

Federal Reserve Bank of Richmond

Market Structure and the Macroeconomy

Richmond, Virginia • May 16–17, 2019



Richmond Fed Research Conferences in 2019

April 19

**University of Virginia - Richmond Fed
Research Workshop (Spring)**

May 16–17

Market Structure and the Macroeconomy

May 22–23

**Technology-Enabled Disruption:
Implications for Business, Labor Markets,
and Monetary Policy**

June 7

**Technology Diffusion and Productivity
Workshop**

September 27

Regional Economics Workshop

October 2

Investing in Rural America

October 11

**Richmond Fed - University of Virginia
Research Workshop (Fall)**

For more information, please visit:

www.richmondfed.org/conferences_and_events



**FEDERAL RESERVE BANK
OF RICHMOND®**

Richmond • Baltimore • Charlotte

Market Structure and the Macroeconomy

May 16, 2019

- 9:00 AM** **Quantifying Market Power**
Jan Eeckhout, Universitat Pompeu Fabra Barcelona
- 10:00 AM** **Indivisibilities in Distribution**
Thomas Holmes, University of Minnesota
- 11:00 AM** **Diverging Trends in National and Local Concentration**
Esteban Rossi-Hansberg, Princeton University
- 12:00 PM** **Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States**
Gauti Eggertsson, Brown University
- 12:45 PM** **Lunch Speaker**
Tom Barkin, Federal Reserve Bank of Richmond
- 2:15 PM** **The Rise of Niche Consumption**
Joseph Vavra, University of Chicago
- 3:15 PM** **Markups Across Space and Time**
Arlene Wong, Princeton University
- 4:15 PM** **Labor Market Concentration, Earnings Inequality, and Earnings Mobility**
Kevin Rinz, Census Bureau

May 17, 2019

- 8:45 AM** **From Population Growth to Firm Demographics:
Implications for Concentration, Entrepreneurship and the Labor Share**
Hugo Hopenhayn, University of California, Los Angeles
- 9:45 AM** **Labor Market Power**
Simon Mongey, University of Chicago
- 10:45 AM** **Using Empirical Marginal Cost to Measure Market Power in the U.S. Economy**
Robert Hall, Stanford University
- 11:45 PM** **Concentration in U.S. Local Labor Markets:
Evidence from Vacancy and Employment Data**
Claudia Macaluso, University of Illinois at Urbana-Champaign

Conference Summary

In recent years, U.S. policymakers have voiced concerns that some industries have become too concentrated in the hands of a few large firms. Policymakers and economists alike worry about market concentration because it can be a sign of weakening competition. Over a two-day conference titled “Market Structure and the Macroeconomy” held at the Federal Reserve Bank of Richmond, eleven researchers presented evidence on various signs of market power and discussed what an increase in firms’ market power might mean for the economy.

Gauti Eggertsson of Brown University provided an overview of many of the reasons economists have to be concerned about an increase in market power. Presenting joint work with Jacob Robbins and Ella Getz Wold of Brown University, Eggertsson began by describing a number of recent changes in the economy that have puzzled macroeconomists. These include: the rise of wealth relative to output; the lack of corporate investment despite favorable conditions such as low borrowing costs; the decline in the real interest rate despite steady returns to capital; and the decline in labor and capital shares. Using a model, Eggertsson showed that these trends can all be largely explained by two changes — a decline in the natural rate of interest and a rise in market power. A decline in the natural rate of interest is something economists at the Fed and elsewhere have hypothesized for some time. Eggertsson and his coauthors arrived at the second explanation after examining the growing body of evidence that U.S. firms are increasing in size and power.

Market power is hard to measure directly, so many of the presentations at the conference focused on various indirect ways of identifying changes in market power. One of the most visible signs is an increase in industry concentration. Firms with a higher market share may face less competition, leaving them free to exercise more market power. This has made market concentration a focal point of the debate about changes in market power. Several researchers at the conference pointed to data showing that concentration has gone up across industries in the United States over the past three to four decades. However, not everyone agreed that this trend is necessarily a sign that firms have more market power.



Thomas Holmes of the University of Minnesota argued that the rise in concentration may be tied to returns to scale in shipping. In work with Ethan Singer of Compass Lexecon, he looked at how the use of shipping containers complements firm scale, promoting the growth of large firms. As supply chains have stretched across the globe, shipping costs have become a key factor in the competitiveness of firms. U.S. firms pay the same price for each shipping container they receive from China or elsewhere regardless of how full the container is. Holmes noted that large retailers that sell a variety of products, such as Walmart or Amazon, are able to ship full containers more frequently by consolidating shipments across their various product lines. Smaller firms often ship containers that are only half full, incurring greater costs. In principle, smaller firms could share containers to reduce their costs, but Holmes said that in practice few do so because of the transaction frictions involved. Thus, shipping technology seems to favor the development of large, diversified firms over smaller, specialized sellers.

Joseph Vavra and Brent Neiman of the University of Chicago examined retail concentration from the perspective of consumers. They began by identifying two seemingly contradictory trends: households have concentrated their spending on a smaller number of products since the early 2000s, but the variety of products sold at stores has grown over the same period. Vavra explained that households are increasingly engaging in “niche” consumption. Each household has a small number of top products in each category that the household prefers to buy, and those preferences vary by household. Using a model, Vavra



showed that the rise of niche consumption has had a neutral effect on changes in market power, as the different consumer preferences balance out in the aggregate.

Hugo Hopenhayn of the University of California, Los Angeles presented a paper with Julian Neira and Rish Singhania of the University of Exeter that attempts to explain the changes in industry concentration as a consequence of demographic changes. Since the 1980s, firms have become larger and older on average. The share of firms in the economy that are more than ten years old grew from about 30 percent in 1990 to about 50 percent in 2014. Over the same period, labor force participation has declined. An increase in the labor force increases the entry of new firms to employ workers, while a decline in labor force participation reduces the entry of new firms and promotes the growth of larger, older firms. Notably, Hopenhayn showed how this relationship between population and firm demographics can explain the rise in industry concentration while remaining agnostic about changes in market power.

Another paper presented at the conference by Esteban Rossi-Hansberg of Princeton University also focused on understanding the rise in concentration across U.S. industries. Alongside Richmond Fed economists Pierre-Daniel Sarte and Nicholas Trachter, Rossi-Hansberg observed that while national markets are becoming more concentrated, most markets are actually local. For example, restaurants and coffee shops compete for customers locally rather than na-

tionally. Using data from the National Establishment Time Series, Rossi-Hansberg and his coauthors found that in many industries with increasing national concentration, concentration at the local zip code level was actually falling. Moreover, these diverging trends appear to be driven by large national firms expanding into new markets and increasing the number of competing firms at the local level. To the extent that local markets matter most for competition, Rossi-Hansberg argued that rising national concentration may not be a sign of increasing market power.

The divergence of national and local concentration was also a theme in the presentations of three researchers who examined labor market concentration. Many commentators have argued that rising monopsony power is to blame for slow wage growth in recent years. Kevin Rinz of the Census Bureau found that average employer concentration has increased sharply since 1990 and that rising concentration reduces earnings and increases income inequality. This would seem to lend support to the idea that firm power in labor markets is increasing. But Rinz also found that employer concentration at the local level has declined since 1976. As a result, he argued that labor market concentration has contributed very little to rising inequality. In fact, he found that inequality would have been higher in the absence of falling local employer concentration.

Simon Mongey of the University of Chicago presented similar findings from his work with David Berger of Northwestern University and Kyle Herkenhoff of the University of Minnesota. Using a model, Mongey showed that an increase in labor market concentration would reduce the aggregate labor share across industries, explaining one of the puzzles that Eggertsson identified in his presentation. But similar to Rinz, Mongey found that local labor market concentration has declined since 1976. This decline should have increased the labor share, suggesting that labor market concentration cannot explain the decline in labor share observed in the data.

Claudia Macaluso and Chen Yeh, at the University of Illinois at the time of the conference and now economists at the Richmond Fed, focus on quantifying labor market power in the U.S. using a two-pronged approach. Along with Brad Hershbein of the Upjohn Institute, they first propose to estimate a direct measure of employer market power: the

“markdown.” Markdowns express the percentage of revenue going to the worker in the form of wages for every additional dollar produced. They find that workers at the average manufacturing plant earn fifty-six cents on each additional dollar of revenue and conclude that the U.S. labor market is far from the perfect competition ideal. Under perfect competition, wages and marginal revenue would have a one-to-one relationship. Furthermore, an employment-weighted average of markdowns shows a mild increase since the mid-1970s. When looking at concentration, on the other hand, Macaluso (who presented their findings at the conference) noted that fewer than 5 percent of the jobs created in the last decade have been in highly concentrated labor markets and that local labor market concentration has diminished over time. Thus, the extent to which concentration may be interpreted as employer market power appears limited. Furthermore, while labor market concentration has a negative correlation with local wages, it has a positive one with jobs’ skill requirements, implying that policy solutions like increasing the minimum wage may have both positive and negative effects on workers.

Economists have also looked to markups, or the prices firms charge for goods and services relative to their costs, as another signal of market power. Markups are difficult to measure, particularly across the entire economy. Doing so requires knowledge of firms’ marginal costs, which are often not publicly reported. Robert Hall of Stanford University presented a method of estimating marginal cost as the ratio of the change in cost to the change in output, for which data are available. Hall also discussed the value of expressing market power in terms of the Lerner index, which is the ratio of price minus marginal cost relative to price. The Lerner index ranges in value from zero to one, whereas markups have no upper bound, making econometric analysis more difficult. Using this method, Hall found that market power has been rising across a group of sixty U.S. industries since the 1980s.

Jan Eeckhout of Universitat Pompeu Fabra Barcelona also focused on markups as a sign of market power. He presented research with Mongey and Jan De Loecker of Katholieke Universiteit Leuven that showed average markups have been rising under a variety of assumptions and measurement tech-

niques. Eeckhout noted that this increase has been driven by a few large companies. For most firms, markups have remained flat. But sales have become concentrated at high-markup firms, and their larger market share means that they drive the change in average markups observed across the economy. Eeckhout identified two causes for this change in market structure. First, markets have become more concentrated through mergers and acquisitions and less rigorous antitrust enforcement. Second, changes in technology have led to higher fixed costs that create barriers for new entrants. Using a model, Eeckhout showed that the rise in market power indicated by the increase in markups may lead to higher profits for high-markup firms, wage stagnation, and a decline in business dynamism.

Arlene Wong of Princeton University also presented data on markups. She and her coauthors, Eric Anderson and Sergio Rebelo of Northwestern University, used price scanner data to calculate markups in the retail industry. They looked at data from 2006 to 2009, encompassing the period before and during the Great Recession. Wong noted that markups are relatively stable over time and only slightly procyclical. However, she and her coauthors found a large dispersion in markups across space. Retail markups vary by region, depending on local income levels. Wong said that this difference is driven by variation in the types of products purchased in each region rather than differences in prices for the same products. In regions with higher incomes, households tend to buy more products with higher markups. This suggests that an increase in markups could be partly driven by consumer choice rather than firms exploiting a lack of competition.

Taken as a whole, the presentations at the conference reflect the state of the market power debate in the economics profession. Some researchers argue that the evidence points to increased market power in the hands of large firms, while others offer alternative explanations. Undoubtedly, this topic will remain an important area of research in the years to come as policymakers and economists continue to discuss changes in market structure. ■

For more information on the research discussed at the conference, please see the interviews that follow. Interviews have been edited for length and clarity.

Quantifying Market Power

by Jan De Loecker, Jan Eeckhout, and Simon Mongey



High markups (the prices firms charge for goods and services relative to their costs) can be a sign of market power. Jan De Loecker of Katholieke Universiteit Leuven, Jan Eeckhout of Universitat Pompeu Fabra Barcelona, and Simon Mongey of the University of Chicago

find that average markups have been rising over time. This increase has been driven by high-markup firms increasing their market share across the economy. The increasing market power of such firms may be contributing to a variety of economic phenomena, such as wage stagnation and a decline in business dynamism. Eeckhout discussed these findings at the conference.

