

July 22, 2008

Comments on
“Firm-Size Distribution and Cross-country Income Differences,”
by Laura Alfaro, Andrew Charlton and Fabio Kanczuk

Jeffrey Frankel
Harvard University and NBER
31st annual International Seminar on Macroeconomics
Bank of Slovenia, Ljubljana, June 20, 2008

The general motivation of the paper by Alfaro, Charlton and Kanczuk is an important one. Variation in income levels across countries is known to be explained less by inputs (physical & human capital), than by TFP (Total factor productivity). The difference is not necessarily technology. To take one set of examples, during the hey-day of communism, the countries of the Soviet Bloc tended to have high rates of investment in physical and human capital, and to have advanced technology – at least in certain areas, such as aerospace. Yet real income per capital lagged increasingly behind the West, due to the inferior model of social organization.

Laura Alfaro and her co-authors work in the promising new line of research that focuses on firm heterogeneity. At least it is new in trade theory (or “new new” to distinguish it from Helpman-Krugman New Trade Theory based on trade in imperfect substitutes). The question is how efficiently different countries allocate resources across firms. Some distort the allocation of resources across firms, reducing overall TFP.

In the Introduction, there is a good statement describing the policy areas that bear examination:

“These distortions should be interpreted as the different types of policies that might generate these effects such as non-competitive banking systems, product and labor market regulations, corruption, and trade restrictions. For example, specific producers might be offered, by governments, special tax deals and contracts financed by taxes on other production activities, and by non-competitive banking systems, favorable interest rates on loans based on non-economic factors, leading to misallocation of credit across establishments. Corruption and trade restrictions might also result in less productive firms obtaining a larger share of the market.”

One could add to the list industrial policy and the ability of government to address externalities, public goods, etc., though I am inclined to agree that corruption and intervention to benefit special interests are more likely to be problems. Anti-trust policy (competition policy) is the place where the public sector is most likely to be able to reduce distortions, as opposed to creating them. There is already a lot of research on

these topics, of course, but most of it is probably not phrased in terms of firm heterogeneity.

The authors have a huge data set, WorldBase, from Dun and Bradstreet. It includes 20 million firms in 79 countries. I always say, one can learn more from a big data set and simple techniques than from a small data set and fancy techniques. Whatever its flaws, it must be capable of some important contributions to the state of our knowledge.

I would love to see tests of such propositions as

- (1) Firms that are owned by well-connected elites tend to get easier credit, growing in size but not necessarily earnings, with an outcome that is either good or (more likely) bad for national productivity.
- (2) Firms in export industries tend to be especially productive and those that are protected against import competition tend to be especially unproductive.
- (3) In the high-technology sector, where innovation is all-important, being small is especially useful (firms need to stay nimble and competitiveness) or else the opposite (pharmaceutical companies need to be large and patent-protected, so that they can use profits from successful drugs to pay for long-term research in other areas that may not pay off).
- (4) Firms in certain sectors, like agriculture or heavy industry, are favored over others like services, with negative effects on aggregate TFP.
- (5) Firms that are highly regulated tend to be less productive than those that are not.
- (6) Firms that are privately owned tend to be either more or less productive than those that have to justify themselves to stockholders.
- (7) Firms in highly concentrated industries, although profitable, tend to have low productivity.

The reader who might be led by the beginning of this paper to expect something along these lines will be disappointed. These are not the sorts of hypotheses to be tested here.

The authors are concerned exclusively with one dimension of firm heterogeneity: the size distribution of plants. Lack of data explains why they don't test directly the proposition of interest: that large firms are less productive.¹

The authors could well point out that their data set does not come with the variables that would be necessary to test propositions like the seven above. But their narrow focus appears to be by choice, derived from some recent literature, which in turn originated specifically in two papers of outstanding pedigree: Lucas (1978) and Melitz (2003). I can see how a model of good versus bad size distributions could be interesting; but I object to leaping over all other kinds of firm heterogeneity to a framework in which firm size is assumed to be the *only* relevant dimension of heterogeneity. Imagine that all

¹ The first time through, I could not tell if this was the hypothesis, or that large firms are more productive. Or both, or neither? Or some more complicated hypothesis about the desirable size distribution. I still have trouble telling for sure.

firms are the same size, but some are well-connected or corrupt while others aren't; some are exporters while others are protected; some are in heavy industry while others are in services; some are in small oligopolistic industries while others are in large competitive industries; some regulated, some not; some privately-owned, some publicly traded, etc. No variations in productivity along these dimensions are admitted in this paper.

