

- Solow, R. 1957. Technical change and the aggregate production function. *Review of Economics and Statistics* 39 (3): 312–20.
- Tomer, J. F. 1987. *Organizational capital: The path to higher productivity and well-being*. New York: Praeger.

Comment Timothy F. Bresnahan

The goal of this paper is an important one. The authors set out to define “organizational capital,” to measure it, and to see whether capital markets properly value it in large companies. The importance arises from several conjectures that I and many others share with the authors. Few doubt that the organization of firms is among the determinants of productivity and productivity growth. Few doubt the value of a set of systematic measures in this area. Few doubt that reorganization of firms arising from use of information technology will leave customers better off and successful firms with higher market shares. And finally, few doubt that the widespread assertion of a “new economy” form of organizational capital gave capital markets an extraordinarily difficult valuation problem in the late 1990s. All these conjectures lead us toward goals like those in this paper.

The paper’s methods to measure organizational capital center on production function residuals at the firm level. The authors focus on large, publicly traded firms. In the first approach, a Cobb-Douglas production function, with R&D as one of the inputs, is estimated in first differences with firm dummies. The residual is cumulated and called organizational capital.

In this approach a high rate of growth of firm sales, above and beyond growth in employment, physical capital, and R&D capital, measures the rate of growth of organization capital. The firm-specific effects are modeled as moving over time. At year t , the measure of organizational capital growth comes from the model run on years $t - 4$ through t . The next year, the measure of organizational capital growth comes from the model run on years $t - 3$ to $t + 1$, and so on.

In a second approach, the firm’s sales, general, and administrative costs (SGA) are taken to be a measure of expenditures to gain organizational capital (OC). A production function more general than the Cobb-Douglas is estimated; all the Cobb-Douglas parameters are allowed to vary with SGA. The measure of OC comes from comparing predicted sales at $SGA = 0$ to the actuals. Thus, in the second version, a productivity residual is projected onto SGA and the interaction of SGA with other inputs and called OC.

Timothy F. Bresnahan is the Landau Professor of Technology in the Economy at Stanford University and a research associate of the National Bureau of Economic Research.

With those two sets of estimates in hand, the paper examines several economic questions. The authors seek to document the size and growth of measured OC in the class of firms they study. They investigate whether measured OC is correlated with any of several variables associated with theories of organizational improvement. These include market share, computer investment, and administrative expenses. Finally, the authors examine the stock market valuation of measured OC using standard finance methods. They are interested in both the degree to which the stock market values companies that have high measured OC and the question of whether the stock market over- or undervalues such companies.

The particular methods choices determine the empirical content of the measures of OC. Much trouble arises because it is difficult to become convinced that production function residuals at the firm level measure OC.

I will begin with a discussion of what we have learned from earlier research that uses methods related to this paper. I will then turn to the main results and their interpretation. Potential worries about whether the object that has been measured might be precisely something called “organizational capital” grow as we examine the authors’ tests of organizational theories, their growth accounting, and their capital market investigations. I will go on to talk about the conceptual and data limitations that are holding back this entire field, not just this paper.

Methods I: Residuals

The most important modeling choice in the paper is to measure OC either as a production function residual at the firm level or as the portion of a production function residual that is explained by variation in SGA.

In general, estimating production functions is not a happy task. It is very difficult to succeed in that effort either at the aggregate level or at the individual firm level. At the aggregate level, it is difficult to get the output deflators right, so that real output—the dependent variable in the production function—is always suspect. At the firm level, that particular problem is somewhat ameliorated. Firms that have some hard-to-measure competitive advantage, such as better products, will tend to have larger market shares in industry equilibrium. They will thus tend to have larger nominal sales. To the degree it comes from a superior production technology that is uncorrelated with other factors of production, production function estimates will capture it.

That is the basic logic under which this paper might be right. Organization capital is likely to be reflected in superior products, especially in the service sectors. Thus, looking at a firm-level productivity residual may well capture OC if it is there.