What have you found in your research on markups, and what do those findings say about changes in market power in the U.S. economy?

Eeckhout: What we find is that over the last four decades, since 1980, the distribution of markups and of profits has changed. The median, that is, the fiftieth percentile of the distribution, has remained roughly unchanged both for markups, which is the price above the marginal cost, and for profits, which is basically sales minus all the different costs. Those have remained unchanged for the median in the distribution, but they have changed dramatically for the upper percentiles. So a firm that had high markups in the 1980s now has even higher markups. And that's driving changes in average markups. But it's important to keep in mind that what is changing is the distribution.

Now, if you ask most business people, they will

say that business is tougher today than it was a few decades ago. And I think that's consistent with what we find in the data. There are a few business leaders who are extremely successful — Jeff Bezos, Mark Zuckerberg, or Bill Gates, for example — but most smaller businesses face stiff competition from those giants. And those high-profit firms are not just in the IT sector where you see this kind of disparity, it's everywhere. Part of the explanation for the change in markups is the change in the retail sector. There are a lot of big nationwide or even global chains. Whereas in the past you would buy your clothes from the local tailor, now you buy them from a chain.

Now, is that good or bad for the economy? It depends. The bad thing is there is less competition, which is not good because firms with fewer competitors charge prices that are too high relative to their costs. But it's not necessarily bad. Think about why Amazon is so successful and has such a large market share. It's because they can do things very cheaply. Same with Walmart. All of these chains that we see around the country, the reason they are so good for customers is because they provide goods at a lower cost and that's beneficial.

So there is a trade-off. I call this the "Amazon paradox," which says that, on the one hand, Amazon is so productive, which is great. It gives us low prices. But because they are so productive relative to other firms, they drive out the other firms and that kills the competition. Considering the costs that they face, they really should be selling at even lower prices than they are. But they don't do that because none of their competitors can get anywhere close to where they are. And so one of the challenges is how to deal with that.

How clear is it that markups are increasing for big firms? How hard is it to accurately measure markups?

Eckhout: It's really hard to measure markups. There are different methods, and we measure them from accounting data. But it's clear that no matter how you measure them, whether it's price over marginal costs or profits over sales, these things have changed. You can measure markups and profit rates in many different ways, and I think you want to look at all of them. We've tried to do it from as many angles as possible. By having different angles, we can come to a conclusion that's fairly robust. And what we find is, whichever way we measure these things, in every case we see a change in the upper tail of the distribution. We see increasing markups for the really high-markup firms.

Now, it is a mistake to conflate markups and profits. Profits take into account all costs, while markups only take into account the cost of producing. The cost of producing a bottle of water, for example, might be some plastic and the cost of water and filtering. But there are some other costs having to do with research and development that are not directly related to the production. That is another change we see in the economy compared with forty years ago. The costs of development, marketing, and compensation for executives have become larger shares of the total cost. These costs also often include intangibles. That's what's driving the wedge between what we see as markups and as profits, and that's why we cannot interpret them in the same way. But the important thing is that both markups and profits have gone up.

How does the efficiency of big firms affect how we interpret their higher markups and profits?

Eckhout: This gets back to the Amazon paradox, which suggests that the way these firms have achieved their market power is by cutting costs. It's not by raising their prices; it's by making their products cheaper. A markup just tells you the ratio between price and cost. It could be that the price has gone up for the markup to go up, but it also could be that the cost has gone down. In that second case, firms have become more efficient, which is good for the customer. But I think that's good for



customers only as long as there is some competition to put pressure on the firm to charge a price that's low enough.

One issue is that these firms might be charging low prices now, but maybe when their competition is decimated, then they will raise prices. That's a possibility. Even today you see from the flows of income that Amazon generates that they're pricing too high. And that is surprising because the products they sell to consumers are ridiculously cheap compared with any competitor. And yet, they could sell them for even less.

Could it be that in order to become so large and efficient, these firms have had to make investments that we haven't fully accounted for and that helps explain the markups?

Eckhout: It is definitely the case that the firms that invest more in intangibles are also the ones that can produce at a lower cost. Because Amazon has invested in its delivery network, it can cheaply and efficiently deliver products to my doorstep. That does mean that they need to price a little bit higher to cover that fixed cost of investment. But what we see in the data is that these firms still earn profits that are over and



above what they would need to cover such costs. This indicates that while their size helps them to become so efficient and cheap, it also allows them to kill their competition in the process. We see that from two angles: their large market share and the profits that firms with large market shares earn.

The interesting question is, how do you deal with this? Just chopping up these firms is not necessarily going to be the right solution. If you chop them up, you lose the reason why they are so cheap. The reason why they are so cheap is because they are so big. But because they are so big, they also kill the competition. So we have to find a way to boost competition and at the same time allow those firms to keep their efficiency from being big. It's really like a catch-22.

We have done some work looking at this, and we find that we could do better by introducing a little bit more competition. The question is, how do we do it? There are a number of things we can think about, such as forcing Amazon to have competitors on its network. In Europe, if you cross out of your cell phone network, telecom providers who have towers in that area must let other operators use their tower at a regulated price. It's a form of competition, and it lowers prices. We see that in most European countries the prices charged by telecom providers are much lower and profits are much lower than they are in the United States. So maybe there is a similar solution for introducing competition in other industries.

To me, there are simple solutions to other problems. For example, Facebook's acquisition of Instagram, in hindsight, probably shouldn't have been allowed. These were two direct competitors for advertising

money in social networking. They shouldn't be owned by the same entity. Likewise, InBev probably shouldn't have been allowed to buy Anheuser-Busch. They now have a virtual monopoly in the U.S. beer market. This is not good for the customer. This is not a case like Amazon that makes distribution cheap. This is you going to a bar where you have a choice of many beers, but they are made by the same owner. You don't have competition between owners.

Some researchers have argued that while national concentration is going up, local concentration is falling. So in the case of InBev, for example, they may own a large share of the national market, but in any given local market, they must compete with dozens of microbreweries. Does that change things?

Eckhout: Part of the reason why you see local concentration going down is because national chains are generating market power by competing in local markets. They compete, which is good for the customers, but profits at the local level have gone up. Concentration has gone down because something has changed in terms of the scale. These chains are all much bigger than the local firms. Like Amazon, the chains are so efficient that they lower costs at the national level. I want to stress that there's a difference between setting high prices and lowering costs. These guys do it because they are really cutting costs. That's their strategy, and they're very good at it. I do agree that it's an interesting challenge to address because you don't want to get rid of the lower concentration at the local level. You would like the profit margins to be lower at the local and at the national level, but at the same time, you don't want to kill the large chains' efficient provision of goods.

Has there been any research on whether large efficient firms do actually raise their costs once they have driven out competition? That argument has been made about Walmart in the past, for example.

Eckhout: Thomas Holmes, who presented at this conference, was one of the first people to do research on Walmart's strategy to conquer America.

I think Walmart became successful eventually not by jacking up prices. They made huge profits mainly because they were so efficient nationwide. To me, it's a precursor of what Amazon did in the last decade. Walmart did this in the 1980s and 1990s. I don't think it's the case that in many of these local markets, they destroyed all the competitors. You know, people say Walmart destroyed all the small stores, and it's true that many small stores have disappeared. But at the same time, you had all these large chain retailers that entered these local markets. There is a danger of a Walmart or of an Amazon jacking up their prices in the future, but what we've seen is that Amazon became a threat to Walmart and put even more competitive pressure on Walmart. So there was no possibility for Walmart to become the sole dominant player in the market.

Then what is the likelihood of another firm coming along to challenge, say, Amazon if they get too big?

Eeckhout: If it happens, it will occur with a change in technology. At the beginning of the twentieth century, there were Sears and the other mail-order companies. That was a big change because we had railways, and companies could deliver products anywhere. Later on, you had Walmart and Amazon competing to deliver goods to customers in different ways. It wasn't a copy of Walmart that was competing with Walmart, it was Amazon, which is a different model. So maybe in twenty years, some other type of model will allow a new firm to eat away the power that Amazon has now. But we won't know until it happens.

Are there any other questions related to markups or market power that you're looking at right now?


Eeckhout: In addition to finding the correct measurement of markups, one of the things that I have been interested in from the start is what kind of macroeconomic effects they have. My coauthors and I have looked at wages, labor force participation, labor market dynamism, and how fast firms change their labor force. One thing we haven't looked at yet, but are starting to look at now, is the effect of

higher markups on wage inequality for workers with different skills. We've seen that the wages for skilled workers have gone up. The conventional wisdom is that this is due to technological change, specifically skill-biased technological change. That is, new technologies are much better for skilled workers than for unskilled workers. If you drive a car for a living, that technology hasn't changed much over the last several decades. But if you are a product designer, the advances in computers and software have made a huge difference. That makes you so much more productive. So what people have been arguing is that the changes in wage inequality we are seeing are completely driven by pure technological change. But is there something about market power that exacerbates that?

In trying to measure that, we do find that market power can in fact make wage inequality stronger. Think of Apple, which has its headquarters in Cupertino. They attract these brilliant coders all there together. But Apple outsources its customer services for troubleshooting, its call centers, and its production facilities. And those workers get paid rock bottom wages, whereas the guys at headquarters get these phenomenal wages. Apple can only do that because it has so much market power. It's that interaction between the two that we're looking at — the effect of market power on wage inequality. ■

Indivisibilities in Distribution

by Thomas J. Holmes and Ethan Singer



One factor that can affect industry concentration is the ability of large firms to exploit efficiency gains from economies of scale. Thomas Holmes of the University of Minnesota and Ethan Singer of Compass Lexecon examine how large, diversified

firms can more efficiently make use of shipping containers. Firms pay essentially the same price for each shipping container they receive regardless of how full it is. Large retailers that sell a variety of products, such as Walmart or Amazon, can consolidate shipments across their product lines to ensure they always ship full containers. Smaller firms often ship containers that are only half full, incurring greater costs. At the conference, Holmes discussed how the “indivisibilities” of shipping containers favor large firms over smaller ones.

Could you describe what you mean when you use the term “indivisibilities” in the context of this paper?

Holmes: An “indivisibility” is an economics term used to describe a situation where the cost of doing a job at half scale is essentially the same as doing it at full scale. In the distribution sector (wholesale and retail), let’s say goods need to move from point A to point B by a truck pulling a trailer. Sending that trailer on the road half full costs essentially the same as sending it full. In today’s economy, there is competitive pressure for firms to offer a large variety of goods and to deliver them fast. This means that for any particular good, the shipping firm might not have enough shipment volume to fill up a box (a trailer or

ocean container). The big firms like Walmart, Target, and Amazon are combining their large-scale and advanced information technologies to overcome these indivisibility issues.

How much lower are indivisibility costs for firms like Walmart and Target relative to smaller firms? And is the increase in indivisibility costs relatively precipitous once you get beyond these kinds of firms?

Holmes: Indivisibility costs gradually increase as retail firms get smaller. But once retail firms get small enough, they tend to offload distribution to other firms, such as wholesalers. To get a sense of the magnitude of the costs involved, in our paper we look at the contents of ocean containers. The first point is that big firms ship boxes that are filled to the brim virtually all the time, while small firms often pay a shipping premium, using more expensive half-size containers that they, not infrequently, ship half full. So right there we see a big difference in transportation cost. But that is only part of the story. It’s not just about how full the boxes are, it is also about speed of deliveries and use of information technology. Small firms could fill their boxes if they chose to have very infrequent deliveries, for example. The key advantage of the big firms is they don’t have to make this trade-off. They enjoy the benefit of frequent deliveries with maximal use of information technology and ship full boxes all the time.

How do large online retailers like Amazon fit into this picture given that the variety of goods they sell is much greater and their distribution strategies are much different?

Holmes: In its early days, Amazon tended to focus on high-value goods that could be shipped



relatively cheaply. But as they have expanded, they have moved into things such as gas grills and microwaves that are bulky and have high transportation costs. In many ways, they have followed the Walmart script in the way they have vertically integrated into distribution. But they are also innovating in big ways. Online retail presents both challenges and opportunities compared with traditional brick-and-mortar stores. On the one hand, online has to deal with the disadvantage of getting goods to the consumer's doorstep, the "last mile" problem. Everyone talks about this logistics challenge. But there is less discussion about a key advantage of online retail: while Walmart basically has to have all of its products on the shelves of all of its 3,500-plus stores in order to sell a good, Amazon gets to put high-turnover goods in fulfillment centers close to consumers and low-turnover goods in maybe one place in the whole country. This provides a great deal of flexibility.

