Section 2, if small investors effectively have no control, then firms with large retail shares of investors see any common ownership effect magnified. Figure 6 shows that after controlling for year fixed effects, there is clearly a relationship between profit weights and the retail share of a firm. The more of a firm held by retail investors, the more magnified is the implied common ownership effect of investment managers.

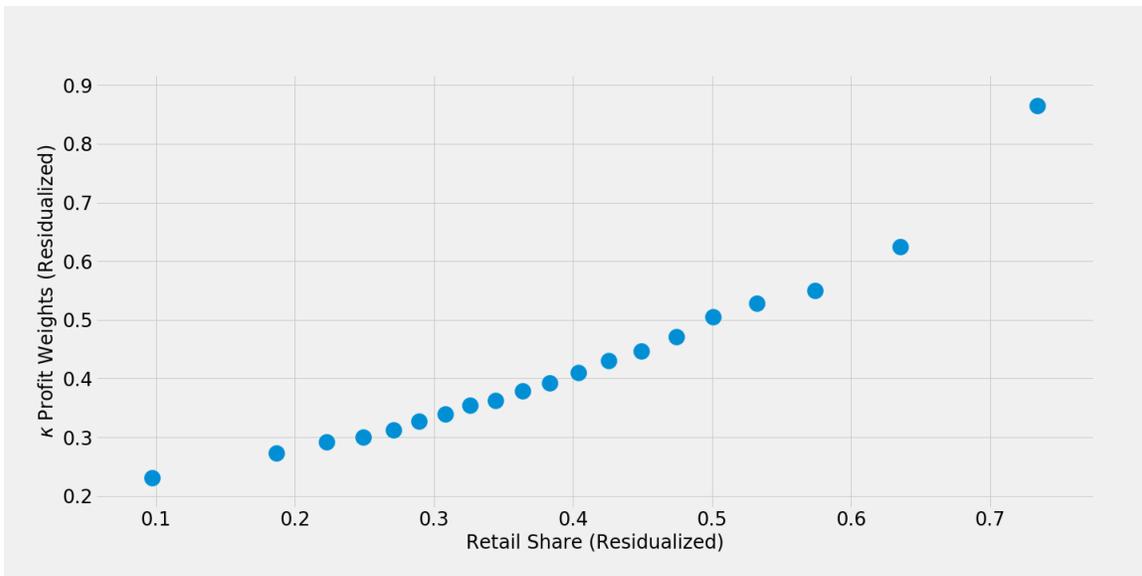
Finally, we will examine three specific product markets that have received attention in the empirical literature on the common ownership hypothesis: airlines, banks, and ready-to-eat cereal. For the first two, shown in Figure 7, we compute pairwise profit weights for 1980-2017. For banks, we use SIC code 6020 to select firms from the S&P 500 sample. For airlines, as many airlines are not S&P 500 constituents, we manually created a set of 23 securities associated with nationwide airlines over the relevant time period. Given the number of mergers and bankruptcies, the time-series unfortunately reflects a large amount of entry and exit and significant cleaning of the data was required (many institutional investors continued reporting holdings of non-existent securities long after bankruptcies). We compute the average profit weight on rivals and average within quarter. For the third product market, presented in Figure 8, we borrow from [Backus et al. \(2018a\)](#) and show all pairwise profit weights for 2003-2017. The fact that Kellogg's has large, undiversified owners is immediately apparent.

Figure 5: Average Profit Weights for S&P 500 Firms, 1980-2017



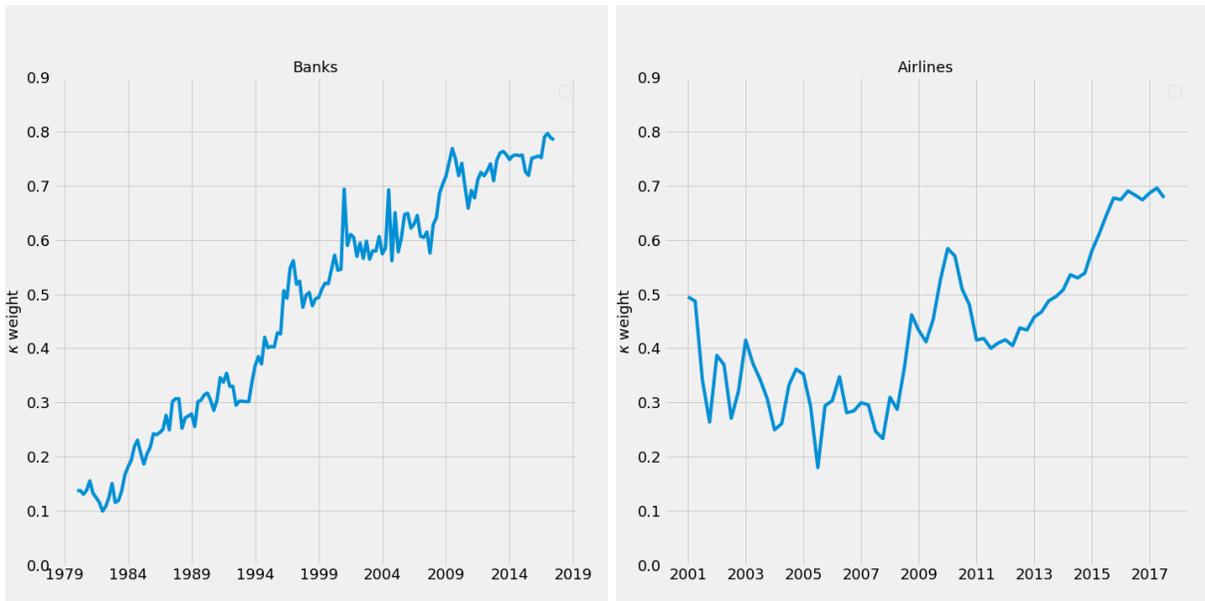
Notes: This figure depicts the average common ownership profit weight for pairs of S&P 500 firms by quarter from 1980 through 2017.

