Theory and Measurement of Common Ownership

By Matthew Backus, Christopher Conlon and Michael Sinkinson*

The theory of common ownership posits that investors, by taking non-controlling ownership stakes in competing firms, effect a partial merger. Azar, Schmalz and Tecu (2018) documented empirical evidence for this claim in the US airline sector, evidence that inspired a growing and controversial empirical literature on common ownership. Backus, Conlon and Sinkinson (2019*a*) reviews that literature. Here, we focus on the measurement of common ownership itself. Differing approaches to measurement drive controversy in the growing common ownership literature, as researchers describe historical patterns, attempt to test the predictions of the model, and use it to generate counterfactual predictions.

We consider three approaches, which are sequentially nested by modeling structure. The first is descriptive; it measures ownership patterns and the extent to which investors overlap between firms. The second maps ownership into primitives of the managers' objective functions, which are assumed to aggregate the preferences of their investors. The third maps these primitives into equilibrium outcomes of specific strategic settings, resulting in measures such as the "Modified Herfindal-Hirschman Index" (MHHI) of Bresnahan and Salop (1986).

I. Measuring Investor Overlap

In the United States, institutional investors with over \$100M in assets are required to file quarterly 13(f) forms with the US Securities and Exchange Commission (SEC) listing publicly traded securities. These filings measure ownership at the level of a legal entity, which may or may not correspond to the level at which decisions are made. These filings also suffer from additional shortcomings: short position investments are not distinguished from long ones; the entity that controls voting rights is often ambiguous; dual-class shares complicate (or sometimes obviate) investor influence for some firms; and reporting errors are common.

The commonly-used Thomson Reuters database of 13(f) filings introduces a number of additional errors and coverage issues, which we document in Backus, Conlon and Sinkinson (2019*b*). For the period 2000 to the present, we scraped data directly from the SEC and have made our data available to the public.

Suppose each shareholder $s \in S$ has a portfolio in which they own a fraction of firm $f \in F$ denoted by β_{fs}. Measurement of common ownership, then, is finding ways to characterize the potentially large $F \times S$ matrix β, which summarizes holdings across all shareholders and firms. A challenge of the descriptive approach is that one must choose among the arbitrarily many ways to reduce β, a high-dimensional object, to a reportable statistic. It is possible to correlate ownership statistics of the form $f(\beta)$ with various outcomes, but economically meaningful claims require placing additional structure on the problem.[1]

II. Firm Objectives: Profit Weights

With two additional assumptions, the theory of common ownership maps overlapping ownership positions as measured above

* Backus: Graduate School of Business, Columbia University, 3022 Broadway, Uris Hall 611, New York, NY 10027, CEPR and NBER, matthew.backus@columbia.edu. Conlon: Stern School of Business, New York University, 44 West 4th St, Kaufman Management Center 7-76, New York, NY 10012, cconlon@stern.nyu.edu. Sinkinson: School of Management, Yale University, 165 Whitney, Evans Hall Suite 3473, New Haven, CT 06511, and NBER, michael.sinkinson@yale.edu. Any errors are our own.

[1] As an example of this approach, He and Huang (2017) count the number of common blockholders who own $\beta_{fs} \geq .05$ and $\beta_{gs} \geq .05$ in both firms and correlate these measures with growth in market share.

into primitives of a firm's objective function.

ASSUMPTION 1: *Investor returns/ portfolio values are given by:* $v_s \equiv \sum_s \beta_{fs} \pi_f$.

ASSUMPTION 2: *Managers maximize a* γ_{fs} *weighted average of investor returns:* $Q_f \equiv \sum_s \gamma_{fs} v_s$ *(Rotemberg 1984).*

The first assumption defines investor portfolios as the sum of corresponding cash-flow rights $\beta_{fs} \geq 0$, multiplied by the value of each firm π_f.

The second assumption is more controversial and states that managers maximize a weighted average of their investor payoffs. Because investors hold heterogeneous portfolios, they may disagree about their preferred objective for the firm. Managers aggregate preferences of heterogeneous investors using a set of Pareto weights γ_{fs}.