You consider what would happen to ocean freight costs if Walmart were to be split in half. What would that mean for Walmart's profit rate and for consumer welfare?

Holmes: In our analysis, we project that if Walmart were split in half, its ocean import costs would

increase by about 4 percent. This includes costs of not packing boxes quite as full, but also the costs of not being able to order as frequently on account of the lower volumes. We focused on just the import aspect of the distribution, but we expect that a cutback in Walmart's size would have analogous impacts throughout Walmart's supply chain. Walmart has a thin profit margin, and its spending on transportation is roughly of similar magnitude as profit. A 4 percent hit to profits would be a big deal. In our paper, we don't work out how consumers would be affected. However, Walmart's business model throughout its history has been to turn its cost savings into price savings for consumers so as to generate high volumes. So I would expect that some of the higher costs would be passed to consumers.

Related, it seems like splitting up Walmart would have nontrivial effects on ground transport costs as well. Do you have an idea of what the division of Walmart into two firms would mean for total transport costs?

Holmes: We suspect that the same kind of indivisibilities that we are highlighting in the initial movement of goods from factories in China to import distribution centers at ports also happen, in mirror image, in getting goods from regional distribution centers to

stores. So I suspect the cost effects over the system as a whole would be of similar magnitude.

What role did transportation play in the rapid success of Walmart?

Holmes: I think it's a huge chunk of it. Just to be able to generate the volume so that you can have this very dense network of distribution centers; that was a novel idea. Interestingly, if you look at something similar like Target, they didn't get it. They scattered. Their first store was in my hometown of Minneapolis, but then they put another one in Denver, and they just sort of scattered these around. They weren't thinking about distribution. They were focused on marketing and so forth, but they didn't get these economies in density that Walmart did.

Walmart came at it completely differently. Instead of scattering stores around the country, they started from Bentonville and radiated out, never jumping ahead to get to the next market. They get there because they have this scheme where you add more stores, more stores, more stores, and then you can sort of skip forward and build the next distribution center. But all along the way, you never have a distribution center that doesn't have stores thick around it. You're building the stores out, and for a little while your stores have to go a little farther than you'd like to. But then you build the next distribution center and switch those stores over to that one.

What are the frictions that are causing shipments coming from India to be less efficient than those coming from China?

Holmes: The issue here is scale. Walmart's volume of imports from India is only a small fraction of its volume from China. This makes it difficult for Walmart to execute the playbook it is using in China. Everything is scaled up, and it's easier to fill the boxes and to make frequent deliveries. One thing I really want to stress is that this isn't just about filling boxes. Anybody can fill a box. If you're a small company, you could just get one delivery a year or something and your box is full. But then you've only got one deliv-

ery a year. The trick is to get frequent deliveries that can utilize information technologies on consumer demand and to keep the boxes full. So when you compare Shenzhen versus Mumbai, or China versus the whole Indian subcontinent, all the stuff is spread out. And Walmart is wedded to this idea that they're going to break it up five or six ways. That's sort of putting pressure on how you fill the boxes. If you buy something from some small factory in the middle of India somewhere, now you have to cut it up into five boxes. Sometimes to do that, they send it to Mumbai, and sometimes it's just easier to put everything in its own box and send it on six different ships. Whereas in China, things are more concentrated. Trucks come down to these consolidation stations, and they're just easy to stash full.

Some critics of the current transport system might look at it and say, this looks awfully antiquated. For instance, we are using ships, containers, and cranes that are not optimal. Do you care to comment?

Holmes: I think the system works pretty well. It is true that many of the standards date from the 1960s, but I think the reasons they worked well then are generally why the standards work well today. In particular, a key benefit of a container is how it can be lifted off a ship and be attached right to a trailer for a truck to haul away. A full-size container is a pretty good size for taking on our highways. As with anything, there might be some gains to making changes at the margins if the standards had to be set all over again. However, there is great benefit from the high degree of standardization all over the world. This makes it easy for firms to share equipment and keeps transportation costs down. ■

Diverging Trends in National and Local Concentration

by Esteban Rossi-Hansberg, Pierre-Daniel Sarte, and Nicholas Trachter



Several studies have highlighted the fact that market concentration is rising across industries at the national level. But many markets for goods and services are local, not national. Customers planning to eat out at a restaurant typically limit

their choices to local establishments, for example. Esteban Rossi-Hansberg of Princeton University and Pierre-Daniel Sarte and Nicholas Trachter of the Richmond Fed find that while national markets are becoming more concentrated, local market concentration is actually falling. These diverging trends appear to be driven by large national firms expanding into new markets, probably increasing competition at the local level. Rossi-Hansberg presented their findings at the conference.

Market concentration has been the topic of recent public policy debates. What role have economists played in that debate?

Rossi-Hansberg: I think economists have played a crucial role in the debate. There is obviously a constant concern in all economies about market concentration because, to the extent that it leads to an increase in market power and pricing power for a company, it can be detrimental for consumers. Along those lines, there are antitrust enforcement institutions that are there to try to prevent abuse from happening and intervene whenever it's needed. The United States has a long tradition of intervening to break up big firms, and this has been an important mechanism to prevent concentration.

Lately, what economists have found is that concentration in the economy has been growing. This has

triggered an exploration of what is underlying those changes and whether they can tell us something about the evolution of the economy. This is a bit different from the actual enforcement of antitrust law, which happens at a much more detailed industry level compared to the more aggregate macro trends that economists have calculated in a series of papers. The results have raised an interesting and controversial question: Is the fact that we're seeing big firms becoming more important in many sectors a problem or not?

That's definitely an important question to ask. There are a lot of related facts that people have built into a common narrative. Those facts include the decline in the labor share, the decline in firm dynamism or firm creation, the mobility of workers across firms, increasing profits, and increasing markups. So there's a range of facts out there, and it seems they are all very consistent with a narrative that the economy is being dominated by large market players and that these big players may be exerting market power.

What does your paper add to that narrative?

Rossi-Hansberg: A cornerstone fact of that narrative is the increase in firm concentration. That's a very easy thing to calculate. You don't need a lot of assumptions to measure an increase in concentration. Hence, it is not controversial whether an increase in concentration has happened or not. We know that national concentration has increased, more in some sectors than in others, but overall it definitely has increased. The other facts that I mentioned, such as the increase in markups and profits, require a lot more assumptions. And so depending on what assumptions you make, you can find big increases or not. But I think what lent credibility to the whole narrative was that everything was consistent. Maybe each fact is not 100 percent proven, but if I show you that there are ten facts or five facts that are all consistent with each other, and I can

build a narrative around it, that seems very credible.

What our paper does is look at one aspect of the narrative: the fact that concentration has increased. One problem is that, yes, concentration has increased at the national level, but most markets are not national. Most markets are local. For some things, it's obvious. When you're choosing a restaurant, or a coffee shop, or a hotel, you're going to use establishments that are local. You're not considering all the hotels in the country or all the coffee shops in the country, you are only considering the coffee shops, hotels, or restaurants in a small geographic area. Given that, we decided to ask whether we find the same increase in concentration when we look at a more narrow geographic area. What we find is something completely different. We find a secular decline in local concentration since the 1990s that is very large; it is even larger than the increase in national concentration. The decline is also very pervasive. The share of employment or the share of sales in industries where we see this type of decline in local concentration amounts to about three-quarters of the economy. What that tells us is, to the extent that competition happens at the local level in these sectors, the whole narrative of increasing concentration and the accompanying increase in market power crumbles completely.

Where did you get the idea to look at local concentration? Did you start with a theory about the proper definition for markets, or did you look at the data and see this diverging trend and decide to investigate?

Rossi-Hansberg: We had a basic notion that a lot of markets are local, and because a lot of markets are local, we should be looking at local measures of concentration. When we saw that the results for local concentration were very different from the trends in national concentration, we started digging further. One question we asked ourselves was, how could concentration be increasing at the national level but declining at the local level? The key factor, we discovered, is the role of multiestablishment firms. When you look at the firms that have been growing in industries where we see increases in national concentration, pretty much all of that expansion is happening at the

extensive margin. That means firms are replicating geographically rather than increasing the size of the establishments that they already have. So it's not just that they're increasing employment, they are actually increasing the number of establishments they have.

When a firm like Walmart expands, for example, it's adding new establishments in different places, which decreases local concentration. For example, think about an industry in which every location has a monopolist. For example, the local hardware store is a monopolist in every town. Now, Home Depot enters and takes over essentially the whole market in those locations. But the local hardware store that used to be the monopolist is still there. It just has a smaller share of total sales, say 5 percent. If that happens at the national level, we see a big increase in concentration because, in this example, Home Depot took 95 percent of the market. But at the local level, where there used to be one store, now there are two. So concentration at the local level has declined.

How have your findings been received, given that they contradict the common narrative of rising concentration?

Rossi-Hansberg: Well, some people have received it better than others, I think. But in general, the notion that we need to look at these local measures and that trends in local concentration may look different from trends at the national level is pretty embedded in the profession now. Everyone working in this literature acknowledges that these trends are different, and I believe this is the important result of our paper. But there are two essential frontlines for further discussion. One question is, what is the right geographic scope of markets to consider? Our paper doesn't really answer that question. We just calculate changes in concentration over time for different levels of geographic aggregation. For all local levels, whether for counties, cities, or zip codes, we find big declines in concentration. But that doesn't answer the question about the right market definition. And, of course, the right market definition is going to depend on the industry. For some industries, the market is national. As in manufacturing, where goods are easily shipped from one place

to the other. For other industries, like restaurants, it's ridiculous to think about a national market. So there's clearly variation across industries in terms of how markets should be defined. Our paper is fairly descriptive in that sense. It just says that there are different national and local trends in concentration, and the creation of new establishments by top firms decreases local concentration while increasing national concentration. But one could ask, how do we know whether we should care about national concentration versus local concentration for a particular industry? That's a fair question, and I don't exactly know the answer. We have some idea of how tradable these industries are. And we find in the data that the diverging trends in local and national concentration are less pronounced in sectors like manufacturing. But how to determine exactly the right market is not something that we address in the paper. Moving forward, it's an important question to address.

The other issue people have raised about the paper has to do with the data we use. We use the National Enterprise Time Series (NETS), which is a private dataset. That's not the standard data people use for these types of studies. The standard data come from the Census. The reason we use these data is because they don't have the disclosure restrictions that Census data have. We were very interested in understanding the role that top firms play in these changes in concentration, and if you cannot disclose the role of top firms or of one particular firm, then you cannot do these calculations. NETS allows us to do that, but it's not the standard Census dataset that people have used in the past to think about firms and firm dynamics. Because of that, there has been some hesitation about whether or not the NETS data are reliable enough to make our findings credible. We think it is, and we've done a lot of checks to demonstrate that. By now other researchers have found similar trends using Census data.

Are there any other questions that you hope to address in future work?

Rossi-Hansberg: This paper is just saying that there is increasing national and decreasing local concentration and that top firms are responsible for it. But it

doesn't say what's underlying these secular trends in concentration. That is, why are these firms expanding by adding establishments rather than growing their existing establishments? I have some work with Chang-Tai Hsieh ("The Industrial Revolution in Services") that argues that this may be the result of what we call a new industrial revolution in services. So, Ford's assembly line was a big new way of producing manufacturing goods. It worked by scaling up the plant and streamlining production so that you could produce many cars in the same location. It was a fundamental revolution for the manufacturing sector, and pretty much every modern manufacturing process works that way today. But it did nothing for services, since producing many services in the same location is useless. Most services have to be produced close to consumers. What we've seen since the 1970s is something like an industrial revolution but in the service sector. Essentially, information and communication technologies have facilitated the geographic replication of service firms. Firms have invested in building a brand and model establishment, as well as in building a distribution, communication, and information network to operate many establishments. These investments amount to switching to high fixed-cost, low marginal-cost technologies that encourage firms to replicate rather than build up existing establishments. The result is that service providers were able to scale up by expanding geographically and are now present in many more locations. This phenomenon can, therefore, explain the trends in local concentration, which are particularly pronounced in the retail and service sectors.