Figure 6: Average Profit Weights on Other Firms by Retail Share
S&P 500 Firms, 1980-2017, Year FEs



Notes: This figure reports the mean common ownership profit weight for twenty equally sized groups of firms ranked by retail share, which is given by 1 minus the sum of the institutional holdings observed in our dataset. We have residualized on year fixed effects in order to remove any mechanical relationship introduced by the downward trend in retail share.

Figure 7: Average Profit Weights on Other Firms For Banks and Airlines



Notes: This figure depicts the average common ownership profit weight for pairs of airlines (left) and banks (right) by quarter from 1980 through 2017.

Figure 8: Average Profit Weights on Other Firms For RTE Cereal



Notes: This figure depicts the common ownership profit weight for each major ready-to-eat cereal manufacturer on each other manufacturer from 2003 to 2017. Note that the own-profit weight is normalized to 1.

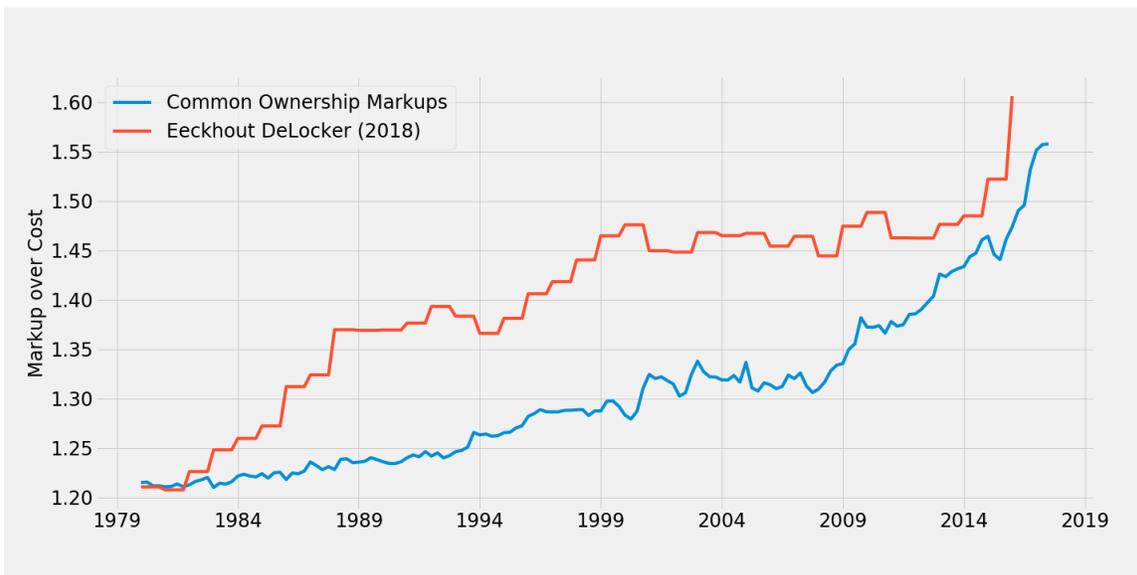
3.4 Calibrating Price Effects

As we discussed above, the manifestations of common ownership profit weights are myriad. However, a first-order question one might ask is whether the scale of the increase in profit weights depicted in Figure 5

could generate large increases in markups. In order to answer this question, we calibrate a simple model of Bertrand competition in [Backus et al. \(2018b\)](#).

Consider a market with eight symmetric firms, which would imply an HHI of 1250. We assume that these firms sell a single product each, that products are substitutes, and that consumers choose the product which gives them the most utility $u_{ij} = \alpha - \beta p_j + \varepsilon_{ij}$, where ε_{ij} is an i.i.d. logit preference shock. We calibrate α and β so that the outside good share is approximately zero, and we match the markups of [Eekhout et al. \(2018\)](#) of 1.21 in 1980. This yields a price elasticity of -7 .

Figure 9: Simulated Markups: 1980–2017



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

With this calibration in hand, we then simulate prices forward to 2017 in Figure 9. We stress that this is a back-of-the envelope calculation. We find that while we are able to explain approximately 90% of the rise in markups found in [Eekhout et al. \(2018\)](#), the timing is off – that the rise of common ownership incentives seems to substantially lag the rise of markups depicted in [Eekhout et al. \(2018\)](#). It is difficult to conclude that common ownership is driving markups; however, we do find that if we take the theory seriously, and it remains to be shown that we should, the effects would be economically large.

