The Average Location of Impact: A New Index and How to Use It

Research Summary

Loren K. Fryxell[1] and Charlotte Siegmann[2]

[1]Department of Economics and Global Priorities Institute, University of Oxford
[2]Department of Economics, MIT

July 24, 2023
(Latest Version Here)

Suppose you would like to communicate where most of the impact of an action or policy change accrues along some spectrum. This is an important characteristic for that policy in particular as well as for policies in general. For example, if most of the impact accrues to the rich, we may question if the poor are being unduly neglected in policy decisions; if most of the impact accrues in the short run, we may question if the future is being unduly neglected in policy decisions; and if most of the impact accrues in the long run, we may put less emphasis on empirical studies which measure that policy's impact in the short run.

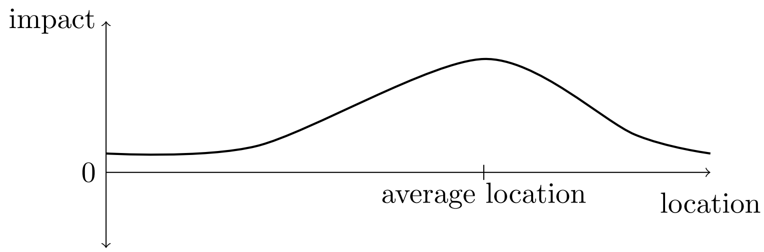

For this distribution, the *average location of impact*—with each location weighted by how much impact occurs there—is a good measure. If the spectrum is time, this figure communicates that most of the impact of the policy occurs in the future. If the spectrum is income level, this figure communicates that most of the impact of the policy is experienced by the rich.

However, an important feature about impact is that it may sometimes be *negative*. For example, if the spectrum is time, the following policy imposes a negative impact on the near future and a positive impact on the far future, and if the spectrum is income level, the following policy imposes a negative impact on the poor and a positive impact on the rich.

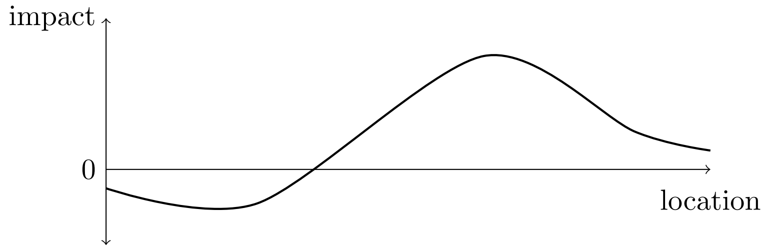

But an important feature of a standard average is that weights are *non-negative*.[1] The question then becomes:

> How can we compute an average location when some locations are weighted negatively?

The answer builds upon the following insight. There is nothing inherently special about an impact distribution which is everywhere non-negative. This does *not* mean that no costs were imposed. This simply means that at every point along the spectrum, the impact which accrues at that point is, on *net*, non-negative. For any policy which has a net positive impact overall, if we zoom out and accrue impact coarsely enough, the impact distribution will be everywhere non-negative.

We propose a generalized notion of the average, which we call a *trimmed mean*, which accepts negative weights. The trimmed mean simply zooms out the minimum amount, accruing impact just coarsely enough such that the distribution is everywhere non-negative, and then takes the standard mean of the resulting distribution.

Formally, we define the *trimmed mean* as follows. For any non-monotonic cumulative distribution function F (where at some locations impact is negative), find the non-decreasing cumulative distribution function F^* (where impact is everywhere non-negative) which comes *closest* to it in the following sense,

$$F^* \in \underset{\hat{F} \in \mathcal{F}^{\text{CDF}}}{\arg\min} \int_0^\infty \left| F(x) - \hat{F}(x) \right| \, \mathrm{d}x.^2$$

The trimmed mean, denoted $\mathbb{E}^T$, is the mathematical expectation of this function F^*, i.e., $\mathbb{E}^T(F) = \mathbb{E}(F^*)$. Our main result is that, if F is absolutely continuous, F^* exists and is unique—and hence the trimmed mean exists and is unique.

Graphically, the trimmed mean looks like this.

[1] For example, imagine computing the average height in a room with two people who are 6 feet tall and one person who is 5 feet tall. We weight 6 feet by 2 and 5 feet by 1 and all other heights by 0, for an average height of 5.66 feet. No heights are weighted by a negative number. For the standard notion of an average, negative weights are not allowed or considered.

[2] $\mathcal{F}^{\text{CDF}}$ is the set of all (non-decreasing) cdfs. We discuss why this notion of distance is appropriate in the paper.

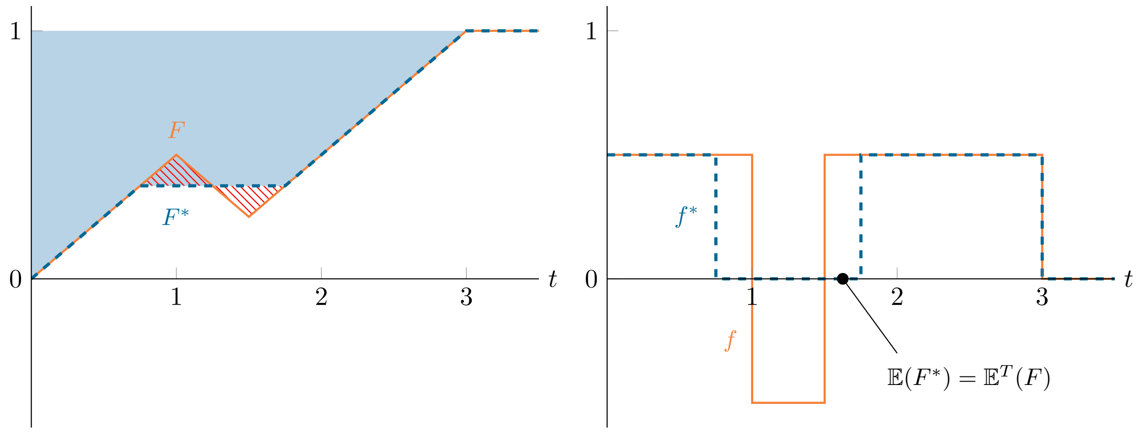

The area shaded in red is the distance between F and F^*. F^* is the non-decreasing function which minimizes the distance to F. The expected value of a (non-decreasing) cdf is equal to the area above the cdf and below 1, shaded in blue. This is the mean of F^* and the trimmed mean of F.

We can now return to our motivating example.

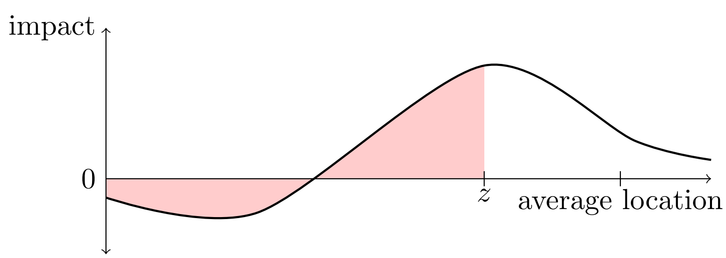

The trimmed mean zooms out just enough to consider the area in red, which contains zero net impact overall, as containing zero net impact at each location along the way. It then takes the mean of the resulting distribution, which is zero from 0 to z and unchanged everywhere else, giving us the average location of impact.