The downside of working at the firm level is that firms are highly heterogeneous. Unobserved heterogeneity in firm-level productivity or demand makes estimating the production function difficult. Even with panel data on

firms, it can be hard to learn the firm-level production function, especially if one is hoping to tell apart selection, endogeneity in the inputs, and true advantages to the firm in better productivity.¹ The firm-level productivity residual can measure, inter alia, unobserved demand and supply movements at the firm and industry level and the firm's response to those shocks.

This is the downside of the logic behind the measurement strategy of this paper. The firm-level production function residual may well contain the effects of OC if it is there, but it will also contain everything else about how firms are different, or in different circumstances. The fundamental idea of the paper is a *relative* one. Those firms with more output growth per unit input growth relative to other firms are labeled as high-measured-OC firms.

On balance, I suspect firm-level analysis is likely the least bad for the purposes of this paper, which is about firms, but it is important to recall what it can and cannot do. The paper's calculations, which make no effort to control for the most familiar measurement problems in this area, are likely suspect.

Methods II: Organizations

The second fundamental point to make about methods relates to the measurement of organizations. The paper does not measure firm organization.

There are no measures of hierarchy or its absence; no measures of centralization or decentralization; no measures of the degree to which workers are on dynamic incentive contracts; no measures of the divisional structure of the firm; no measures of the workplace organization in productive units; no measures of corporate culture; and no measures of the role of management. In short, there are no measures suggested by the theory or empirical literatures in empirical or theoretical organizational behavior or the economics of organizations.²

The paper does use the firm's expenditures on SGA as a predictor of the production function residual. The implicit story is that a firm that spends more money on administration (the "A" in SGA) is organizing itself better.

For a number of reasons, the dollars spent by a firm on SGA have only a very distant linkage to "organizational capital" in the sense of the paper. The most direct point is that a badly organized firm may need more managers, not fewer. A separate point is that SGA includes marketing expenditures as well as managerial ones.

1. While these arguments are longstanding, perhaps the clearest sense of the trade-offs can be found in Griliches and Mairesse (1998). That paper reports an interesting effort to deal with the problem, as does one by Olley and Pakes (1996), with a nonparametric selection model.

2. A summary of the relevant theory can be found in Milgrom and Roberts (1992). Recent empirical and theoretical developments were reported at the NBER Organizational Economics Conference: see <http://www.nber.org/~confer/2002/orgec02/program.html>.

There are also serious econometric issues associated with SGA. Of course it is possible that the SGA coefficients measure managers' organizational contribution. To the extent that organizing the firm is costly and these costs are reflected in SGA, that would be the right interpretation. There is a problem with the implicit identifying assumption in that (most positive) case. If firms with higher SGA are more productive, are firms with lower SGA making mistakes?

More serious problems arise when either SGA is not the only driver of improved organization or it is not given by some exogenous process. Suppose that there are organizational improvements that do not lead to more expenditure on SGA, such as improvements in workplace organization, streamlining bureaucracy, improvements in the quality of decisions of their communications, and so on. The paper will not measure these.

The paper will, incorrectly, measure correlations between SGA and the firm residual that are not causal. If firms with looser budgets spend more on management salaries, for example, that will be measured as OC in this paper. Thus, an uptick in sales, part of which becomes management rents, is called OC.

My real point is that there is only so much that can be done with production function residuals. This paper labels a production function residual as OC. Sometimes, it uses only the part of the residual that can be predicted by SGA, or SGA interacted with the factors of production. It is always using a productivity residual, however.

An alternative approach would be to link outcomes at the firm level empirically to data that actually measure something about organizations. There are many approaches to that (see note 3). In recent work related to this paper, Erik Brynjolfsson, Lorin Hitt and I (Bresnahan, Brynjolfsson, and Hitt, 2002) have tested the hypothesis that investments in information technology are complementary to workplace organization. In our work, the organizational variables we use are measures of the way production workers are connected to the firm—how they are supervised, whether they work in teams, the nature of their in-house training, and so on. In my view, much of the reason to believe the conclusion that the results are specifically organizational depends on the use of these data. We are hardly alone in this. For example, Cockburn and Henderson (1998) have interesting measures of the organization of laboratories in pharmaceutical firms, related to the firms' ability to absorb scientific knowledge and turn it into profitable products.