4 Empirical Results in the Literature

The preceding sections have documented how the true primitive of common ownership —the profit weights κ_{fg} —have increased over time. We now turn to empirical evidence on the competitive effects of common ownership. Much of the initial empirical literature that addresses the common ownership hypothesis has its roots in the so-called *Structure-Conduct-Performance* (SCP) paradigm in Industrial Organization which dates back to [Bain \(1951\)](#). The central conceit of the SCP literature was to regress measures of firm performance (profits, prices, and investment) on measures of market concentration (market shares, Herfindahl

Indices (HHI),¹⁸ etc.), often using data across industries. Much of the recent literature measuring the impact of common ownership on prices (such as Azar et al. (2018b), Azar et al. (2016), Antón et al. (2018b), Kwon (2016)) can be viewed as an extension of the SCP approach.

Indeed, the regressor of interest in a number of these papers is what is known as the Modified HHI Delta (MHHD). It amounts to the product of the two market shares of (s_f, s_g) in market m and period t , interacted with the κ_{fg} profit weight.¹⁹ The motivation is to extend the typical notion of concentration to include common ownership and to run SCP-style regressions of the form:

$$\log p_{fmt} = \beta_0 + \beta_1 \underbrace{HHI_{mt}}_{\sum_f s_{fmt}^2} + \beta_2 \underbrace{MHHD_{mt}}_{\sum_f \sum_{g \neq f} \kappa_{fg} s_{fmt} s_{gmt}} + \beta_3 \log s_{fmt} + \varepsilon_{fmt}.$$

These recent papers tend to find both that concentration is associated with higher prices ($\beta_1 > 0$), and also that higher effective concentration from common ownership is associated with higher prices ($\beta_2 > 0$).

As early as the 1970's, there were a number of theoretical criticisms of the SCP approach. The relationship between prices and measures of market concentration (including HHI) is predicated on very strong assumptions about market structure, the cost structure of the industry, and conduct (such as Cournot competition with homogeneous products and symmetric marginal costs). Demsetz (1973), a seminal paper that anticipated the direction of modern IO, suggests the following example to highlight the problem: suppose we reallocate market share from a smaller, high-cost firm to a larger, low-cost firm. While this would increase most measures of market concentration (such as HHI) and market power (price cost margins), this might also lead to lower prices and higher surplus for both consumers and producers.²⁰

A second critique, often associated with Bresnahan (1989), is that the relationship between prices and concentration (market shares) is an equilibrium relation that may represent either a supply curve or a demand curve (or more likely, neither). Put simply: *Which “causal relationship” between price and quantity are we trying to identify?* Absent specific exclusion restrictions, regression analysis might not identify any meaningful economic relationship. This problem is further exacerbated by the fact that exclusion restrictions required by instrumental variable (IV) approaches are hard to come by. The classic argument for exclusion of an instrumental variable is “something that appears in another equation.” The challenge arises because both prices and market shares are equilibrium outcomes (linked by first-order conditions), and most relevant variables would affect both supply and demand; thus, almost nothing can be excluded.²¹

The demise of the original SCP literature is summarized in Schmalensee (1989), which concludes by noting that this line of research...

“... has taught us much about how markets *look*, especially within the manufacturing sector in

¹⁸The HHI is computed as the sum of squares of market shares, or $HHI = \sum_f s_f^2$.

¹⁹The MHHD was originally conceived of in the context of homogeneous good Cournot competition. It is computed as $MHHD = \sum_f \sum_g \kappa_{fg} s_f s_g$, which can be rewritten as $MHHD = HHI + \sum_f \sum_{g \neq f} \kappa_{fg} s_f s_g$. The latter term is referred to as the MHHD-Delta, or the difference between HHI and MHHD.

²⁰Another famous (and simple example): in the perfectly competitive model with upward sloping supply there exists a mechanical relationship between HHI and price even though all firms are price takers, and so HHI cannot have a causal effect on prices charged by firms. That is, industry concentration would have no causal effect on prices even though a regression would show a statistically significant relationship.

²¹Fixed costs have been proposed as a potential instrument for concentration or the number of firms, but are notoriously difficult to measure. Other proposed instruments for concentration include mergers undertaken for unrelated reasons.

developed economies, even if it has not shown us exactly how markets *work*.”

It is useful to note that a number of criticisms of the SCP literature do not apply to all of the recent literature on common ownership. One early critique of the SCP literature was that performance was often measured using unreliable measures of accounting profits rather than prices or economic markups. Another criticism of that literature was that identification often relied on variation in concentration across disparate industries rather than looking at variation in concentration across geographic markets within an individual industry. [Gutiérrez and Philippon \(2016\)](#) provide a recent example of the former approach while [Azar et al. \(2017\)](#) provide an example of the latter.

The prior SCP literature (and its demise) highlights some important points. In any attempt to study the relationship between common ownership and market power, there will always be multiple endogeneity issues. The first endogeneity issue is the endogeneity of the κ profit weight. If successful firms attract more common investors, or become a larger share of the index, we might expect a positive correlation between κ and market outcomes such as prices or markups. It may be possible to address this endogeneity problem with instrumental variables strategies, which we discuss below. The second endogeneity problem is that quantities, market shares s_{fgt} , or concentration measures HHI_{mt} are determined simultaneously with prices. Here, instruments are difficult to come by – because quantity (or market share) enters both the equation for supply and demand. For this reason, if we are trying to identify causal effects of common ownership on market power, then using indices that conflate κ , which captures useful variation in ownership structure, with market shares, which capture problematic variation that biases our estimates, seems particularly problematic. This is exactly what MHHI-based concentration measures do, and why we believe the literature on common ownership should move past them.