Under these two assumptions one can rearrange the manager's objective such that:[2]

$$Q_f \propto \pi_f + \sum_{g \neq f} \kappa_{fg} \cdot \pi_g, \tag{1}$$

$$\text{where} \quad \kappa_{fg} \equiv \frac{\sum_{\forall s} \gamma_{fs} \beta_{gs}}{\sum_{\forall s} \gamma_{fs} \beta_{fs}}.$$

This implies that managers maximize their own profits π_f plus some κ_{fg} weighted sum of the profits of other firms π_g. These κ_{fg} terms are known as *profit weights* and have a long history in economics.[3]

The Pareto weights γ_{fs} stand in for the influence of investors on firm decisions (e.g. corporate governance). Absent an assumption on γ_{fs} the expression in (1) is sufficiently general to accommodate a host of behaviors. For example, the manager m might place weight on his own private benefit π_m with $\gamma_{fm} > 0$ and potentially ignore his investors completely: $\gamma_{fs} = 0; \forall s \neq m$. Alternatively, the manager may place equal weight $\gamma_{fs} = \gamma_{fs'} > 0$ on his largest two shareholders and ignore the rest; or place equal weight on all shareholders $\gamma_{fs} = c; \forall s$.

Nearly any model of corporate governance (with or without agency frictions) can be written using Assumptions 1 and 2.[4] In other words, the controversy arises from the specific choice of γ and not Assumption 2 itself. Unfortunately, there is little guidance from the corporate governance literature about how to measure or specify γ.

One might be inclined to estimate γ from data on market outcomes. As a general problem this is somewhat futile, as there are often many more investors S (several thousand) than there are firms F (a handful) in an industry. Even if we knew the profit weights κ, we would not be able to recover the Pareto weights γ.

The empirical literature proceeds by assuming $\gamma = \beta$, or "proportional control." In Backus, Conlon and Sinkinson (2019*b*) we consider a generalization to $\gamma = \beta^\alpha$ for $\alpha \in \{1/2, 1, 2, 3\}$, a parameterization that offers some flexibility in the relative influence of large and small investors. Considering the period 1980–2017, we found convergence in the averages of pairwise κ across these specifications – by 2017 there is little difference – however α does matter for measuring the frequency of extreme values (e.g., $\kappa_{fg} > 1$).

With an assumption on γ, the primitives of the manager's objective function are fully specified, and those primitives may vary over time with changes in the observed ownership β. This variation across time and pairs of firms provides a way to compare different assumptions on γ.

However, it is difficult to map κ to market outcomes without making additional assumptions on the nature of interactions between firms (e.g. that they are horizontal competitors engaged in selling substitutes). Absent these assumptions, it may be possible to develop reduced form, correlation-based tests, even though the magnitudes

[2]This expression first appeared as such in O'Brien and Salop (2000).

[3]Dating as far back as Edgeworth's "coefficients of effective sympathy".

[4]For example, the measure of common ownership and investor attention proposed by Gilje, Gormley and Levit (2019) can be shown to be mathematically equivalent to our Assumptions 1 and 2 under a formulation of $\gamma_{fs}(\beta_s)$ that places less weight on investors as they become more diversified. They rule out strategic interactions among firms, which makes their measure inappropriate for examining the common ownership hypothesis.

of coefficients may not be interpretable. For example, Gramlich and Grundl (2017) don't find a strong direct relationship between prices and functions $f(\kappa)$ in the market for retail banking (a setting similar to Azar, Raina and Schmalz (2016)). These reduced-form tests may not be as simple as they look. For example, O'Brien (2017) points out that equilibrium outcomes depend not only on κ_{fg} but on the entire $F \times F$ matrix κ; the precise relationship depends on the form of the strategic game.

III. Fully-Specified Strategic Games

If one begins with the firm's objective function from (1) and fully specifies both the Pareto weights γ and the form of the strategic game played among all firms F, it is possible to derive relationships between common ownership and equilibrium outcomes such as prices, quantities, investment, and entry or exit.

Perhaps the most common example in the literature is to assume that the firms engage in symmetric Cournot competition (simultaneous Nash-in-quantities) so that $Q_f(q_f, q_{-f}) = \pi_f(q_f, q_{-f}) + \sum_{g \neq f} \kappa_{fg} \cdot \pi_g(q_f, q_{-f})$. When one solves for the first-order-conditions of the resulting game, it is possible to derive a relationship between share-weighted average markups and the MHHI:

$$\sum_f s_f \frac{p_f - c_f}{p_f} = \frac{1}{\epsilon}\left[MHHI(\kappa)\right], \tag{2}$$

$$MHHI(\kappa) = \underbrace{\sum_f s_f^2}_{HHI} + \underbrace{\sum_f \sum_{g \neq f} \kappa_{fg} s_f s_g}_{\Delta MHHI}. \tag{3}$$

$MHHI(\kappa)$ is not a measure of common ownership, it is a modified concentration index and an equilibrium outcome itself. Scholars sometimes portray $MHHI(\kappa)$ and the profit weights κ as different ways to measure common ownership, although they are not. The profit weights, κ, are a primitive object in the manager's objective function; while $MHHI(\kappa)$ or $\Delta MHHI(\kappa)$ are equilibrium outcomes of a Cournot game where the market shares depend on κ.