Economists have traditionally struggled to find empirical evidence of the gains from information and communication technology. We believe that some of the gains from computers and software have been in allowing many service industries to invest in fixed-cost-saving distribution and communication networks, for example. These fixed-cost technologies then allow them to replicate geographically. The result is that smaller towns and more cities are getting access to the services of top firms. Part of the decline in concentration that we have discussed is due to the new entry of top firms into new markets. ■

Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States

by Gauti Eggertsson, Jacob Robbins, and Ella Getz Wold



A number of recent changes in the macroeconomy cannot be explained by traditional models. These puzzles include an increase in wealth relative to output, low corporate investment despite favorable interest rates, a decline in the real

interest rate despite steady returns to capital, and a decline in the labor and capital shares of production. Gauti Eggertsson, Jacob Robbins, and Ella Getz Wold of Brown University find that these puzzles can be largely explained by two changes in the economy: a decline in the natural rate of interest and a rise in market power. Eggertsson discussed their findings at the conference.

What puzzles in macroeconomics have you been studying?

Eggertsson: The first puzzle relates to the work of Thomas Piketty, who has talked about how capital is back. He has pointed to a great increase in the amount of capital income relative to output. We have a little twist on that suggesting that it's actually not really capital per se, that is, an increase in capital as a factor of production (things like factories), but it's really wealth more broadly defined that has risen relative to GDP. The traditional measure of capital over output, which is basically just a weighted average of past investments, has actually been relatively constant. But investment has been relatively constant, so there seems to be a rise in what people measure as wealth relative to output over the last thirty years for reasons that cannot be captured by the traditional

investment series. Traditional models have a difficult time accounting for this because in the traditional neoclassical model, wealth is synonymous with productive capital. So something else is going on. Now, I think most people don't necessarily associate wealth with productive capital. But in the models economists write down, they are very closely tied together. So this suggests that we need to move beyond that.

Another puzzle relates to Tobin's Q, which states that if you take the market wealth of corporations, that is, how they're valued on the stock market, and then you divide that by the replacement cost of capital, it should be equal to approximately one in the long run. That's because if firms are valued a lot higher than what it costs to put them together, then people will start investing in more firms and that will bid down the price until it's about the same as the replacement cost of capital. But what we've seen in the past thirty years or so is that there has been an increase in the valuation of corporations without the corresponding increase in investment. So that's a second puzzle: Given that there has been a huge increase in the valuation of corporations, why are we not seeing more people investing in productive capacity?

Another feature of the basic neoclassical growth model, at least in its most stripped-down form, is that the average return on capital should be the same as the marginal product of capital. You can get one measure of marginal returns, the real interest rates that firms face when they make investment decisions, from financial markets. Meanwhile, you can measure average return directly from the national accounts. You just take the value-added of corporations, subtracting things like cost of capital, labor costs, and taxes, and then divide by the capital stock. This gives you a measure of average return. The puzzle is that the marginal and average returns on capital

should be the same in the neoclassical model, but the average return has remained relatively constant while the marginal return has been declining persistently over the last thirty years or so.

A fourth puzzle relates to the share of production allocated to the factors of production — capital and labor. In the standard neoclassical model, the shares paid to capital and labor are roughly constant over time. Instead, what we're seeing is that the shares paid to labor and capital appear to be declining over the past thirty or forty years. What seems to have been occurring is what some people have called growth in "factorless" income, which is income that is not really tied directly to either capital or labor. In our paper, we call that pure profits.

The fifth puzzle we look at is: despite the fact that Tobin's Q is high and despite the fact that real interest rates are at record lows, we see very little investment.

Our paper argues that you can explain all five puzzles with two things: an increase in monopoly power and a decline in the real interest rate. There is little disagreement that there has indeed been a sustained fall in real interest rates. This can be directly observed in financial markets. It is much more controversial whether monopoly power has increased. As has been discussed at this conference, various measures from the micro data suggest that there has been an increase in the monopoly power of firms, which of course has various implications. Our paper is not so much providing new micro evidence for the rise in monopoly power. We are just pointing out that if you want to make sense of these broad patterns in the macro data, you need something that at least shows up in a similar way as a rise in monopoly power. To me, that's complementary evidence to those researchers who are presenting microeconomic data supportive of a rise in market power. There has been some discussion about whether or not the micro evidence is compelling. Our paper would then pose the question, how could you then account for these puzzles in the macro data if there were no change in monopoly power?

What led you to consider a rise in market power as a possible explanation for these puzzles?

Eggertsson: It makes sense to move a little bit away from a model of perfectly competitive markets. The model we had been working with when we were addressing different questions already had those features. But then we realized that not only do you need imperfect competition, you need the degree of imperfection, or monopoly power, to be increasing over time. In some sense, it's easy to make the case that we don't live in perfectly competitive markets. You only need to look out of your window. But the second proposition, which is that the degree of competition appears to be declining over time, is a bit more surprising.

We came to the idea of declining competition because we were exploring secular stagnation, which is the idea that the declining real interest rate we have seen over the past thirty years is a permanent feature of the economy. And as we were writing down models to make sense of that, we found it difficult to do so without some complementary force that prevented things like an increase in investment. I think for those who do not believe that there has been an increase in monopoly power over the past thirty years, they are left trying to find an alternative explanation because it is clear that interest rates have declined. It is clear that investment has not increased. So if rising monopoly power is not the explanation, then what is? It's not enough just to say that productivity has declined, at least within the context of modern models. We need another force to check it. And to me, this is by far the most plausible, especially given the additional empirical micro evidence people have been putting forth.

That said, we sort of came to this independent of any evidence. We felt that it was the only way we could rationalize these things. We were writing this paper on secular stagnation, and this was more of a footnote in that paper just to make the model work with the data. Then all of these papers came out with

evidence of increasing market power, and so it made sense to try to illustrate why we felt that we had needed it in our previous work.

A key aspect of your model is that you need both higher market power and low interest rates to solve the puzzles. Do you see a causal link between those two developments?

Eggertsson: We certainly thought about writing down a model that would try to hit all five birds with one stone — to have a common feature that explained every puzzle. The prime suspect we have for the falling interest rate is demographics. And it could be that the aging of the population is also feeding into a lack of business dynamism. New entry tends to be driven by young entrepreneurs. So to the extent that the population is getting older, you might imagine that there's less firm entry, which feeds into higher concentration and market power. There's some evidence of that, so I think in principle one could write down models that rationalize those two phenomena with a common factor. We felt we didn't need to quite take a stand on it now, but I think certainly that's something we might do at some later point.

Are there other explanations for the puzzles you look at?

Eggertsson: There has been discussion of the fact that the capital stock could be mismeasured due to intangibles. The main problem I see with that explanation is that the mismeasurement would have to be so big that I think it's somewhat implausible. Also, researchers do try to take into account intangibles in the data we already use. There is always the possibility that somebody figures out a better way to measure things and finds that indeed that is a reasonable explanation, but I haven't seen a compelling case made for it yet.

The other thing people have noted is that the real interest rate we are looking at is the risk-free rate. Maybe the risky rate is a lot higher than that and there has just been an increase in risk premia. Again, I think the main problem with that explanation is that

it's hard to make those numbers match the data. It would require too big of an increase in risk premia. I suppose if you put mismeasurement of risk premia and capital stock together with a variety of other things, you might be able to cobble together a different explanation. It's certainly something we'll think harder about and see if we can put together other compelling cases, but I think it is going to be difficult. Of course, there is also the possibility that this is a temporary phenomenon that reverses itself. So I'm not 100 percent invested in this position; it just seems like the most plausible explanation at this time.

Does your work suggest any policy responses?

Eggertsson: I think it's a bit premature at this point. In our model, we don't really have a theory for the degree to which rising market power is efficient or inefficient. So I can't really say one way or the other because we haven't explicitly modeled it. But my sense is that rising monopoly power is a cause for concern, and it is hard not to tie this development at least in part to changes in antitrust enforcement over the last three decades or so. These developments seem to coincide quite a bit with the deregulations in the 1980s and after. So my hunch is that there might be implications for policy, but I certainly wouldn't be willing to start giving policy prescriptions just yet.

I work on this from the theory side, and I can think of models that one could write down where these changes could be fully efficient versus models where they correspond to something more problematic. So while my suspicion is that this is probably not a good development, that's just a suspicion at this point. At this conference, for example, there were people making interesting cases for one position or the other. Some talked about how it's less obvious that there is more concentration in terms of the choices available to consumers. I think we'll know more in a couple years ■

The Rise of Niche Consumption

by Brent Neiman and Joseph Vavra



Whether they shop online or at brick-and-mortar stores, consumers today have a growing number of products to choose from. Despite this increase in variety, Brent Neiman and Joseph Vavra of the University of Chicago find that households have concentrated their

spending on a declining number of products since the early 2000s. At the conference, Vavra explained that households seem to be engaging in “niche” consumption. Households buy specific products within each category of consumption, and preferences vary by household. Vavra argues that this increase in niche consumption is likely driven by an expansion of product availability and substantial associated welfare gains as households are better able to pick products tailored to their specific tastes. He notes that niche consumption has had a neutral effect on changes in market power as different consumer preferences balance out in the aggregate.

What was your motivation for this paper?

Vavra: As this conference highlights, there has been an explosion of work over the past five or so years looking at market structure changes on the production side of the economy. But I would say that we know comparatively very little about household product spending patterns and how changes in market structure on the production side have potentially translated to household spending. Based on my own anecdotal experience, it feels like there are more and more products and choices available across stores. In the paper, we focus primarily on grocery store spend-

ing, where we can really accurately measure detailed product spending in several different categories. And we find two divergent trends, which together we call the rise of niche consumption. Basically, across time there are more products available, which is consistent with my anecdotal experience, yet households are increasingly sorting into distinct niches from one another and purchasing different products.

So, even in this traditional sector that I think of as still being primarily dominated by brick-and-mortar retailers, we’re seeing this growing importance of long-tailed products and product niches, which has received a fair amount of attention in digital markets. We’re seeing similar phenomena showing up in traditional retail, and most of the paper is dedicated to drilling into those trends empirically and thinking about their theoretical implications for welfare and market power.

How much of this growth in niche consumption is due to online retailers like Amazon that sell a huge variety of products versus brick-and-mortar stores?

Vavra: The categories we’re studying are primarily things that are sold in brick-and-mortar stores, so in a mechanical accounting sense I think the rising importance of online platforms isn’t particularly important for our results. It’s certainly possible that some combination of greater online competition or the internet fostering greater knowledge about product availability has influenced brick-and-mortar retailers’ decisions to stock a greater variety of products. But the direct effect of online shopping is not a first-order thing for the trends we’re finding.

Could you give a few examples of this niche consumption? How big is it empirically?

Vavra: If you look at a product category like tortilla chips, you see that something like the original Tostitos chip lost a fair amount of market share over our sample period. But if you look at some other products that are gaining market share, Tostitos Scoops and Tostitos Cantina were two products that had fairly large increases over this time period. So, if you look at chips in the aggregate, it seems like we're adding new products that weren't there before. If you look at overall aggregate spending, it used to be concentrated on a small number of products and now it's spread over these different products. But at the same time, if you look at individual household spending, before that product innovation, households were just splitting their spending roughly equally between, say, Tostitos and similar Mission brand chips. And now there are these more specialized products available, and households seem to be sorting into just one of these products and they tend to not switch back and forth very much. So even though aggregate concentration has gone down, household concentration has gone up.

In terms of the magnitudes of these things, overall we're finding that household concentration has increased by around 12 percent or so over the years in our sample and aggregate concentration has declined by 25 or 30 percent. Whether those are big or small numbers is a little hard to assess without a model, which is one reason we build a model to try to think about these things. And through the lens of the model that we build, I would say these trends are fairly sizable and are associated with pretty substantial welfare gains.