4.1 Empirical Approaches

Based on the above, a primary issue in empirically studying the common ownership hypothesis either within the SCP framework or outside of it is that one would need some kind of natural experiment that randomly changes the level of common ownership so that one could measure the effect on an outcome of interest. This need is worsened by the fact that ownership patterns and therefore κ values are at the firm-quarter level, and so exploiting geographic variation in common ownership is not possible unless some firms do not operate in all markets. Even an instrument that did randomize common ownership would not be sufficient in many cases, however, as in differentiated product markets, any such shocks would have theoretically ambiguous effects.

Many papers of the SCP style regress outcome variables on measures such as MHHI (or MHHID, or GHHI). Since these measures vary with market shares, it is possible to get geographic variation within a time period (for example, if banks compete locally, or if airline routes can be considered separate markets). However, as explored in [O’Brien \(2017\)](#) and [Backus et al. \(2018a\)](#), if products are differentiated, there would be spurious correlations between those market shares and outcome variables, and so simple regression analysis is unreliable. It is trivial to generate examples in which, under differentiated product competition, when firms price competitively and ignore common ownership, there is a spurious, mechanical correlation induced between prices and MHHID. This correlation can be positive or negative.

Therefore, researchers would like to have instruments that affect the level of common ownership in a market without directly affecting prices (or whatever the object of interest it). One instrument that has been used in several papers to date is the 2009 acquisition of Barclays Global Investors by Blackrock (see, for example, [Azar et al. \(2018a\)](#)). This combined the \$1.7 trillion in assets that Blackrock had been managing with the \$1 trillion that Barclays had been managing, creating a much larger common owner, with arguably zero impact on, say, airfares. Therefore, this potentially creates variation in the time series of common ownership that can be used to analyze prices in a product market. Another instrument that has been used is the 2003 mutual fund scandal, which revealed that a large number of mutual funds had engaged in “late trading” and “market timing.” In the months and years following this revelation, vast amounts of capital were withdrawn from the affected mutual funds. This created plausibly exogenous variation in common ownership, depending on how much of a firm was owned by the affected funds prior to the scandal (see, for example, [Antón and Polk \(2014\)](#)).

An additional challenge is that concentration measures like MHHI vary at the *market level* while the theory of common ownership implies asymmetries in profit weights across firms within the same market. In [Backus et al. \(2018a\)](#), the authors study the ready-to-eat cereal market using a structural model and exploit the fact that Kellogg’s has large, undiversified owners for historical reasons; separately, Post changed hands many times and was at different points in time included or excluded from the S&P 500; and Quaker is part of a much larger organization (Pepsi), which has implications for profit weights. They find that SCP-style regressions lead to spurious findings (in particular, a negative coefficient when regressing price on MHHID and controls), and that using a structural approach finds no support for the common ownership hypothesis in ready-to-eat cereal pricing. It should be noted, however, that the authors are able to quantify the magnitude of the effect of “turning on” the common ownership hypothesis: they estimate that price effects from common ownership, if firms behaved according to the model, would be larger than those from any possible 4-to-3 merger in the industry. This implies that the common ownership incentives may in fact be very strong, even if firms do not seem to respond to them in that context.

4.2 Initial Evidence and Replies

The recent wave of empirical attention on the common ownership hypothesis began with several papers doing SCP-style regressions. In particular, a large amount of attention was drawn to this question because of papers such as [Azar \(2011\)](#), [Azar et al. \(2016\)](#), and [Azar et al. \(2018a\)](#), which documented common ownership and regressed prices on measures of concentration. These papers attracted several responses, such as [Kennedy et al. \(2017\)](#) and [Dennis et al. \(2017\)](#), that claimed to find no common ownership effects using alternative approaches in the same industries. The early empirical papers deserve much credit for bringing attention to this question, which had been dormant for many years. However, it is difficult to draw strong conclusions given the concerns about the empirical methods, as discussed above. In addition, while theory predicts that the common ownership hypothesis should manifest itself throughout the economy, even if evidence were found in a particular sector either in support of or against the hypothesis, one could not simply extrapolate to other product markets and draw broad conclusions.

Airlines and Banks Azar et al. (2018a) exploits the acquisition of Barclays Global Investors by Blackrock as a shock to common ownership of airlines to investigate whether there is a price effect. A market is defined as an airline route. The authors find that routes with increases in MHHI saw higher fares and argue for a causal interpretation. In Azar et al. (2016), some of the same authors examine the various fee and interest rate schedules of banks and regress these measures on “GHHI,” a version of the MHHI measure that also allows for inter-ownership (e.g. Firms A and B compete, and Firm A owns shares of Firm B, while Firm B also owns shares of Firm A). Since common ownership is national, the variation in GHHI comes from the relative market shares in different markets, as well as differences in levels of institutional ownership (due to, for example, a bank being included or excluded from a financial index). The authors find an increase in prices associated with GHHI. These projects seem to date back to Azar (2011), which notes the increase in institutional investor ownership of public firms over the 2000-2010 timeframe, while focusing on models of how common ownership may manifest in different kinds of product markets and different models of competition (homogeneous Cournot and Bertrand, differentiated Bertrand, etc.).