If the strategic game is something other than symmetric Cournot, there need not be any relationship between $MHHI(\kappa)$ and equilibrium outcomes. For example, if the strategic game is Bertrand Nash-in-prices in differentiated products so that Q_f is a function of p instead of q, then $Q_f(p_f, p_{-f}) = \pi_f(p_f, p_{-f}) + \sum_{g \neq f} \kappa_{fg} \cdot \pi_g(p_f, p_{-f})$. This gives a different first-order condition:

$$p_f = \frac{\epsilon_f}{\epsilon_f - 1}\left[\underbrace{c_f + \sum_{f \neq g} \kappa_{fg}\, D_{fg}\,(p_g - c_g)}_{PPI}\right]. \tag{4}$$

The bracketed expression is the Price Pressure Index or $PPI(\kappa)$: This relates prices to the own elasticity of demand ϵ_f, and substitution to rival's products as measured by a "diversion ratio" D_{fg} (O'Brien and Salop 2000).

Backus, Conlon and Sinkinson (2018) show that mis-specifying the form of the game can lead to spurious results. For example, regressions of prices on $MHHI(\kappa)$ can yield spurious positive or negative results when the strategic game is actually Bertrand and there is no effect of common ownership.

Testing the theory of common ownership by regressing prices on $MHHI(\kappa)$ leads to additional challenges. First, the implied relationship in (2) is between share-weighted average markup and $MHHI(\kappa)$, not prices and $MHHI(\kappa)$. Second, any analyses using $MHHI$ require the researcher to compute market shares in addition to κ, thus introducing the myriad difficulties of proper market definition.[5] The $PPI(\kappa)$ is hardly better, as it requires estimates of the diversion ratios D_{fg}. Finally, because $MHHI(\kappa)$ is a market-level measure, it is unable to exploit variation across firms.

[5]This is further complicated when researchers construct market shares from publicly available databases such as COMPUSTAT which are limited to publicly traded US firms. The $MHHI(\kappa)$ measure implicitly requires the appropriate geographic and product market as all products must be equally good substitutes within the relevant market.

IV. The Role of Measurement

The descriptive approach to measuring common ownership is limited by the lack of interpretation. For this reason, Backus, Conlon and Sinkinson (2019*b*), which measures common ownership in the US from 1980 to 2017, advocated for a focus on κ, the objective function of the firm. Because it is generic to the formulation of firms' interaction, there is no need for compromises on market definition.

In Backus, Conlon and Sinkinson (2019*b*) we also show that the key to testing the theory of common ownership hypothesis is really the profit weights κ. While one might learn about κ through equilibrium outcomes like $MHHI(\kappa)$ or $PPI(\kappa)$ or other outcomes such as entry, R&D, or investment, testing common ownership — like testing collusion — is about the profit weights firms put on each other.

These reflections on measurement are critical to structural testing, but can also guide reduced form work. For example, seemingly innocuous financial market transactions could have potentially large impacts on product markets through κ. Boller and Scott Morton (2019) provide some encouraging evidence here. They find that when firms join stock indices, there are pricing anomalies for the stock of rival firms. These pricing anomalies are correlated with the theoretically-motivated κ — consistent also with asymmetries of κ between firms — but not other, purely descriptive overlap measures.

REFERENCES

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

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

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2018. "Common Ownership and Competition in the Ready-To-Eat Cereal Industry." Unpublished.

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

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2019*b*. "Common Ownership in America: 1980–2017." NBER Working Paper 25454.

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

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

Gilje, Erik P., Todd A. Gormley, and Doron Levit. 2019. "Who's paying attention? Measuring Common Ownership and its Impact on Managerial Incentives." *Journal of Financial Economics.*

Gramlich, Jacob, and Serafin Grundl. 2017. "Estimating the Competitive Effects of Common Ownership." FEDS Working Paper 2017-029.

He, Jie, and Jiekun Huang. 2017. "Product Market Competition in a World of Cross-Ownership: Evidence from Institutional Blockholdings." *The Review of Financial Studies* 30.

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

O'Brien, Daniel P. 2017. "Price-Concentration Analysis: Ending the Myth, and Moving Forward." https://ssrn.com/abstract=3008326. Accessed on 2020-01-08.

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