One could go on in this vein for some time. Recent research has made some progress in measuring organizational features at the firm and industry level and relating them to economic outcomes, but there is still much work to be done. Much of the most promising work has used case study or historical methods. Statistical methods have crashed up against fundamental data limitations. This is not a problem of measuring residuals.

Scholars who seek to measure organizational effectiveness using data on the structure of organizations, on incentives within the organization, on management, and so on—all the things that would be the beginnings of empirical correlates of organizational capital—do not typically use the same data used to study national economic growth. Indeed, most of the research literature in organizations goes forward without using any of the data resources of the federal statistical system at all, but is rather based in the hard work of scholars in business schools.

The issue is not, I hasten to add, that the people working in the federal statistical system are unaware of the changes in the organization of firms and industries that matter for modern economic growth.³ Nor is the problem that the federal statistical system fails to understand the importance of making data resources available to researchers for purposes of generating new knowledge. Like me, they recognize the value of a thriving empirical economic research program for future economic policymaking. The issue is simply one of resources.

Results

While the paper has a number of subsidiary results, the main results are as follows:

1. Measured organizational capital is large and makes a substantial contribution to the growth of firm output.

The size of measured OC is different in two methods used by the authors, but it is in either case a substantial fraction of physical or R&D capital for the firms they study.

2. Measured OC is persistent at the firm level.

The paper reports the persistence of its measures of OC and finds that it is high. The paper also examines the underlying production function residuals themselves, finding them to be highly persistent. Thus, the persistence of measured OC is not an artifact of the particular time series process assumed by the authors.

3. Measured OC is correlated in the cross section with

- a. Firm market share in its two-digit SIC industry
- b. Firm expenditures on information technology
- c. SGA (for the first OC measure)

4. Value in the stock market at the moment the authors' OC measure could be calculated.

5. Stock market gains *after* their OC measure could be calculated.

3. The federal statistical system has recently undertaken a very wise overhaul of the way it classifies establishments, of which an important part is recognizing some organizational issues.

These last two results come from a standard financial markets valuation regression. The finding that OC predicts firm value is in a regression in which several other assets and analysts' forecasts are held fixed.

Interpretation of Results

The authors have made the tables, and any discussion naturally begins with their interpretation. It should not end there, however, as many of the substantive interpretations we find in this paper involve further unverified assumptions.

Findings 1 and 2, substantial measured organizational capital and substantially persistent measured OC, seem to me to be better labeled "assumptions" than "results." The paper reports that the contribution of organizational capital to the growth of sales is large. Since OC was measured as a residual (or a residual projected onto SGA) from a first-differenced production function, the contribution of OC to growth in output is an identity, not a result.

Similarly, the finding that measured OC is persistent at the firm level is the finding that a firm production function residual is persistent at the firm level. That well-known fact establishes that firms are different in ways that move slowly, not necessarily that the reason for that is OC.

The three subfindings under result 3 are interpreted by the authors as evidence that their economic interpretation is right. I am less convinced.

Findings 3a, 3b, and 3c establish that the production function residual is correlated with market share, IT budgets, and SGA. I have already talked about the SGA correlation. The interpretation the authors offer for the first two correlations is that improvements in OC improve firm competitive position and that OC and IT are complements. I agree with both of those economic conclusions but do not see any evidence for them in this paper.

The finding about market share seems to me to be simply a restatement of what the paper does. Positive production function residuals arise when a firm's sales rise more rapidly than its inputs. A firm will have a larger value of those residuals when, relative to other firms, its sales/input is rising rapidly. That could happen because the firm has grown more productive through a positive improvement in its organization, the authors' preferred interpretation, or because there has been an outward shock to the firm's demand for any reason, which would cause output and measured productivity to rise. There is a substantial literature that attempts to sort out these two different flows of causation in order to identify the part of the firm-level productivity residual that is in fact a shock to the production function. This paper interprets the entire correlation as productivity.