Airline Replies and Responses Azar et al. (2018a) attracted two “reply” papers that revisit the question of airline routes and common ownership. The first, Dennis et al. (2017), applies existing methods from the airline pricing literature to clean the data in ways that eliminate certain types of fares that are not representative and also address the many bankruptcies in the data. They also restrict attention to shares that investment managers report to have “sole” voting authority, as opposed to “shared,” and revisit weighting in the regressions. They critique MHHI regressions and find no evidence of the common ownership hypothesis in airline pricing. Azar et al. (2018b) directly replies to this paper by showing discrepancies in data construction on ownership and robustness of those authors’ earlier paper.

Second, Kennedy et al. (2017) attempts to recreate the Azar et al. (2018a) dataset and main results and then regresses prices on profit weights and instruments using the same Barclays-Blackrock merger, but also membership in the Russell 1000 index. They also estimate a structural model with nested Logit demand and a parameterized ownership matrix. Overall, they find no support for the common ownership hypothesis in airline prices. In Azar et al. (2017), the authors of the airlines paper critique this response, noting that there are differences in how the two papers construct instruments as well as two large concerns with the Kennedy et al. (2017) analysis: unexpected signs in the first-stage regressions that call in to question the validity of the instruments, and the structural estimates of costs are nonsensical.

Other Early Evidence One pair of papers examines the effect of common ownership on inter-firm relationships. In Cici and Rosenfeld (2015), the authors examine the market for syndicated loans. Looking at pairs of lenders and borrowers who had previously contracted, they show that loans are more likely in the future if the pair gain a shared investor that holds more than 5% of both. In Freeman (2018), the author uses Compustat data on supplier relationships to examine how they respond to changes in Common Ownership. The paper exploits the 2003 mutual fund scandal and ensuing changes in common ownership. The result is that supplier relationships with more common ownership are “stronger”, in that they have a lower propensity to end. Both of these papers avoid regressing price on measures of concentration (the first uses an indicator for having a common owner with at least a 5% stake in both parties; the second uses two different ad-hoc measures of shared ownership), but as transaction events are also highly endogenous, the

bar is high for regression-based evidence.

He and Huang (2017) examine market share growth (as measured by sales within SIC code) of firms as a function of common ownership, using bank mergers as a natural experiment that alters common ownership. The authors find through difference-in-difference estimates that market shares increase from an increase in common ownership.

Broader Criticisms, Discussions, and Alternatives **Gramlich and Grundl (2017)** and **Backus et al. (2018b)** argue that the correct primitive to examine is the profit weights instead of measures such as MHHI and GHHI. The first paper examines the banking sector and finds that results are sensitive to different specifications, but that overall evidence of the common ownership hypothesis is mixed. The second paper looks at the S&P 500 in its entirety from 1980–2017 and finds that implied profit weights have increased significantly. Relatedly, **O’Brien (2017)** shows that regressions of price on measures of concentration are fundamentally problematic as they have no basis in theory. There may be spurious relationships as concentration is an equilibrium object, and simple models show non-monotonicity of price in concentration measures. The author suggests structural approaches to the empirical challenge.

In **Gilje et al. (2018a)**, the authors micro-found a voting model of common ownership, allowing for inattentive investors, to understand how changes in common ownership affect actual managerial incentives. If investors are inattentive to firms they own, then firms have fewer incentives to respond to investor portfolios. Diversified funds, as they spread out their holdings, are necessarily less attentive all else equal to their holdings. This is offered as a potential reconciliation between the rise in common ownership and weak managerial incentives to act on it.

In **Backus et al. (2018a)**, the authors estimate a structural model of demand in the ready-to-eat cereal market, incorporating profit weights implied by common ownership in certain specifications. The authors test alternative models of conduct and conclude that there is no statistical evidence for the common ownership hypothesis in prices when looking at historical data.

Mechanisms A first criticism against the common ownership hypothesis is why an investment management firm would want this to occur. **Lewellen and Lewellen (2017)** directly investigates the financial motive of investment managers to engage with their portfolio constituents. The authors show that even small increases in portfolio value generate significant increases in management fees for managers. Thus, an incentive exists to engage in any activity that increases the value of firms in a portfolio.

A second criticism was that a large missing element of the literature is how would this effect actually occur between investors and firms. **Antón et al. (2018a)** examines how executive compensation may be affected in a world of common ownership as a link between common ownership and actual product markets. In particular, if firms in an industry exhibit a high level of common ownership, these common owners might not want strong managerial incentives as total industry profits could fall. Instead, they might prefer softer competition between the firms. Empirically, the authors test how sensitive managerial wealth is to their firm’s performance, and find a negative effect of common ownership on the sensitivity of manager wealth with respect to firm performance, implying common ownership leads to weaker managerial incentives. This

may relate to a wider literature on tacit collusion. A recent paper (Aryal et al. (2018)) noted outside of the context of common ownership how airlines may use earnings calls to signal cooperation through “capacity discipline”.

Finally, there is the question of whether shareholder votes are a channel through which this effect may take place. He et al. (2017) study proxy voting by institutional investors and finds that diversified investors are more likely to vote against management. They conclude from this that diversified investors internalize corporate governance externalities, improving overall corporate governance.

