

Commentary on Trading in the Clouds

Gary Pulsinelli

Thank you, Professor Baker, for that very thought-provoking presentation. I commented, as she noted, on her paper when she was at this symposium a few years back, when she explored some related issues of banks migrating to the cloud. In preparing for this response, I went back and reread that paper from the symposium *Banking on The Cloud* and one paragraph in particular struck me when I read it: “A limited number of significant cloud service providers exist. This is typical of SIFMUs. Most SIFMUs, such as designated clearinghouses, are essentially natural monopolies.”

My reaction to that would have been, wait a minute—how is that a natural monopoly? Don’t natural monopolies typically arise out of natural conditions? When we think of a natural monopoly, we think of things like utilities, where it doesn’t make sense to have a lot of people build transmission lines or pipes or all those kind of things. It makes sense to have one person do it, so duplicative effort doesn’t make any sense. That wouldn’t really seem to fit cloud service providers. The ‘lines’ or ‘pipes’ for delivery are already there; that’s what the Internet does. It might be somewhat expensive to set up all the servers and security needed to make it work, but nothing on the scale of setting up a whole new electric grid.

It might even be possible for them to start small and get bigger over time.

How then could they be natural monopolies?

However, I've recently been introduced to a broader concept of that term through the auspices of my Law Science and Technology Seminar. One of my guest speakers, Micah Beck—a computer science professor, of all things—was exploring the concept. His particular focus was the possibility that web search might turn out to be a natural monopoly, because in theory, it shouldn't be that hard to set up a web search engine. All the data is out there; you've just got to pick it up and index it. So that would suggest it's not a natural monopoly under the traditional definition.

He, though, introduced me to the concept of 'subadditivity,' which is economists' current take on monopoly. The basic idea was originally expressed by William Baumol who defined a natural monopoly as “[a]n industry in which multi-firm production is more costly than production by a monopoly.” Expressed mathematically (because I know how much lawyers love mathematical equations):

If q^1, q^2, \dots, q^k are output bundles that sum to q , then a single firm is superior on efficiency grounds to a multi-firm industry if the following condition holds:

$$C(q) < C(q^1) + C(q^2) + \dots + C(q^k)$$

Where $C(q^x)$ represents the cost of producing commodity bundle q^x .

So all this really means is that, if the inequality holds, then a single firm can produce bundles q^1, q^2, \dots, q^k more cheaply than if they were produced by two or more firms.

[And yes this *will* be on the final exam.]

The basic point here is that in some industries, one firm can more efficiently produce the amount required to meet the entire demand than can multiple firms each meeting a part of the demand. If that's the case, then over time, one producer will tend to gradually drive out the others, resulting in a monopoly even in the absence of any monopolistic conduct, and thus it's a natural monopoly.

My guess is this might well apply to cloud service providers, much as Professor Baker suggested. Thus, the market is at best likely to remain in the hands of only a few providers and potentially over time even be reduced to only one. If the market does tend to natural monopoly in this way, that would support Professor Baker's call to regulate SIFMUs in both the banking and exchange context, as she argues today. With so few producers, a problem at any one of them has the potential to take down a significant portion of the exchanges business, with potentially devastating market consequences.

While regulation is a good idea in any case, it might also be useful to figure out if there is a way to break up those monopoly—can we do something

about them? Professor Beck's article has some interesting ideas in that line for web search, but they rely on the fact that the data Google accumulates comes from the public. So he says the public therefore deserves some say in what happens with it. That would be harder to apply to cloud services generally, although perhaps in the present instance the public interest enters through the concerns highlighted by Professor Baker.

Alternatively, I think maybe part of the solution lies in redundancy. If the exchanges (and banks) want to move stuff to the cloud, perhaps they should be required to keep it all in more than one place, in the hands of multiple cloud service providers. That way, if one of them goes down, everything is still preserved on the others. Now that concept already exists on a smaller scale in the form of the RAID arrays of disk drives. The idea is that we store identical sets of data on multiple disks at the same time, thus, ensuring that the failure of one disk doesn't compromise the integrity of the data, because the same data is stored in other places. The cloud providers thus wouldn't just store data on one cloud but store it on multiple clouds.

Now you might expect the cloud providers to have been opposed to that type of system—they don't want to deal with their competitors, they want to keep business to themselves. However, Professor Baker sent me a couple of interesting articles on this point. One is an article from the *Wall*

Street Journal headlined “Oracle and Microsoft Agree to Deepen Interoperability of Cloud Platforms,” suggesting that the providers are concluding that their clients, many of whom use multiple platforms, want them to work together and they’re going to have to do with their clients want, which should make a multi-provider scheme more feasible. The second article, from *Financial Times*, “Google Bets on the Cloud Breaking Up,” observes that Google is working on a ‘multi-cloud’ approach that “involves harnessing the resources of a number of different public clouds to handle a computing task. For customers, it reduces the risk of lock in by single cloud supplier.” If the regulators require redundancy, the providers will figure out a way to make it work.

One final point on redundancy: This analysis primarily applies to data storage, keeping all the data together, but the bigger problem, as Professor Baker noted, is functionality. That’s a little bit trickier, because the platforms don’t all run the same way. Maybe one way to look at this is the main cloud gets almost all of the functionality that the exchange needs, while the backup ones just get a much more limited setup, the really fundamentally necessary parts to keep the market going. That way you won’t have to maintain the whole package in multiple places but only a reduced package on one of them.

Finally, I'm not quite sure how these articles reflect on my natural monopoly contention. On the one hand, they might suggest 'commodification' of cloud storage, with the potential for multiple providers, even small ones. On the other hand, perhaps it means this is just the first step in de facto consolidation, if not actual monopoly; it may be that these actions will eventually result in one provider absorbing all the others and becoming the only cloud provider, or perhaps they'll be so intertwined that in effect they're behaving as only a single provider. Only time is going to tell us how that comes out (or perhaps federal regulators).

Thank you.