The logic, taken from other papers such as Restuccia and Rogerson (2007), is apparently that if resources are diverted across firms for any reason other than productivity – whether it is government subsidies or private market distortions – the result will be that bigger firms will on average be less productive. Clever model. Within such a model, I can indeed see how size dispersion might be a “sufficient statistic” for the extent of distortion.

But in reality, as opposed to inside a model, this leaves out a lot. It leaves out, for example, the seven questions listed above. Nor is this list anywhere near exhaustive. I am not sure that firm size would even be on my list of top ten factors that determine national productivity. Korea and Taiwan, two countries with similar starting points, both accomplished rapid productivity growth. Yet Korea did it via large firms and chaebols, while Taiwanese firm size was typically much smaller. That difference in size distribution hardly seems to have mattered.

It seems to me that the authors are not sufficiently aware of the leap, from a special model that focuses on heterogeneity of firms by size, to a reality in which size dispersion is the only kind of heterogeneity that matters. I infer this because the leap from firm-heterogeneity to size-dispersion still occurs in a single inconspicuous footnote, footnote 3: “The model implies a direct relation between productivity and size...” It isn't until after equation (13) that we are told the key assumption: “For the purposes of this paper, the distribution of plant-size is a summary statistic of the resource misallocation for each country.” After equation (17) the very strong assumption underlying the calibration exercise is revealed, “Our measure of success is based on the question: What would the dispersion of incomes be...?”

Their preliminary look at the data shows that most countries have a higher share of the economy in small firms than large ones, and that this is especially true of the United States. The authors take the US distribution to be the desirable norm. This decision reflects a general belief that the US tends to do things right, and a lack of obvious alternatives. But presumably it also reflects a belief that (i) the large role of small firms in the United States specifically reflects a relative absence of government subsidy and oligopolistic structure, which is desirable, (ii) while size might be desirable in some cases, such as a Microsoft or a Plough, the US economy has the right number of such firms, which is why the ideal structure is not taken to be atomistic, and (iii) conspicuous failures in the US system (Enron, Worldcom, mortgage brokers, Bear Stearns, subsidized farmers, etc.) are sufficiently rare compared to other countries that they do not negate the results.

The exercise appears successful. The paper summarizes its conclusions “We calculated the implicit distortion needed to generate size distributions consistent with size

histograms for a sample of 79 countries. We found the loss in output caused by these distortions to be quantitatively important. For our preferred calibration, the model explains 0.58 of the log variance of income per worker. This figure should be compared to the 0.42 success rate of the usual model.”

Even if we were to grant the proposition that size is a sufficient statistic for distortions, I still see two more difficulties. One is that the data are at the establishment level (plants), due to data availability -- whereas the taxes, subsidies and other distortions with which we are concerned are overwhelmingly at the company level.

The other difficulty is that calibration strikes me as a very weak tool for testing such a hypothesis. There is a big difference between saying that (i) a simulation based on countries’ observed deviations from the US size distribution is capable of explaining the observed variability in national incomes, and saying that (ii) the variation in size distribution is in fact associated with variation in income. To take a potentially alarming example, what if big firms are good, not bad, and the extent of deviations of countries’ size distribution from the US distribution is capable of generating the right income variance because it captures how *good* countries are, rather than how *bad* countries are? Would the authors’ calibration methodology be able to tell the difference? I hope so ! Otherwise, one might as well test the theory “crime doesn’t pay” by looking at the absolute difference between the wealth of criminals and the wealth of law-abiding citizens, and ignoring the sign.