Do those numbers surprise you? Or do they seem about what you would expect from observation?

Vavra: I would say that the aggregate concentration trends didn't surprise me all that much. I sort of suspected that the number of varieties available in the aggregate had been increasing, and then I would expect that as there are more products available, the concentration would go down. I don't think I had any strong feeling one way or another about the particular size of the trend. I think the increase in house-



hold concentration and the fact that households are purchasing fewer varieties is maybe a little bit more surprising to me ex ante. But I don't think I had any strong priors one way or another going into the project about what these trends would look like.

How do consumption patterns look across and within demographic groups?

Vavra: There are some obvious things, like the fact that larger families have less concentrated spending. Slightly less obvious, but maybe not surprising, is the fact that older and poorer families have more concentrated spending. But we see this trend toward niche consumption happening in every demographic group that we observe in our data. So this is mostly a within-group phenomenon. It's not primarily a rich versus poor thing. And I think that's not that surprising. These are very finely defined products, so goods purchased by different groups differ markedly from each other on average. But there's always a lot more variance across households within groups than between, say, a representative rich household and a representative poor household. Once you aggregate up to broad demographic groups, you're already averaging out a ton of heterogeneity in what different people purchase. So we think of these trends as being primarily a within-group phenomenon. That pretty strongly suggests, at least to us, that whatever the root cause of this trend is, it's probably something that's pretty broad-based and is affecting many different people in the economy. It's probably not

a cause that is unique to a particular experience or demographic group.

Do you find any major differences between rural and urban areas?

Vavra: Not in the trends. I would strongly suspect that if you compare cities to more sparsely populated areas, there are going to be more things that you can purchase in Manhattan than in rural Kansas. But in terms of the trends, they look broadly similar in most locations. It doesn't seem like there's any strong correlation with density or anything like that.

Do you see the phenomena you're describing as consistent with some of the fragmentation in other parts of the economy and in society more generally?

Vavra: We view our results as very much complementary to that type of evidence about growing fragmentation in several aspects of society. Our results are showing that households are increasingly segmenting in the products that they purchase, so product consumption is becoming more fragmented. Of course, this doesn't mean that the underlying causes of the trends we're documenting need to be the same as the underlying causes of fragmentation in other dimensions. Indeed, I think the fact that we're finding these trends within narrow demographic groups suggests that this product market fragmentation is not primarily being driven by growing income inequality or political polarization, so it's probably a distinct phenomenon from those other things. But it's yet another dimension of increasing fragmentation in society.

What do your findings imply about market power?

Vavra: I should start by saying that this is the part of the paper that's still a work in progress. We think it's an important question, but we're still actively working on it. For now, our model delivers some preliminary results that we think are interesting and suggestive in this dimension. The first is that, relative to most macro models, our model has fairly compli-

cated yet still tractable implications for the elasticity of product demand and firm market power. In particular, when firms in our model gain some additional spending along the intensive margins from households that were already purchasing their products, it is kind of like the standard elasticity of demand that you would get in any kind of CES (constant elasticity of substitution) macro model. But the thing that's a little more unique to our setup is that firms can also pick up additional demand along what we think of as an extensive margin from households that weren't previously purchasing the good but are induced to purchase the good by lowering prices. So that means market power in our model is endogenous and differs across products.

Thinking about the implications of that, our preliminary results suggest that this rise of niche consumption is associated with basically two offsetting forces for market power. The first is that if you hold the number of products and the level of competition constant, as households become more fragmented or sort more into these narrow niches, it is a force toward greater market power for firms in the economy because the market is more segmented and that lets firms increase markups. But at the same time, as the number of products goes up, holding all else equal, that's a force toward greater competition. So that force pushes toward lower market power. And when you put these two forces together, for now our results suggest that they exactly cancel each other out because of the particular specifications we have. So it suggests that aggregate market power may be relatively unaffected by this increase in niche consumption because of these opposing forces.

More broadly, an implication of the model is that any mapping from measures of concentration to implications for aggregate market power is complicated. Even if we observe an increase in concentration, that doesn't necessarily imply any increase in firms' market power. It depends on the structure of the model.

Do you find anything related to the "paradox of choice," that is, the idea that an increase in product variety might be making people's lives more difficult rather than improving welfare?

Vavra: I think we're highlighting a slightly different force. We're not directly modeling the paradox of choice. Instead, we're highlighting a positive effect on welfare from greater variety of choice, which a lot of models already include. But we have an additional mechanism that most macro models don't, which is that when more products are available, households are able to better select a basket of goods that's more suited to their idiosyncratic preferences. This ability to select a set of products that's ideally suited to their tastes is going to have a positive effect on welfare. And it's a fairly sizable effect through the lens of our model.

So we're not directly modeling the potential negative effects of being overwhelmed by variety. But our model can loosely proxy for that channel if it manifests itself in some greater psychological cost of purchasing varieties, which maybe gets bigger if there are more varieties out there. Maybe it's harder to figure out what you like or find the products that you want to buy in the store. Our model can sort of account for that in a reduced-form way as an increase in the cost of purchasing varieties. And indeed, our model finds a little bit of that kind of force in the background, but it's pretty modest, and it doesn't have much of an effect on welfare. It's overwhelmed by the positive effect of greater variety availability letting households select the things that they particularly like.

Are there any extensions to this research that you are hoping to explore in the future?

Vavra: Right now our model is largely focused on the household demand side. One thing that we'd like to think more about is modeling the firm supply side and putting these two things together. At the moment, our model implies that an increase in the number of varieties is something you really need to fit the data. But we'd like to think about firms potentially wanting to target particular types of customers or explore the technological change on the firm side that lets them actually increase the varieties they offer. That might be something like supply chain management or better information technology that

helps firms determine what households like.

In addition to the trends that we're looking at within the United States, I'm also interested in cross-country differences in these kinds of things because I think you learn a lot about different countries having more or less availability in stores. I know one of the things that people who immigrate to the United States are often struck by is just how many products are available in the store. Quantifying that is something we might be able to do a little bit more carefully with similar scanner data in other countries. Even just from anecdotal data, the number of varieties available in the United States is higher than in stores in other countries. But this does not necessarily mean that individual households are purchasing more varieties than in other countries. Our model suggests that looking at individual household purchases and how correlated or uncorrelated they are across households can be informative for assessing the welfare benefits of differences in varieties across countries. ■

Labor Market Concentration, Earnings Inequality, and Earnings Mobility

by Kevin Rinz



Many have argued that rising monopsony power is to blame for slow wage growth in recent years. Kevin Rinz of the Census Bureau found that average employer concentration has increased sharply since 1990. Moreover, rising concentration is

associated with reduced earnings and increased income inequality. But Rinz also found that labor market concentration has actually declined at the local level since 1976. In his presentation at the conference, he noted that employer concentration therefore may not explain changes in inequality.

What was the motivation for this paper?

Rinz: There is a lot of evidence of rising markups and increasing concentration nationally that suggests that markets of various kinds have become less competitive. In the labor market specifically, there has been some talk from policymakers that declining competition could be behind broad macro trends that we have seen, such as stagnating wages and increasing inequality. But within the labor market, local conditions seem like they're probably what's important, since most job searches are relatively local. When I started looking into this, there wasn't really much evidence on how concentrated local labor markets were, and there was even less evidence about how concentrated they had been in the past. So my goal for this project was to start by providing some basic facts about local labor market concentration over the last four decades and then to estimate how changes in that concentration might affect earnings and inequality.

What do the trends look like nationally versus locally, and are they different from one industry to another?

Rinz: I find that, as opposed to national labor market concentration, local labor markets have generally been becoming less concentrated over time. Most major sectors of the economy have seen fairly similar trends, either small declines or relatively stable concentration over the period I considered. The most notable exception to that is retail trade, which has seen increasing local concentration. But even there, local concentration in retail trade has been increasing less quickly than national concentration in retail trade.

What's driving the differences in national concentration compared to local concentration?

Rinz: I think the thing that seems to be driving the divergence between local and national concentration is the behavior of nationally large firms. It looks like between 1976, when my data started, and 2015, which is the most recent year available in my paper, nationally large firms have expanded their geographic reach pretty substantially. They have been entering new markets that they hadn't been operating in at the beginning of the period, which leads them to compete directly with each other more frequently. So when you measure concentration nationally within industries, this leads to increasing employment concentration, since these big firms are operating more broadly. But when you measure concentration locally, the fact that they're more often in the same place as each other leads to lower concentration because they have more similarly prominent competitors operating in the same market.

I actually had a chance to present this paper last year with Alan Krueger in the audience, and he gave me a good example of an industry that illustrates this dynamic pretty well: concert promotion. It used to be the case that every town a band played in had its own relatively small businesses that would help them put on their shows. But now, larger companies like Live Nation or AEG Live are operating in all sorts of different markets. So they are more dominant nationally and also more frequently operating in the same places. That provides a real-world example of how it might be that national concentration could be increasing while local concentration declines.

How do you define a local market, and how much variance do you see in concentration across markets?

Rinz: There are a lot of plausible ways to go about defining a market. In my paper, I use the intersection between four-digit NAICS industries and commuting zones as my definition of a local market. I also tried a few alternatives, and the main takeaways from the paper are not very sensitive to which particular definition you use.

It turns out there is actually a good amount of variation in concentration across markets. You can see there are some that have measures of concentration on the Herfindahl-Hirschman Index that are relatively high, implying equivalence to a market with only two equal-size firms. And there are other markets where it's very low, suggesting that any given firm has very little market power. The more concentrated markets tend to be in places that are less populous.

In markets where concentration is increasing or was already large, what do you see happening with earnings?

Rinz: In places where concentration goes up, earnings tend to go down, and inequality tends to increase. There are different ways to think about whether the effects I find are large or not. One kind of thought experiment is to compare the experience of the median worker to the experience of a worker at the seventy-fifth percentile of the concentration distribution.



What does that magnitude change in concentration imply for earnings? Making that move would reduce earnings by about 15 percent. Of course, that move amounts to a tripling of concentration, which is a big change and one that you don't generally see within a given market. So that's not a common experience by any means.

A more macro-oriented way to think about the magnitude would be to take the actual experience of changes in local concentration in the aggregate over the period covered by the data and think about what earnings and inequality would look like if they were instead held constant. My estimates imply that earnings would be about 1 percent higher and inequality would be about 6 percent lower in 2015 if local concentration had remained at its 1976 levels. I view that as kind of meaningful but not especially large in macro terms.

Are these effects different across demographic groups?

Rinz: It does look like there are some potentially different effects across demographic groups. The most striking example to me is that men appear to be more adversely affected by increases in concentration than women. I think the broader set of demographic estimates could be useful for kind of guiding future work that's more focused on the mechanisms of what exactly changes within a labor market as it becomes more concentrated. Different groups experience the effects differently.

What do you find when you look at changes in concentration over time by region?

Rinz: It turns out that rural areas are not only relatively highly concentrated in the cross section, they also tend to be the ones that experience larger increases over time. The Great Plains region in particular is home to some of the commuting zones that have the highest average concentration across their industries and also to the zones that saw the largest increases in concentration between 1976 and 2015. It's also the case that if you average across all of these localities, the places that have experienced larger increases in concentration since 1976 are much less densely populated on average than the places that have seen declines. So it's a real phenomenon.

Does this suggest that workers in rural areas are subject to significant and persistent earnings disparities relative to workers in urban areas?

Rinz: I didn't explicitly test urban versus rural in the paper, but some of my estimates do suggest that the consequences of increasing concentration might be more severe in rural areas. When I produce estimates that give each market equal weight, so the smallest markets in the least populous areas are treated equally to, say, the financial market in New York City, I find more negative effects from concentration than estimates that give markets the weight that's proportional to their employment. So that tells me that the smaller markets that are being "upweighted" in this second approach are potentially even more severely impacted by increased concentration. Of course, these are just average estimates, so it is not necessarily the case that every rural worker is adversely affected. But on average I think this suggests that they may be hurt more than workers in urban areas.

Are you considering any extensions to your paper at this stage?