The finding about IT is one of a long series of findings that IT budgets are correlated with success at the firm level. The work of Erik Brynjolfsson, alone and with collaborators, has established this fairly clearly. In work that I did with Erik and Lorin Hitt, cited above, the correlation is very care-

fully unpacked to learn how much of it is related to improvements in organization that lead to productivity improvements and how much reflects other flows of causation. There is no similar effort in this paper.

In short, I agree that improvements in organization at the firm level, often in recent years connected to IT, are an important area for productivity growth. What I don't see in this paper is an attractive measure of that or any evidence for it. The paper's complaint that "there exist no operational measures of firms' organizational capital" is still right.

Capital Markets

The paper goes on to show that high-measured-OC firms are (finding 4) more valuable in the stock market and (finding 5) get predictable stock-market gains. The authors conclude that investors reward firms that have invested in OC before the returns on that capital are visible (4) but undervalue this capital (5).

The finding that high-measured-OC firms are more valuable in the contemporary stock market (4) follows from the measurements. Any economic interpretation of measured OC would likely lead to finding 4—including the authors' interpretation, but also including any other.

Finding 5 is more interesting. Whenever stock market returns can be predicted in a regression, there is something to pay attention to. The undervaluation conclusion follows from the empirical finding that returns in stock prices can be econometrically predicted using the measure of OC.⁴

I am dubious about the economic conclusion. My doubt has its roots in two arguments, one about the sample period and the other about measured OC.

The sample period begins in the 1990s (for this exercise) and ends in 2000. That is a period of surprising sustained macroeconomic expansion. The capital markets appeared to base valuation of companies in this era at least in part on the view that the aggregate economy would continue to have surprising growth, and so that growth companies would continue to have good prospects. This is not an idea that was quite so important in the capital markets either earlier or, as we now know, later.

In general, this was a good period in which to have been holding equities in firms that grew with or faster than the aggregate economy. A bull market, in short. The authors' measures of OC are measures of relative company growth. When company sales per unit input rise, measured OC will rise.

So the evidence in this paper that there is a "misevaluation of organizational capital on the stock market" (or the need for a new risk factor) is this. Holding a portfolio of growth companies during the 1990s would have led

4. An alternative interpretation the authors offer is that this finding might point us toward "a new risk factor associated with organizational capital." I will treat that as part of their main interpretation.

to extraordinary gains. It seems to me that there is a long step from that evidence to the conclusion that the capital markets undervalue “organizational capital.”

Those remarks merely report what I hope is standard applied finance common sense. You could “show” by these methods that many different assets were “undervalued” in the 1990s. U.S. equities in general, for example, could be “shown” to be undervalued in that period by the methods used in this paper. Anyone who held U.S. equities through that time period earned good returns. The problem is, you are looking at a long bull market. You always do well buying stocks before a bull market. You also will always do particularly well buying high-measured-OC stocks before a bull market.

There are many, many plausible ways in which the U.S. stock market was getting valuation wrong in the 1990s. A bubble in which all equities were overvalued is one plausible way, and a bubble in which growth stocks were overvalued is another plausible way. I think that these authors have some more work to do before they make a convincing case that U.S. equities markets were seriously *undervaluing* growth companies in the 1990s.

Two Industries’ and Four Firms’ “OC”

The paper names about a dozen leading firms in several industries as widely cited examples. It attributes the sustained success of these firms to OC, conceptualized as a factor of production unique to these firms within their industries. I know two of the industries from which these firms are drawn, motor vehicles and computing, well enough to use them to consider the broader issues raised by the firm-level OC approach.

In automobiles, there is cross-section variation in the effectiveness of firms in managing design, production, and distribution processes. One important innovation here is undoubtedly organizational, at least in part. I refer to the Toyoda system (sometimes called Kanban or “lean production”) of manufacturing automobiles. The authors link this organizational to Toyota Motor Company’s current market position. That is largely right, but to understand the broader economic issues related to that calls for adding two pieces of context.

