Emulating others and fat-tailed distributions

Ben Klemens

4 April 2013

I finally posted my paper on leptokurtic (fat-tailed) outcomes to arXiv[1]. I started writing it as part of my dissertation, so it took over a decade to put out. I'm still not entirely happy with it—it's two subpapers that don't meet as well as I want them to:

- The first subpaper gathers the literature in finance about how returns are leptokurtic and that people in finance have real incentives to emulate each other.

- The second subpaper writes down a simple model where agents get direct benefit from emulating each other—their utility function is increasing in a private valuation of an action and in the percentage of others who are acting—and finds that leptokurtic distributions result.

I've always thought that emulation is such a blatantly clear facet of human behavior, and was struck by how the economics literature largely ignores it. Your typical econ textbook starts off by saying that agents have a utility function which we will take as the given starting point, making as few assumptions as possible about what's in that function. If people have a strong preference for giving away their money, the textbook advises that our only job as economists is to model the implications of that preference. But then we go to the journals, and we find a rather narrow range of utility functions assumed. They have to be out there somewhere, but I never found a paper in a mainstream econ journal that was willing to assume that agents get direct benefit from emulating others.

Nonetheless, emulative behavior is fundamental to being human: mothers emulate their babies, who emulate their mothers; con men and magicians devise tricks that depend on how it can sometimes take effort to *not* emulate another person; cliques of people who grow to behave similarly are abundant, and not just in high school; people pay good money for fashion magazines that instruct the reader on how to dress like everybody else—even what underwear is fashionable.

We can sometimes explain away certain aspects of emulative behavior as really being something more direct. For example, the actions of others imparts information, and gathering information is expensive. So if somebody else buys product A, that should provide information to you that product A is better. Of course, this works iff the other party has information, and isn't just imitating somebody else or making a random draw. There are *network externalities*, wherein a product you buy is worth more when

---

there are others on the same network or using the same standard, which produce a narrowly-rational motivation to emulate others.

Depending on the assumptions you want to make, the outcomes from the emulation-for-information or emulation-for-network story can be equivalent to the outcomes from the direct emulation story. The only difference is that those stories are special cases, and so applicable to fewer situations.

The broad intended lesson of my arXiv paper is that when you see a fat-tailed distribution, or a widening gap between the winners and losers, then a social emulation story should pop into your mind. Emulation causes a bit more push to the extreme outcomes—if, without any emulation, private tastes would lead to 65% taking some action, we can expect that if people have some slight preference to emulate the majority, then the same private tastes would lead to maybe 67% taking the action. Apply this to every level of action, and you've transformed your private-tastes distribution to have fatter tails.

There are some obvious cases where the herd-following story already comes up, especially in the case of blockbusters and flops, including movies (see the paper by de Vaney and Walls, cited in my arXiv paper) and cities (which are pretty obviously a network). But there are more subtle cases, and cases where a lot of parties want to pretend that emulation doesn't happen. The stock market is such a case, because none of the heroes in any Ayn Rand novels had any interest in emulating others. As I mention in the paper's conclusion, Congressional roll call votes provide another such case.

Further, here in the real world, equity traders and Congressfolk do have their own opinions, and we can't expect a lot of their positions to be based on pure emulation. Nor does the data demonstrate a bifurcated blockbuster/flop dichotomy in equity returns or vote counts. If we had a simple Central Limit Theorem at play, and people independently took positions, we'd expect a Normal distribution of returns or vote counts. Instead, we see distributions that are bell curves with fatter tails than the Normal. The paper demonstrates that we can explain this spreading-out via an emulative component to the decision process.

To give another example, I just ran across an article that finds a blockbuster-flop pattern in citations[2]. It tries to blame it on an information overload story, but it's also pretty easy to explain via emulative behavior, either due to information gathering or simple fashion.

[I found the paper via social means, from the blog[3] of someone who was once a fellow student. He adds a useful point about Gini coefficients (though applying his analysis to the explosion of publications seems to require that we assume that as publication became cheaper, the additional papers that appear are lower-quality and less citation-worthy).]

Alexander Schuessler, in his *Logic of Expressive Choice* (BUY![4]), is my favorite work that models humans as expressing identity by joining groups and emulating others. His title is an homage to Mancur Olson's *Logic of Collective Action*, which bangs its head against the wall for the length of a book asking why it is that people exert ef-

---

[2] http://wp.sigmod.org/?cat=18
[3] http://www.mathofpolitics.com/2013/04/02/inequality-smaller-ginis-can-fit-in-smaller-bottles/
[4] http://www.amazon.com/exec/obidos/tg/detail/-0691006628/qid=1120157199/sr=8-1/ref=pd_bbs_ur_1

fort to work for politically-oriented groups when the Prisoner's Dilemma predicts that they'd leave the work to others. Schuessler's response is that if you are willing to allow the possibility that people have some desire to be part of a social group, then there's not much of a paradox left.

There can be other explanations, and the fact that I've constructed a model which assumes emulative behavior and ends in a leptokurtic outcome doesn't prove that (1) emulative behavior must lead to leptokurtic outcomes or that (2) leptokurtic outcomes are always caused by emulative behavior. But it does demonstrate that the causal story that emulation $\Rightarrow$ fat-tailed outcomes is a possibility for a lot of contexts.

The other thing I hope readers get out of the paper is that agent-based models can be used to generate distributions, which we can then do statistics on as with any other distribution. I'm working on a paper that goes into immense detail on this point, and I'll try to post at least one example here over the next month.