Rinz: One that I think would be interesting is to look more explicitly and thoroughly into the conditions facing rural markets. Because the levels of concen-

tration are higher, the increases in concentration are more substantial, and the responsiveness to increases appears to be larger. I think digging more into that would be interesting and important.

Another issue I am planning to explore more is related to how labor markets are defined locally. A lot of the work on market concentration relies on definitions of convenience, such as industries, occupations, or geographies that are available off the shelf and can be assembled relatively easily. But I think everybody acknowledges that labor markets are a lot more complicated than that. One question I think would be interesting to answer is, do we get a similar picture of how concentration has evolved over time if we define labor markets using observed worker behavior, like job transitions, instead of using these easily accessible market definitions? ■

From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship, and the Labor Share

by Hugo Hopenhayn, Julian Neira, and Rish Singhania



At the national level, concentration has increased across industries in the United States over the last several decades. Hugo Hopenhayn of the University of California, Los Angeles and Julian Neira and Rish Singhania of the University

of Exeter argue that this change may be the result of changes in demographics. They find that firms have become larger and older on average: about 30 percent of firms in 1990 were more than ten years old, compared with about 50 percent in 2014. At the same time, labor force participation has declined as the population has gotten older. Hopenhayn, who presented at the conference, noted that a decline in labor force participation reduces the entry of new firms, resulting in the growth of larger, older firms. He showed how the relationship between population and firm demographics can explain the rise in industry concentration.

How does the aging of companies and of the population as a whole explain the growing size and concentration of firms?

Hopenhayn: Firms' size is increasing with age. That is a fact that holds not only for U.S. firms but for firms globally. The average size of a firm that is between twenty and twenty-five years of age is over four times as large as the average size of a new entrant. The same holds for concentration; average concentration increases with the age of a cohort. The aging

of firms, measured by a shift to the right in the age distribution of firms, will thus translate into higher average firm size and concentration as more weight is put on older cohorts.

In our analysis, the aging of the population does not, by itself, have a direct effect on concentration or average firm size, but it's the driver of the aging of firms. The rate at which firms are created is directly related to the growth of the labor force, as labor is one of the most important inputs of firms. This holds true in the aggregate data for the United States and also when comparing population and firm dynamics across counties in the United States, as shown in the work of Benjamin Pugsley and Ayşegül Şahin. So the aging of the population slows down the growth rate of the labor force and thus new firm creation. Population aging leads to a gradual process of firm aging, which leads to an increase in average firm size and average concentration.

What other factors could be driving the aging and concentration trends?

Hopenhayn: It has been argued that higher concentration and market power might have resulted in increased barriers to entry. While this might be true, if indeed incumbent firms were more entrenched and protected from the competition of entrants, this should translate into an increase in their market shares and into lower exit rates even when conditioned on firm age. This is something we don't see in the data, so that is a challenge for that kind of theory. Admittedly, it might still be true in certain markets.

It has also been argued that initial public offerings have decreased considerably, which means there has

been a decline in the number of public firms, and that this might be due to listed firms acquiring new ones before they become public. This could also lead to increased concentration.

I should also point out that because the dataset that we use (which includes all firms as listed by the Business Dynamics Statistics) started in the late 1970s, the age of all firms that entered the market prior to that is unknown. This is referred to as left-censoring of ages. As a consequence, while we consider that the left-censored group and its average size and concentration have increased, we cannot tell how much of this is the result of aging in this group as opposed to increases conditional on age. So it might be that some of the changes we have seen are not only the result of a change in the age distribution.

What led you and your coauthors to consider changes in firm demographics as an explanation for changes in concentration?

Hopenhayn: The fall in entry rates plus the change in age structure that we observed implies a slowdown in employment growth. This prompted us

to look for a cause and immediately suggested a slowdown in population growth. So we came to this fact indirectly.

The U.S. isn't the only developed economy experiencing slowing birth rates. Does the trend toward older, larger firms also appear in other countries with aging populations?

Hopenhayn: We have not done a serious analysis, partly because we don't have the necessary data. What we have seen for many European countries is that during this period there wasn't a slowdown in fertility and there has not been a marked increase in average firm size. This is quite preliminary research, though, and it's too early to reach any conclusions. Further, these countries had a decrease in fertility that occurred earlier than in the United States, so a more exhaustive study should consider these episodes and their impact on firm demographics. The same is true for Japan, for which unfortunately we could not find the data.

What, if anything, do your findings imply about changes in market power in the U.S. economy?





Hopenhayn: We have no direct measure of market power as measured by markups or the share of profits, so all our inference is indirect. There are studies that show that larger firms have lower labor share and higher markups. The aging of firms means that there is more weight on larger firms, so this should also imply a rise in average markups and a decrease in aggregate labor share, as others have observed. Indeed, the recent literature on growing markups also concludes that most of the change is explained by composition effects, in other words, reallocation. The aging of firms due to decreased entry provides a channel for this reallocation.

Based on your model, do you have any predictions about the future trend in the startup rate and average firm size?

Hopenhayn: Our model predicts a moderate rebound in entry rates and consequentially a slight reversal of the firm-aging trend. This should result in slight decreases in average firm size and concentration, which in turn might imply a slight reduction in markups and a rise in labor share. But, if anything, these effects are very modest since the change in fertility is persistent. The rebound in our model occurs because, aside from long-run effects from the decrease in fertility that we have observed, there has been additional aging in the transition path to this new long run. This component of aging is what we predict will be reversed.

Are there other questions related to firm demographics and concentration that you hope to explore in the future?

Hopenhayn: We plan to extend our analysis to other countries as data become available. Our analysis pertains to the U.S. economy as a whole, but there have been considerable sectoral changes during this period. We plan to extend our analysis to a multisectoral framework to better understand firm dynamics at the sectoral level.

Migration is another important source of labor force growth, so it should also be related to firm demographics. We were planning to study this. Another consideration is the aging of the labor force, which might in itself affect entry rates, as entrepreneurs tend to be concentrated in young to middle-age groups.

Regarding sectoral changes in the United States, what sectors do you think could be driving the overall change in firm demographics?

Hopenhayn: There has been a fall in entry in most of them. Manufacturing stands out since its labor share has been declining for a long time and changes in the global economy might have had a bigger impact on manufacturing. We also did our analysis excluding manufacturing, and the results were similar. ■

Labor Market Power

by David Berger, Kyle Herkenhoff, and Simon Mongey



Product markets have become more concentrated nationally, but what does that mean for labor markets? As David Berger of Northwestern University, Kyle Herkenhoff of the University of Minnesota, and Simon Mongey of the University of Chicago note

in a recent paper, most labor markets are local. At the conference, Mongey presented their model, which divides the United States into many different labor markets. Through the model, they explore the implications of changes in labor market concentration.

What are the costs of higher labor market power that you find in your paper?

Mongey: In our model, labor markets are populated by a small number of firms. That number lines up with the U.S. data. In the United States, there is an average of about forty firms in any given labor market. If it's harder for workers to move from one labor market to another than it is to move between firms within the labor market, the firms that are the most productive will have labor market power over workers in that market. If you think of an isolated market where the largest firm is trying to employ a bunch of workers, those workers are going to want to go to that firm because they are potentially offering higher wages than other firms. Additionally, the workers don't really have any other options because they're isolated from other markets. In that setting, the firms that are the biggest in the market would ideally be the ones that are the most productive, paying the highest wages, and employing the most people.

In our model, the firms that are paying the highest

wages and employing the most people are still employing fewer people and paying lower wages than they would if they were in an environment where they faced many other competitors. So the costs associated with labor market power in the economy are that wages are lower and employment is lower than would otherwise be the case. One of the aspirations of our paper is to think of the U.S. economy as containing many of these labor markets. There might be a labor market for making tires in Ohio, for example, and a labor market for cutting ice in New York. We stitch all of those labor markets together so we can add up the cost in terms of lower wages and lower employment across the entire economy. When we add that up, we end up with a lower share of productivity being paid out to workers in the economy. That is the main cost associated with labor market power in our model.

How do your findings relate to recent claims that monopsony power has increased in the economy?

Mongey: Monopsony is a word that has been bandied about pretty freely recently. In macroeconomics, the idea that firms have some kind of market power in the product market in different environments is very natural. In labor economics, models with a perfectly competitive labor market, where there is essentially one wage and a firm can choose how many workers to employ at that wage, has been the benchmark. But there is a class of departures from that perfectly competitive model, which you could call monopsony. Our model fits in there.

In this environment, if firms want to employ more workers, they're going to have to pay them more as opposed to facing a constant wage. They face an upward-sloping labor supply curve. But we don't think monopsony by itself is really enough to think

about welfare in the labor market, especially in terms of changes in labor market concentration over time. Most of the frameworks you would write down with a very simple monopsony model actually distract from what most people think of when they think about labor market power. That is why it was important for us to write down a model where you have heterogeneity in firms — some are big, some are small — and you actually have a meaningful notion of concentration within the labor market. In that “oligopsony” environment, where there is a small number of buyers for labor, you can really think about how concentration, welfare, and the labor share are all linked through the labor market power of firms.

Some people have made the argument that growing labor market power might be behind some trends that we’ve seen in the economy recently, such as slow wage growth and the falling labor share. What did you find regarding this?

Mongey: One of the contributions of our paper is to provide a tight link between the labor share, which is the share of total revenue paid out to workers in the form of wages, and a measure of concentration. And the particular measure of concentration in our paper comes straight out of the model. So as far as this model goes, we can think about it as the welfare-relevant measure of concentration in the economy. And this welfare-relevant measure of concentration in the economy is a particular weighted average of a particular measure of concentration across all of the labor markets in the United States. One of the things that we wanted to achieve with this paper was to try to understand what measures of concentration are actually important for determining things like the labor share and welfare. And this is the one that pops out of the model.

We can use U.S. Census data to construct this measure from the ground up, and then we can look at how the measure has changed over time. What we find is that this particular measure of concentration has actually declined over time. Given that decline in this measure of concentration, we can compute the counterfactual effects on the labor share, holding

everything else in the economy constant. What we find is that a decline in concentration across labor markets in the United States would have contributed to a counterfactual increase in the labor share of around 6 percentage points. So, if you’re looking for something to explain the decline in the labor share, our paper strongly suggests that this measure of concentration is not it.

Some people have looked at different measures of concentration in the economy that have gone up. Those include a national measure of concentration in a particular industry. The reason our model rejects those measures of concentration as being appropriate for thinking about welfare in the labor market is simply because most labor markets are local. If I think about the labor market for mechanics in Chicago, mechanics can move out of the labor market and move to other places, but most of the job transitions in that labor market are across different employers within Chicago. So when you average local measures of concentration, it undoes trends that are masked when you look at national measures of concentration. And these local concentration trends actually go in the opposite direction. Similar findings have been made by a few other contemporaneous studies as well.

If it’s not labor market concentration that’s driving trends like the fall in labor share, do you have any theories about what it could be?

Mongey: It might not be concentration in the labor market, but it could well be concentration in the product market or in other input markets that is driving these trends. In fact, at this conference, my coauthor, Jan Eeckhout, presented a paper that he, Jan De Loecker, and I have been working on where we took a similar set of tools that I used in this paper with David Berger and Kyle Herkenhoff and thought about how we might understand other trends in the economy through changes in product market power. That paper is still in an earlier stage than this one, but it could be the case that labor markets are becoming less concentrated locally while product markets have become more concentrated nationally. So, take national airlines, for example. As airlines have merged,

an airline like United has more labor market power even though there are still multiple airlines operating out of any given airport. So it could be that.

Other theories regarding falling prices and investment goods have been tested in work by Loukas Karabarbounis and Gabriel Chodorow-Reich. So, there's lots of interesting work to be done on the change in the labor share. One reason we liked our paper was that it could help researchers and policymakers by removing one of the potential explanations for the fall in labor share that we thought people were perhaps jumping to a little too quickly.

Even if labor markets are local, if you have a large national chain operating in many local markets, would they have some labor market power?