First, half a century ago when Eiji Toyoda, Ed Deming, and (especially) Taiichi Ohno were working on this organizational improvement, Toyota Motor Company was not the automobile market leader. Back then, General Motors (GM) was the acknowledged leader in “organizational capital” in that industry, based on the (also undoubtedly organizational, at least in part) innovations of Alfred P. Sloan at GM.⁵ Dramatic improvements in “organizational capital” may create huge gains to society, as they did in this case—twice. General Motors created a new form of organization to leapfrog market leader Ford; Toyota in turn created a new form of organization to leapfrog GM.

5. See, e.g., Womack, Jones, and Roos (1991).

The point is the basic one of creative destruction. Society's stock of OC does not necessarily come from existing dominant firms, but instead from the dynamic incentives of entrepreneurial outsiders. Society has an interest in the organization of markets as well as of firms. Markets in which creative destruction can play out create new OC—by destroying the economic value of old OC.

Second, measuring the stock of OC by comparing leading firms to others is a dangerous business. Successful imitation by other automobile firms lowers Toyota's measured OC but not its actual OC. It increases the social stock of OC. To the extent that prices fall from successful imitation and competition, the imitating firms do not have any measured OC. Similarly, GM's measured OC fell as other firms (Ford, Chrysler, and the ancestors of American Motors) imitated its organization, and fell further when new competitive organizations like Toyota entered.

Again, the basic conceptual point comes from the economics of competitive markets. There is no simple and direct relationship between firm performance and industry performance. Society's stock of OC cannot be calculated by adding up the stocks of firm OC as measured in this paper.

The authors bring up two examples of firms from the computer industry where they see "organizational capital" as the source of sustained leading positions: IBM and Microsoft. This is partly right, and it raises a new issue. Both of those firms originally *built* leading positions in that industry through superior performance.⁶ IBM's superior performance clearly had organizational roots, as Tom Watson, Jr., built a firm that could both invent and market large-scale computers. Microsoft's early superior performance was more related to the brilliance of the entrepreneurs who founded it, Paul Allen and (especially) Bill Gates.

The long-term success of both of those firms has other causes, however. Each had a long run as the dominant firm in its industry (Microsoft's continues today). The idea that something you might call "organizational capital" *sustained* either firm in its long period of industry dominance is false. It was not superior performance that kept IBM in a position of leadership through the 1970s and 1980s, nor is it superior performance that keeps Microsoft in a position of leadership today.

Long after its organization had become a liability rather than an asset, IBM held on to market dominance in the mainframe computer industry. The advantages that sustained it were positional, not organizational. Years later, when a long, slow process finally removed IBM from its prior position, the importance of competition and ideas from outside was universally acknowledged, even from ex-IBMers. While slow, the process of creative destruction ultimately worked.

In the present, Microsoft's dominant position in the personal computer (PC) industry does not arise from superior organization. When Internet

6. See Bresnahan and Malerba (1999) for more detail.

entrepreneurs came along, they outinvented and outmarketed Microsoft—just as the young Microsoft outperformed IBM in the early days of the PC. Microsoft continued its dominance by blocking distribution of competitively threatening technologies.⁷ Where the fall of IBM is an example of successful creative destruction, Microsoft's ability to evade competition is an example of blocked creative destruction.

The broad point is that “what's good for GM is good for the country” is nonsense. What's good for the continued growth of the United States and of the world economy is the competitive system that permits creative destruction. Current public policy in the United States has permitted Microsoft to evade even the threat of creative destruction from new competitors born on the Internet. Current public policy is leaning strongly toward putting the “baby bells” back in the position AT&T once enjoyed. When public policy faces decisions like that, it is important to remember the distinction between the economy's stock of OC, which includes dynamic competition, and leading firms' stock of measured OC, which is larger when there is less dynamic competition.

The paper by Lev and Radhakrishnan focuses on productivity growth in large, existing firms, especially ones that are relatively successful in their industry. The paper makes an argument in favor of focusing exclusively on that topic in the study of productivity growth in the whole economy. The exclusivity is a mistake; the error made by this paper in arguing for exclusivity is a form of an increasingly common error.