Alternative Manifestations The majority of the early literature focuses on prices. More recently a number of papers have suggested alternative manifestations of the common ownership hypothesis. For example, Antón et al. (2018b) examines an upside side of common ownership: the potential for greater R&D incentives. The paper gathers a few different measures of patents as well as R&D expenses and models them as a function of common ownership. Theory makes ambiguous predictions as technological spillovers would be a benefit if firms take rival profits into account when making R&D decisions. However, if innovation simply leads to business-stealing, then the common owners would not want it. Empirically, the authors show a positive correlation between common ownership and both inputs and outputs of R&D. In a related paper, López and Vives (2018) posit a theory model of how R&D incentives change under common ownership in the context of symmetric, Cournot competition. They show conditions under which common ownership may be welfare-enhancing or welfare-negative.

With regards to entry, Newham et al. (2018) examine the incentive for a generic drug manufacturer to begin producing a drug after patent expiration in the US as a function of common ownership. The authors show that an increase in common ownership between generic and branded manufacturers decreases the likelihood of entry in the time period 2004-2014.

Summary This area of research is very active. The initial evidence is very provocative: as pointed out in Backus et al. (2018b) and Backus et al. (2018a), the potential effect of common ownership on prices could be very large. However, we remain skeptical of existing evidence of an effect of common ownership on prices in a few particular sectors. More work is needed to examine different settings with more appropriate methods before any major conclusions can be drawn, and so we hope to see more research in this area going forward. This research is particularly important because our understanding of the scale and magnitude of any purported common ownership effects should precede and guide any discussion of policy responses, some of which is already underway (Scott Morton and Hovenkamp, 2018; Posner et al., 2017; Patel, 2018).

Tables 3 and 4 briefly summarize the papers that either empirically investigate common ownership, reply to such papers, discuss the topic, or make recommendations.

Moving forward, however, we believe that the lessons of modern IO, and the demise of the SCP approach should be taken to heart. Aggregate concentration indices such as MHHI tend to conflate plausible exogenous variation (from instruments such as, e.g., mergers of asset managers) with endogenous variation in market share. Alternatively, focusing directly on the primitive of common ownership – the profit weights that characterize the firm’s objective function – allows the researcher to isolate channels through which exogenous

variation affects prices or other market outcomes. This does not necessarily free the researcher from making assumptions about the form of that competition, however it does relax the constraints imposed by the highly restrictive, stylized models (e.g., Cournot competition with symmetric costs, in the case of MHHI) on which those aggregate indices are typically based.

5 Conclusion

The underlying theory of the common ownership hypothesis is intuitive and straightforward: managers answer to investors, who may have preferences other than static profit maximization. The potential implication of this theory is enormous: common ownership could imply substantial market power throughout the economy. However, evidence of an effect of common ownership on prices is, at this point, suggestive at best. Early methods used to investigate the question are problematic, and so more work is required before broad conclusions should be drawn.

We believe the most fruitful direction for future research on competitive effects of common ownership would focus on attempts to measure the impact of within a single industry, with a focus on pairwise profit weights rather than market-level concentration measures as the variable of interest. The most convincing identification strategy would focus on events in the financial space which provide substantial (and plausibly exogenous) variation in profit weights, while also exploiting cross-firm variation in those weights. The asymmetry with which firms treat the profits of rivals under the theoretical models of common ownership provides a set of clear and testable implications, that have not been sufficiently exploited by the early empirical literature on the topic.

Beyond competitive effects, there exist a number of avenues for additional research. First, the question of corporate governance looms large, as there is little to guide the choice of control weights (γ) used in empirical work. Second, whether under the common ownership hypothesis or the conventional assumption of profit maximization, the mechanism by which investors' preferences translate into strategic decision-making and pricing remains poorly understood. Third, it is not clear that pricing is the most salient variable to test the common ownership hypothesis. Profit weights affect *all* strategic decisions by the firm, so it may be that for some markets, the choice of product characteristics or investment decisions may yield more powerful tests of the hypothesis. Finally and most importantly, however, we need more empirical work in more industries to get more pieces of evidence. No one positive finding could necessarily be extrapolated to the economy as a whole, and so many data points are needed.

References

- Anderson, A. and Brockman, P. (2016). Form 13f (mis)filings. Working Paper.
- Antón, M., Ederer, F., Giné, M., and Schmalz, M. C. (2018a). Common ownership, competition, and top management incentives. Working Paper.
- Antón, M., Ederer, F., Giné, M., and Schmalz, M. C. (2018b). Innovation: The bright side of common ownership?