Mongey: So, take a firm like Amazon. It's a big player in the product market. If you think of an individual who is working in the warehousing sector, it might be the case that locally they still have a bunch of firms they could sell their services to. Yet on the product market side, Amazon is able to take its inputs and deliver products at a very low marginal cost. Amazon sells things for cheap. It has market power in the product market, so it could potentially sell these things for even cheaper. But it doesn't do so because it has this market power.

These two things can, in a sense, be orthogonal. We can think of a consumer buying a good and a worker contributing services to the production of that good as being in completely different markets. So, I buy things off Amazon and they get delivered to me from Amazon's national supply chain. The warehouse workers work in localized labor markets for warehouse workers. One of those markets might be particularly concentrated and the other one might not be. So even if I decide to buy from a small set of national producers, there could still be a lot of competition locally for labor.

Now, these two markets could potentially move in the same direction. If the country were made up of only one labor market, then you could think of labor market concentration and product market concentration as the same. But as soon as I have multiple

labor market regions, the two things aren't as clearly linked. So you wouldn't want to infer measures of labor market competition by looking at changes in the share of sales at the top twenty-five firms in a particular industry.

How does minimum wage policy affect labor market power?

Mongey: In our revisions of the paper, we want to make a few theoretical points about how one can characterize the effects of minimum wage when labor markets are concentrated. By concentrated, I mean that there is a small handful of firms with different levels of productivity producing goods and employing workers in the same labor market. So in the model when you introduce a minimum wage, what tends to happen is that you increase the employment and wages of the lowest-productivity firms first. In our model, the firms with the lowest productivity are small, have low wages, and don't employ many workers. If you push the minimum wage up, they increase their employment and their wages. As they increase their wages, they attract workers away from firms that have higher wages that aren't directly affected by the minimum wage.

This is one of the ways that people think of the minimum wage as being welfare-improving. In our model, when you push up those small firms' wages, their competitors respond by also increasing their wages. So it acts as a rising tide that lifts all boats; wages are pushed up throughout the economy. As wages increase, households are willing to supply more labor and employment increases in the economy as well.

However, one of the effects that you find in our model is that if you increase these wages too much, then these small firms, which aren't particularly productive, don't really want to employ all of this labor. So at some point, these small firms start contracting. As you push up the minimum wage, you hit a low-productivity firm, which indirectly pushes up the wages of higher-productivity firms. So there is this wave going through the economy, and at some point the smallest firms start shrinking. Hypothetically, if

you had a minimum wage of \$100 an hour, then a mom and pop store is not going to employ any labor. But a firm like Amazon might still employ some labor.

So one of the findings we get from our model, which we think is a unique contribution to the literature, is that as the minimum wage increases and these smaller firms start shrinking, you are actually handing more market power to the biggest firms. That's because you're shrinking their competitors, which were keeping them on their toes. As you deliver more market power to the larger firms, that's a force toward lower wages and employment at the larger firms. So what we get quite naturally out of that in the model is an optimal minimum wage: a minimum wage that increases wages and employment at the firms with the lowest wages in the economy but doesn't do so to the extent that they shrink and shift more labor to the bigger firms. I think that is where we are going to leave it in the paper rather than making a quantitative statement about what that optimal minimum wage is. But I think it's a new consideration that policymakers might want to make when thinking about changing the minimum wage.

You've noted that the largest firms in the economy also tend to be the most efficient. Have you measured whether having higher market concentration either in the labor market or in the product market is a net positive or a net negative for the economy?

Mongey: That's another point we wanted to make in this paper. If we consider a counterfactual where we completely get rid of market power, we find that welfare increases because we are now in this competitive economy with no labor market power at all. But we also find that labor market concentration actually goes up.

What we're really interested in as macroeconomists is some measure of welfare and consumer well-being. Labor market concentration by itself doesn't really tell you that. You really need a model to interpret what that measure of labor market concentration means. Importantly, you need to take some kind of stand on what the labor market structure of the economy is. If you have a competitive labor

market with high concentration and an oligopsonistic labor market with low concentration, welfare is going to be higher in the competitive case.

In a competitive economy, the largest firms are the most productive, pay the highest wages, and employ the most workers, but they also have the most labor market power. When you eradicate their labor market power, they suddenly grow even bigger and pay even higher wages. As they grow even bigger and pay even higher wages, the smallest firms get even smaller and pay even lower wages. It reallocates labor to the most productive firms in the economy. But because the bigger firms are getting bigger, labor market concentration is going up. So, it's really hard to go from measures of labor market concentration to welfare.

For example, if there were a very pro-competitive change in policy in the labor market in the United States, it might be the case that labor market concentration goes up. And it could be the case that that's actually good for welfare. So, we really need to think through the market structure of the economy, which is what we try to do in the paper.

Are there any other issues related to labor market concentration that you hope to study in the future?

Mongey: We have a grant from the Washington Center for Equitable Growth that is helping us fund research into the effects of mergers on labor market power, wages, and employment in the United States since the 1980s. We're looking at the effects of national mergers on local labor markets. As I said before, we think of labor markets as inherently local things. If two firms merge nationally, how does that merger affect a labor market where the two firms are big players versus a labor market where the two firms are small players? And with that variation, can we learn something about the effect of firm size and labor market power on wages and employment? And potentially from that we could develop some kind of easy-to-compute statistic that policymakers could use to try to understand whether mergers are going to have a positive effect on the labor market or whether they might have a negative effect. ■

Using Empirical Marginal Cost To Measure Market Power in the U.S. Economy

by Robert E. Hall



Markups, the prices firms charge for goods and services relative to their costs, can provide a signal of market power. Firms with more market power are generally able to charge higher markups. But markups are inherently tricky to estimate,

often due to a lack of adequate data. At the conference, Robert Hall of Stanford University presented a new method for estimating markups and market power. Using this method, he found that market power has been rising across sixty U.S. industries since the 1980s.

What do you think might be driving the renewed interest in the topic of market power within the economics profession?

Hall: I think, first and foremost, it is this observation that has been pushed very hard and is highly relevant, which is the decline in the income share of labor. I'd say there's some consensus that rising market power is the one explanation that survives a serious confrontation with the data. So, for example, one possibility would be that industries with inherently low labor income shares have grown relative to those with high labor income shares. But that's been demonstrated to be incorrect. So essentially all the other explanations didn't survive, but market power does. That points researchers toward finding any confirming measures of market power, and a lot of that work has been focused on concentration rather than on actual measures of market power. That research has reached very mixed conclusions about concentration.

But that's no surprise to people who've read the industrial organization literature, which suggests there's only a very modest connection between concentration as measured conventionally and market power. So the time is ripe to look at market power again.

What sparked your interest in this topic originally?

Hall: I was briefly an expert in an antitrust case where we actually had data. It was in the concrete business, which made it easy to believe that there was market power because of transport cost. I was puzzling through how you might measure marginal cost, and then I gradually realized that the problem could be transformed into something that I had studied carefully in graduate school, which was Solow's productivity measurement. So it all fell into place kind of neatly.

The primary focus of your paper is to look at measuring marginal cost rather than trying to measure market power. Why did you choose that focus?

Hall: Normally, economists don't regard measuring price as a problem. So estimating marginal cost is the hard part of measuring the ratio of price to marginal cost. Nonetheless, I think I learned at this conference that there is more to say about measuring marginal cost than is actually in this version of the paper. I think the next version of the paper will be even more focused on presenting new ideas for measuring marginal cost as opposed to providing new estimates of market power. Because, as I explained in my presentation, statistically the results are not very strong. So it's an idea that's still looking for a body of data that would allow you to get stronger results. It's possible there is some kind of middle

ground between the two approaches that have been pushed recently, which would use a different econometric technique than the one I used and be less reliant on weak instrumental variables.

Are you thinking of the work of Jan De Loecker?

Hall: Yes, exactly. He gets relatively strong results by making some very strong assumptions. So, he uses strong assumptions and gets strong results, and I have weak assumptions and weak results. It's very much a question of what you think of his assumptions.

You described your paper as Robert Solow's 1957 paper updated to consider the possibility of market power. Could you explain that?

Hall: There's an equation in the paper that has the Lerner index as a coefficient. If you set that to zero, then you're assuming competition, and you can interpret this as applying Solow's approach to measuring productivity growth. After Solow came up with that idea, many people observed that there's a specification error, since many people (including Solow himself) working with similar data don't believe that every industry is competitive. So what happens? The answer, which I explored in my 1988 paper, turns out to be very simple. It just adds that one additional term, and the coefficient in that term, if you do it the way I do it in this paper, turns out to be the Lerner index. So it's a neat generalization. It's exactly the sort of thing that is nice econometrically because it's so straightforward.

Could you describe some of the advantages of the data that you use and some of the downsides?

Hall: An awful lot of work over the years has gone into refining these relatively detailed industry-level measures of quite a variety of inputs: capital, labor, energy, materials, and services. These make up the KLEMS index. People at the Bureau of Economic Analysis and the Bureau of Labor Statistics have been refining these data because it leads directly into productivity measurement, which is a very important thing for



their agencies to measure. So as a researcher, you can harness all of that work.

If you use data on individual firms, which is attractive for other reasons, there's no agency that's doing the same thing at the firm level. You have these very approximate measures that appear in databases like Compustat, as well as others that are international and have the same kind of coverage as Compustat. So, you could harness that. But the disadvantage is that it's not completely solid because of potential measurement errors, and it appears that there's a lot of variation among the various firms in an industry. That means that if you just use the industry levels, which is what I do, you can get spurious movements that represent changes in the composition. I believe that, within industry, heterogeneity is important. That pushes you toward individual firms, although many of the firms that are in the data are themselves very heterogeneous. So data on individual firms are not ideal. You'd like line-of-business accounting, and that is even harder to get and even sketchier than firm level. It's not perfect.

Another way to go is to look at industries like concrete where it seems like the product is pretty physically homogenous. The heterogeneity in that industry comes from how far it is from the ready-mix plant to the user. For example, concrete is actually cheap on the island of Manhattan because there's lots of concrete made across the river. Whereas if you're in some remote location that is thinly populated, concrete can be quite expensive because of the transport cost. So



there's lots more to do. We need to push on every margin just to try to improve. Better data would certainly be a big part of it. That's why I think the idea of looking at individual industries is still also very valuable.

What do you find when you measure market power across the industries for which you have data?

Hall: There has been some increase but of course statistical uncertainty is present. There is also heterogeneity across these industries. In the paper, I attempt to illustrate the heterogeneity across the sixty industries. That gets at this notion that there are some industries that are close to competitive and some industries that have Lerner indexes of 0.3 or 0.4, which is pretty high.

Have you thought about the potential policy implications that might come from these empirical findings?

Hall: Well, I started this line of research because I was an expert witness in an antitrust case. And I've served in that capacity in a variety of different places, especially high tech. I was very involved in Microsoft on both sides. But the basic tenet of antitrust policy is that just having market power alone is not a violation of any antitrust law. In fact, it may be a sign that things

are going right. If you have winner-take-all industries, then you have what they call in this line of work "competition for the market" as opposed to "competition in the market." Everyone thought Microsoft was going to ruin the world, and there was tremendous excitement about it having market power. That gradually subsided, and now it's Facebook and Google that are getting attention. Of course, they didn't displace Microsoft in the computer desktop field, the desktop just turned out not to be as important as people thought.

The area that's been most affected by the availability of market power measures is merger analysis. The tools for saying what will happen if a merger takes place are often controlled by measures of markups. Agencies will always use a variety of methods to decide whether to challenge a merger or not, but certainly a lot of those methods are rooted in measuring the existence of market power. So there are very direct policy implications for this research. ■

Concentration in U.S. Local Labor Markets: Evidence from Vacancy and Employment Data

by Brad Hershbein, Claudia Macaluso, and Chen Yeh



Claudia Macaluso and Chen Yeh, at the University of Illinois at the time of the conference and now economists at the Richmond Fed, focus on quantifying labor market power in the U.S. using a two-pronged approach. Along with Brad Hersh-

bein of the Upjohn Institute, they first propose to estimate a direct measure of employer market power: the “markdown.” Markdowns express the percentage of revenue going to the worker in the form of wages for every additional dollar produced.