Productivity, Economic Growth, and Organization

Is individual firm OC the growth pole of the U.S. economy? Is sustained firm-level success caused by superior organization pervasive in the economy, observed in almost every industry? Is firm-level OC “the major source of competitive advantage” that we should be careful to value? Is the central locus of this OC in large, established, successful firms? With or without the “organization capital” label, the view that all these questions should be answered “yes” is growing more common in discussion of economic policy formation. That view makes an error at two levels.

The view omits most of the forces that matter to productivity growth at the economywide level. It leaves out fundamental advances in science and engineering, for example, and their later conversion into commercially valuable products and processes. It leaves out firm growth, birth, and death, and the importance of market selection in productivity growth. Perhaps most important, it leaves out entrepreneurship and the possibility of creative destruction.

At a second level, that view makes an important conceptual error. Around 1970, the world looked to firms like IBM, GM, and AT&T as ex-

7. While these claims are controversial, they should not be. They are what the Microsoft documents showed and what the courts, both in trial and on appeal, have found.

emplars of excellence in organization. Public policy listened to the idea that they were an excellent form of organization, as did business policy. Later, in the 1980s, many studies of productivity growth looked to international comparisons. The higher rate of productivity growth in Japan than in the United States led to the advice that business and public policy should emulate Japanese *keiretsu*. My point is that closely linking analysis at the firm level to analysis at the national or world economic growth level is a common mistake.

Indeed, the single most common usage of the phrase “new economy” that forms part of the title of this conference was “new economy company.” That meant something very particular in the firm valuation theory that drove the stock market bubble of the late 1990s. It meant that firm-level organizational capital of a new form was the engine of economic growth. A particular form of startup was supposed to be the key exemplar of superior performance.

And now we are returning to the view that large-scale successful enterprises have valuable lessons for public policy and business policy—the view of this paper and much other recent analysis. Like the earlier version of that view (1970) or the view that *keiretsu* were the way to go (1980s) or that overcaffeinated startups had all the organizational capital (1990s), this new view is shortsighted.

The error in all four analyses is the same. The error is *not* the specific one that we paid “too much attention” to large established firms in the 1960s or today or that we paid “too much attention” to startups in the late 1990s or “too much attention” to Japan in the 1980s. The mistake lies in thinking that it is a sensible approach to productivity growth analysis to have a system where we “pay attention” in this sense to anything at all. The view in which scholarship has the job of figuring out what kind of firm organization is a good kind, and business and public policy have the job of “paying attention” to the scholarship and thus to the form of firm, is dangerous nonsense. It reduces economics to a kind of second-rate industrial engineering, and injects an element of central planning into policy formation.

Ignoring markets is a serious oversight, both conceptually and for measurement. While organization at the firm level is economically important, focusing too closely on success at the level of the firm ignores much of what is important for growth in a *market* economy. Measures of firm success that are based in comparison to other firms must take market competition into account.

References

- Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity, Microeconomics*: 187–249.
- Bresnahan, Timothy, Erik Brynjolfsson, and Lorin Hitt. 2002. Information tech-

- nology, workplace organization, and the demand for skilled labor: Firm level evidence. *Quarterly Journal of Economics* 117 (1): 339–76.
- Bresnahan, Timothy, and Franco Malerba. 1999. Industrial dynamics and the evolution of firms' and nations' competitive capabilities in the world computer industry. In *Sources of industrial leadership: Studies of seven industries*, ed. David C. Mowery and Richard R. Nelson, 79–132. Cambridge: Cambridge University Press.
- Cockburn, Iain, and Rebecca Henderson. 1998. Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery. *Journal of Industrial Economics* 46 (2): 157–82.
- Griliches, Z., and J. Mairesse. 1998. Production functions: The search for identification. In *Econometrics and economic theory in the twentieth century: The Ragnar Fisch centennial symposium*, ed. Steinar Strøm, 169–203. Cambridge: Cambridge University Press.
- Milgrom, Paul, and John Roberts. 1992. *Economics, organization, and management*. Englewood Cliffs, NJ: Prentice Hall.
- Olley, Steve, and Ariel Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64 (6): 1263–98.
- Womack, James P., and Daniel T. Jones, and Daniel Roos. 1991. *The machine that changed the world: The story of lean production*. New York: Harper Collins.