- Antón, M. and Polk, C. (2014). Connected stocks. *The Journal of Finance*, 69.
- Aryal, G., Ciliberto, F., and Leyden, B. (2018). Public communication and collusion in the airline industry. Working Paper.
- Asker, J., Collard-Wexler, A., and De Loecker, J. (2014). Dynamic inputs and resource (mis)allocation. *Journal of Political Economy*, 122(5).
- Azar, J. (2011). A new look at oligopoly: Implicit collusion through portfolio diversification.
- Azar, J., Raina, S., and Schmalz, M. C. (2016). Ultimate ownership and bank competition. Working Paper.
- Azar, J., Schmalz, M. C., and Tecu, I. (2017). The competitive effects of common ownership: Economic foundations and empirical evidence: Reply.
- Azar, J., Schmalz, M. C., and Tecu, I. (2018a). Anti-competitive effects of common ownership. *Journal of Finance*, 73(4).
- Azar, J., Schmalz, M. C., and Tecu, I. (2018b). Reply to: “common ownership does not have anti-competitive effects in the airline industry”.
- Backus, M. (2014). Why is productivity correlated with competition? Technical report, Columbia University.
- Backus, M., Conlon, C., and Sinkinson, M. (2018a). Common ownership and competition in the ready-to-eat cereal industry. Working Paper.
- Backus, M., Conlon, C., and Sinkinson, M. (2018b). Common ownership in america: 1980–2017. Working Paper.
- Bain, J. S. (1951). Relation of profit rate to industry concentration: American manufacturing, 1936–1940. *The Quarterly Journal of Economics*, 65(3):293–324.
- Banzhaf, J. (1965). Weighted voting doesn’t work: A mathematical analysis. *Rutgers Law Review*, 19(2)].
- Ben-David, I., Franzoni, F., Moussawi, R., and Sedunov, J. (2018). The granular nature of large institutional investors. Swiss Finance Institute Research Paper Series.
- Bertrand, M., Mehta, P., and Mullainathan, S. (2002). Ferreting out tunneling: An application to indian business groups. *Quarterly Journal of Economics*, 117(1)].
- Brav, A., Jiang, W., and Li, T. (2018). Picking friends before picking (proxy) fights: How mutual fund voting shapes proxy contests. Working Paper.
- Bresnahan, T. (1989). Empirical studies of industries with market power. In Schmalensee, R. and Willig, R., editors, *Handbook of Industrial Organization*, volume 2, chapter 17, pages 1011–1057. Elsevier, 1 edition.
- Bresnahan, T. F. and Salop, S. C. (1986). Quantifying the competitive effects of production joint ventures. *International Journal of Industrial Organization*, 4:155–175.
- Christoffersen, S. E. K., Geczy, C. C., Musto, D. K., and Reed, A. V. (2007). Vote trading and information aggregation. *The Journal of Finance*, 62.

- Cici, Gjergji, G. S. and Rosenfeld, C. M. (2015). Cross-company effects of common ownership: Dealings between borrowers and lenders with a common blockholder.
- Conlon, C. and Mortimer, J. (2019). Empirical properties of diversion ratios. Technical report, NYU Stern.
- Crawford, G. S., Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*.
- De Loecker, J. and Eeckhout, J. (2017). The rise of market power and the macroeconomic implications. Working Paper 23687, National Bureau of Economic Research.
- Demsetz, H. (1973). Industry structure, market rivalry, and public policy. *Journal of Law and Economics*, 16(1):1–9.
- Dennis, P., Gerardi, K., and Schenone, C. (2017). Common ownership does not have anti-competitive effects in the airline industry.
- Eekhout, J., De Loecker, J., and Unger, G. (2018). The rise of markups. Working Paper.
- Farrell, J. and Shapiro, C. (2010). The b.e. journal of theoretical economics,. *The B.E. Journal of Theoretical Economics*, 10(1).
- Fink, L. (2018). Larry fink’s annual letter to ceos: A sense of purpose.
- Freeman, K. (2018). The effects of common ownership on customer-supplier relationships.
- Gilje, E., Gormley, T. A., and Levit, D. (2018a). The Rise of Common Ownership.
- Gilje, E., Gormley, T. A., and Levit, D. (2018b). Who’s paying attention? measuring common ownership and its impact on managerial incentives.
- Gramlich, J. and Grundl, S. (2017). Estimating the competitive effects of common ownership. FEDS Working Paper 2017-029.
- Gutiérrez, G. and Philippon, T. (2016). Investment-less growth: An empirical investigation. Technical report, National Bureau of Economic Research.
- He, J. and Huang, J. (2017). Product market competition in a world of cross-ownership: Evidence from institutional blockholdings. *The Review of Financial Studies*, 30.
- He, J., Huang, J., and Zhao, S. (2017). Internalizing governance externalities: The role of institutional cross-ownership. Working Paper.
- Johnson, S., Porta, R. L., Lopez-De-Silanes, F., and Schleifer, A. (2000). Tunneling. *American Economic Review: Papers and Proceedings*, 90(2):22–27.
- Kennedy, P., O’Brien, D. P., Song, M., and Waehrer, K. (2017). The competitive effects of common ownership: Economic foundations and empirical evidence. Working Paper.
- Kwon, H. J. (2016). Executive compensation under common ownership. *Department of Economics, University of Chicago (November 29, 2016)*.
- Lewellen, J. (2011). Institutional investors and the limits of arbitrage. *Journal of Financial Economics*, 102.