They find that workers at the average manufacturing plant earn fifty-six cents on each additional dollar of revenue produced and conclude that the U.S. labor market is far from the perfect competition ideal. Under perfect competition, wages and marginal revenue would have a one-to-one relationship. Furthermore, an employment-weighted average of markdowns shows a mild increase since the mid-1970s. When looking at concentration, on the other hand, Macaluso (who presented their findings at the conference) noted that fewer than 5 percent of the jobs created in the last decade have been in highly concentrated labor markets and local labor market concentration has diminished over time. Thus, the extent to which concentration may be interpreted as employer market power appears limited. Furthermore, while labor market concentration has a negative correlation with local wages, it has a positive one with jobs’ skill requirement, implying that policy solutions like increasing the minimum wage may have both positive and negative effects on workers.

What was the motivation for this paper?

Macaluso: The main motivation was to understand two things. First, the connection between employer market power and employment concentration. Second, how workers are affected by employer market power and whether market power can explain trends we are worried about, such as stagnant wages or reduced opportunity for some categories of workers, like low-skill workers.

The first part is concerned with measurement. We wanted to explore if market power in the labor market is related to concentration, and if this is not the case, how are we supposed to measure it? After we have the measurement down, the second part is to investigate the relationship between both direct and indirect measures of market power and workers’ wages.

How do you measure employer market power, and what is the markdown?

Macaluso: At the cost of simplifying a little, there are basically two ways to measure employer market power: markdowns and concentration. We do both.

Let’s think about a firm that has market power in the labor market — a monopsony. If the firm is a monopsony, its choices on the quantity of labor it employs will affect the price of labor (the wage) in the whole market. That’s exactly the same as a monopoly in the product market when the choice of the quantity produced affects the product’s price. When you solve the problem of a monopsonistic firm, what you see is that, contrary to the competitive case, the wage is not equal to the marginal product of labor. There’s a wedge between the two: workers are not compensated one-to-one for every additional dollar

of revenue produced. This wedge is what we call a markdown. You can think of it as something that reduces the wage and lets the firm keep higher profits. The markdown is a natural direct measure of employer market power. In fact, I would say it is *the* measure of employer market power. So we estimate markdowns at the establishment level for the manufacturing sector in the U.S. to draw conclusions on how important monopsony is in the U.S. labor market.

Estimating markdown requires detailed data. Specifically, we need data on capital, labor, and intermediate materials. These data are only available in the United States (and other countries) for manufacturing plants. Nonetheless, what I think is really nice about our estimation procedure is that we maintain very general assumptions when we estimate the markdowns and we are able to uncover the full distribution of markdowns. Then we consider what happens if you don't have the data to calculate markdowns. The data just don't exist for services or retail, for example. Can we use concentration as a proxy?

That's where the comparison with the employment shares comes in. Concentration is based on a firm-level employment share. That's how the Herfindahl-Hirschman Index is constructed, for example. And we find a mildly positive relationship between a firm's employment share and the firm-level markdown rate. It's good to see that there's this positive relationship because it reassures us that even if you can't do a sophisticated estimation of markdowns, you still can draw some conclusions by studying the

firm's employment shares, which we can do for every single firm in the economy. On the other hand, while the ideal measure (markdowns) shows widespread and mildly increasing monopsony, concentration paints a different picture. Thus, the extent to which concentration may be interpreted as reflecting only employer market power is limited.

What data did you use to measure markdowns and concentration?

Macaluso: The paper uses the most comprehensive data we could find, the Census of Manufactures (CM), the Annual Survey of Manufactures (ASM), and the Longitudinal Business Database (LBD). The manufacturing databases, CM and ASM, are extremely rich and detailed data sources that record all inputs a firm uses to produce and their costs. They give an essential picture of how firms operate and help us investigate how workers' compensation relates to revenue produced. The LBD contains the universe of U.S. employers. That means every single establishment in the country with at least one employee. So it is a very comprehensive dataset. All three datasets come from a source that is extremely reputable, the Census Bureau.

We also supplemented the Census sources with Burning Glass Technology (BGT) data. BGT contains data on job openings, or vacancies. Vacancies are a measure of labor demand, the intent to hire, if you will. If we are concerned about employer market



power, we want to look at labor demand as well as equilibrium labor quantities because that's what the job seekers are facing. BGT also has some advantages over the Census datasets. One is that BGT contains occupation-level measures. The literature on the mobility of workers is mostly focused at the occupation level. So just looking at the industry can be misleading. The other advantage of BGT is that we don't have just the number of vacancies by industry or by location. We also have characteristics of those vacancies. With those, we are able to look inside the job openings and look at the skills, education, and experience that are required. That gives us a more complete picture of the jobs and how they might be affected by higher employer concentration. The main disadvantage of BGT, on the other hand, is that it has a limited time span (it starts in 2010 versus the Census data that are available since 1976), and it does not contain vacancies that are not posted online.

How do you think about geographical boundaries and the issue of sectoral mobility in your study?

Macaluso: Whenever you're trying to establish employer market power through concentration, the first step is to define the boundaries of the potential labor market. The larger the market is, the less an employer is able to keep employees there. If you define a market very narrowly as one firm, that firm almost mechanically has very high market power according to concentration. But if the boundary of the market expands, there will be more firms to choose from, which reduces the measures of market power that we pick up in the data. So deciding which set of firms should be included in a market is very important when studying concentration: the fewer we include, the more market power we are likely to find.

When we study markdowns, we don't have this problem because this measure of monopsony is independent of the number of firms we consider. In fact, we concentrate on the narrowest market, so to speak: each manufacturing plant and their workers. We then study how the wages earned by workers compare to the revenue produced. This is to establish

how many workers in the manufacturing industry are currently in a monopsonistic work relationship; that is, their wage does not increase one-to-one with each additional dollar of revenue produced by their employer. We find that the average worker earns fifty-six cents on the dollar, so monopsony is widespread when we look at existing employment relationships through the lens of markdowns.

When we look at local labor market concentration, we don't want to define the market using our measurement instrument, which is concentration. So we take data on mobility of workers along two dimensions that have been shown in the literature to be important for wages. The first dimension is geography. One way to escape a big firm is to move somewhere else. There is quite a bit of literature on how labor markets are very local. There is work that shows that about 80 to 85 percent of all applications that unemployed job seekers send out are within the same metropolitan area. And this is not just in the United States, but also in other countries. So we feel pretty confident that a metro area is a good geographical boundary.

The second dimension that we thought was important is sectoral mobility. Another way to escape your current market is to change your occupation. That would be one way to get out of the grip of an employer that is exercising market power. The literature, at least the micro literature, has favored very narrowly defined labor markets. We think that approach is a little bit too restrictive. In my previous work, I've shown that there is a lot of occupational mobility, especially for workers who go through a layoff. About 50 percent of job changes are at the two-digit occupational level. That would include a change from a server in a restaurant to a janitor in a hospital, for example. So very different industries, very different firms. What I've shown in my previous work is that although these are very different jobs, the skills that are involved in each are similar. In our study, we take a middle approach in how we think about labor market geography and occupational level, and we end up with 356 metro areas and 108 occupations.

Given that mobility between professions is substantially greater in the United States than in other countries, would your results be quite different if you were to look at this internationally?

Macaluso: That's a good question. I think there are many additional constraints in other countries that have made mobility much lower. Many other countries have much more expansive safety nets. France, Italy, and Spain, for example, have much less fluid labor markets. Granted, the fluidity of the labor market has been going down in the United States, and this is another reason why we were interested in studying monopsony and concentration. Can a rise in concentration explain different phenomena like the lack of fluidity? The answer turns out to be no, but I can see how the answer might be different if we carried out this study in other countries.

What is the relationship between national and local labor market concentration?

Macaluso: Some papers have proposed a national measure of concentration that considers the whole industry as a market. So, for example, a manufacturing firm in Ohio and a manufacturing firm in Texas are in the same market. There's another strand of the literature that looks at job mobility data and contends that this is not the right way to define markets. People generally do not look for jobs that are very far from home, and so we should consider a local measure of concentration.

We compute both national- and local-level concentration. When I say national, I mean industry-based, and local is geography-based, where geography is a metro area or a county. This is to check that our results do not hinge on a specific definition of geography, which they don't. The first thing we noticed is that these two measures have different trends. National concentration is increasing or flat depending on the measures you are taking, and local concentration is strongly decreasing. So there's a divergence there, and we asked ourselves, "How can this be true?" If you accept that these are measures of similar things, employer market power, how can one be trending up



and one be trending down?

We find that the relationship between market size and concentration has been weakening over time. It is still true that larger markets are less concentrated, but it's less true today than it was in the 1970s. Instead, we find that markets are becoming more and more alike across space. This also sheds light on why people aren't moving as much. There seems to be less incentive to do so. The firms in the industries that are present in different markets and different geographies are increasingly equal to each other.

You find that the average labor market is somewhat concentrated, but the average job is in a pretty competitive labor market. Could you explain how both of those could be true?

Macaluso: The policy debate on concentration in the labor market has been influenced quite a bit by the fact that most markets in the United States are highly concentrated. We already talked about how if you define a market very narrowly, it almost necessarily has to be concentrated. We also bring attention to another point, which is that when you look at the average market, you are looking at concentration on the land. If you divide the United States into administrative units, let's call them counties or commuting zones, and then you want to see how many of these units are concentrated, you draw a distribution over the land. But the land doesn't have employment market power. You have to look at where the workers

are. To some extent it's much less important that a geographic area's labor market is very concentrated if not so many workers are there. If it is densely populated, then there are a lot of workers who would be potentially impacted by employer concentration. That's why we look at the distribution of both jobs and workers. Once we do that, we find that very few jobs are in concentrated markets. In fact, we estimate that fewer than 5 percent of new jobs are in concentrated markets. That's because in places with very few firms, there are also very few workers. So the impact on workers is smaller.

I want to stress that this does not diminish the importance of monopsony and concentration on specific workers. First, we have the evidence on markdowns for manufacturing jobs. Furthermore, there are categories of workers who may be impacted quite a bit in specific geographies. The fact that these locations may be remote and less densely populated does not mean this is not a problem for them. What I would say, though, is that the data do not support the hypothesis that a majority of workers in the United States are subject to high concentration.

Does your research have any implications for policy?

Macaluso: We had policy in mind when researching this topic. First, we wanted to establish clear facts on the prevalence of monopsony in the U.S. labor markets. In our opinion, it is hard to come up with sound policies if they are not based on facts, so we wanted to provide solid ground to decision makers. That's why we use the most comprehensive data available, different measurement strategies, and state-of-the-art estimation techniques.

We also want to think about different types of interventions. Traditionally, there are three ways economists have proposed to limit employer market power. One of them is unions, which would add another distortion to the market by giving labor some market power as well. That has not been very successful in the United States. It's been more successful in Europe, perhaps for cultural reasons.

The other policy instrument that has come back quite a bit in the public debate is the minimum

wage. The claim is that because of employer market power, wages are depressed. That gets back to the markdown wedge we talked about before. One way to decrease that wedge is to increase the minimum wage. But our paper shows that you have to be careful; this has to be a local minimum wage, not a federal one. There are different levels of markdowns across geography, industries, and firms, so a federal minimum wage would probably not address the heterogeneity we find in the data. We also find that the negative association between concentration and wages is largest at the high end of the skill distribution, where there are high-wage workers who are not likely to be subject to the minimum wage anyway.

Even more surprising, we found that employers that are in concentrated markets tend to request more skills from their workers. An increase in the minimum wage would, if anything, exacerbate this problem because as the cost of labor goes up, it's much more likely that employers impose higher skill requirements to fill their vacancies.

The final policy instrument is antitrust. I personally think antitrust is a blunt instrument because of the variation in markdowns we see even across plants within the same industry and the difference between national and local concentration. Some labor markets are more like New York City, while other markets are more like Fargo, North Dakota. If we change the rules at the federal level, we are ignoring the difference between the two. And our paper shows that these differences are very important. So when thinking about policy, I think we have to be careful to do a local analysis, not a national analysis. ■



**FEDERAL RESERVE BANK
OF RICHMOND®**

Richmond • Baltimore • Charlotte