- Lewellen, J. and Lewellen, K. (2017). Institutional investors and corporate governance: The incentive to increase value.
- López, A. and Vives, X. (2018). Overlapping ownership, r&d spillovers, and antitrust policy. Working Paper.
- Nevo, A. (2000). Mergers with differentiated products: The case of the ready-to-eat cereal industry. *RAND Journal of Economics*, 31(3):395–421.
- Newham, M., Seldeslachts, J., Banal-Estañol, A., et al. (2018). Common ownership and market entry: Evidence from the pharmaceutical industry. Technical report.
- O’Brien, D. and Salop, S. (2000). Competitive effects of partial ownership: Financial interest and corporate control. *Antitrust Law Journal*, 67(3):559–614.
- O’Brien, D. P. (2017). Price-concentration analysis: Ending the myth, and moving forward. Working Paper.
- Patel, M. (2018). Common ownership, institutional investors, and antitrust.
- Porta, R. L., Lopez-De-Silanes, F., and Schleifer, A. (1999). Corporate ownership around the world. *The Journal of Finance*, 54(2):471–517.
- Posner, E. A., Scott Morton, F. M., and Weyl, E. G. (2017). A proposal to limit the anti-competitive power of institutional investors. *Antitrust Law Journal*, 81(3).
- Rotemberg, J. (1984). Financial transactions costs and industrial performance. Working Paper.
- Schmalensee, R. (1989). Chapter 16 inter-industry studies of structure and performance. volume 2 of *Handbook of Industrial Organization*, pages 951 – 1009. Elsevier.
- Schmalz, M. C. (2017). Common ownership and competition: Facts, misconceptions, and what to do about it—background paper for oecd hearing on common ownership by institutional investors and its impact on competition.
- Schmalz, M. C. (2018). Common ownership concentration and corporate conduct.
- Scott Morton, F. M. and Hovenkamp, H. (2018). Horizontal shareholding and antitrust policy. Working Paper.
- Securities, U. and of Inspector General, E. C. O. (2010). Review of the sec’s section 13(f) reporting requirements. Report No. 480.
- Seldeslachts, J., Newham, M., and Banal-Estanol, A. (2017). Changes in common ownership of german companies. Technical report, DIW Economic Bulletin.
- Shambaugh, J., Nunn, R., Breitwieser, A., and Liu, P. (2018). The state of competition and dynamism: Facts about concentration, start-ups, and related policies. Technical report, The Hamilton Project: Economics Facts.
- Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy*, 112(6).
- Werden, G. (1996). A robust test for consumer welfare enhancing mergers among sellers of differentiated products. *Journal of Industrial Economics*, 44(4).

Table 3: Empirical investigations of the Common Ownership Hypothesis

Paper	Context	Instrument	Outcome Variable	Effect of CO	CO Harms Competition?
<i>Initial Evidence</i>					
Azar (2011)	Public Firms	None	Markups	Positive	Yes
Azar et al. (2016)	Banks	Cross-Section Variation	Fees, Interest Rates	Positive	Yes
Azar et al. (2018a)	Airlines	Blackrock-Barclays Merger	Fares	Positive	Yes
Cici and Rosenfeld (2015)	Syndicated Loans	Discontinuity in Holdings	Transaction Likelihood	Positive	Yes
Freeman (2018)	Compustat Supply Chain	Mutual Fund Scandal	Relationship Ending	Negative	Yes
He and Huang (2017)	Public Firms	Bank Mergers	SIC Market Share	Positive	Yes
<i>Responses, Replies</i>					
Kennedy et al. (2017)	Airlines	Blackrock-Barclays, Russell 1000	Prices	No Effect	No
Dennis et al. (2017)	Airlines	Airport Market Shares	Prices	No Effect	No
Azar et al. (2018b)	Airlines	(Reply to Dennis et al. (2017))	-	-	-
Azar et al. (2017)	Airlines	(Reply to Kennedy et al. (2017))	-	-	-
<i>Other Manifestations</i>					
Antón et al. (2018a)	Managerial Incentives	Blackrock-Barclays	Manager Wealth	Steeper	Yes
Antón et al. (2018b)	Innovation	None	Patents	Positive	No
Newham et al. (2018)	Generic Drug Entry	Patent Timing	Entry	Negative	Yes

Table 4: Discussions, Criticisms, Recommendations of the Common Ownership Hypothesis

Paper	Context	Notes
<i>Documentation</i>		
Backus et al. (2018b)	S&P 500	Document Rise of Profit Weights 1980–2017
Seldeslachts et al. (2017)	German Companies	Document increases in Common Ownership
<i>Criticisms and Alternative Approaches</i>		
Gramlich and Grundl (2017)	Banks	Suggest profit weight instead of MHHI.
O'Brien (2017)	Theory	No rationale for regressions of prices on any measure of concentration.
Schmalz (2017)	Theory	Compares industry talking points to data.
Gilje et al. (2018a)	Mechanism	Derives alternative measure of CO effect from inattention.
Lewellen and Lewellen (2017)	Management Fees	Shows that investment managers have incentives to reduce competition.
López and Vives (2018)	Theory	Shows that common owners may incentivize innovation.
Backus et al. (2018a)	Methods	Structural testing for CO effect in cereal market.
<i>Policy Recommendations</i>		
Scott Morton and Hovenkamp (2018)	Policy	Current antitrust laws would allow enforcement action.
Posner et al. (2017)	Policy	Proposes restrictions for institutional investors
Patel (2018)	Policy	Argues for case-by-case analysis.

B Economic Studies

at BROOKINGS

The Brookings Economic Studies program analyzes current and emerging economic issues facing the United States and the world, focusing on ideas to achieve broad-based economic growth, a strong labor market, sound fiscal and monetary policy, and economic opportunity and social mobility. The research aims to increase understanding of how the economy works and what can be done to make it work better.

Questions about the research? Email communications@brookings.edu.
Be sure to include the title of this paper in your inquiry.