Corrigendum to: Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg. 2016. "Star Wars: The Empirics Strike Back." American Economic Journal: Applied Economics 8 (1): 1-32

April 6, 2026

By David Roodman, Abel Brodeur, Mathias Lé, Marc Sangnier, and Yanos Zylberberg

This corrigendum was cooperatively written by a commenter (Roodman) and the authors of the original article (Brodeur, Lé, Sangnier, and Zylberberg).

Section III of Brodeur et al. (2016) contains some errors in reasoning and code, which have minor effects on the paper's conclusions:

i. On pages 15–16 ("CLASS 1"), t tests with finite degrees of freedom are described as yielding "a mix of student distributions." They simply yield a single student's t distribution.

ii. In Panel A of Figure 2, the "Student(100)" density is plotted incorrectly because of a mistake in the code: The complementary cumulative distribution function is plotted rather than the probability density function. In fact, the t_{100} density is practically indistinguishable from that of the standard normal—whose density is correctly plotted.

iii. On page 16 (CLASS 2), the text states that the "Cauchy distributions […] extend the Gaussian/student distributions." While the standard Cauchy distribution is identical to the t_1 distribution, the full Cauchy family gains a scale parameter, such that it consists of scaled copies of the t_1 distribution. It does not, however, generalize the *full* t family; instead, it can be considered as a subset of the family of scaled t distributions. As a result, Figure 2 and the discussion thereof make an apples-to-oranges comparison between the t and Cauchy families. The finding that the Cauchy fits the empirical distribution better than the t is premised on giving scale flexibility only to the former.

iv. The text states incorrectly that a double-humped distribution of published results cannot be obtained if (a) the density of working paper statistics is monotonically declining, (b) the probability of publication is monotonically rising, and (c) there is no inflation (*p*-hacking). A counterexample is the case when the working paper distribution is standard normal, and the probability of publication is 0.1 for $|z| < 1.96$ and 1 otherwise. The distribution of published z statistics then has local maxima at 0 and 1.96.

What are the implications of these corrections? First, within the context of the paper, it remains the case that the scaled Cauchy distributions provide a better model for the published results than the unscaled t distributions, so there is no reason to change the interpretive weight placed on the results obtained using either family, e.g., in Table 2 and Figure 5.

Second, while the estimates in Table 2 are interpreted as floors on the amount of inflation (*p*-hacking) present in the published corpus, it remains possible that selection alone explains the

double hump in the distribution of published results. That said, we believe this is unlikely to be the case; and the published estimates, as floors, already err on the conservative side.

Finally, we note a minor labeling inconsistency in Table 1 of the original publication. The counts reported for the *Journal of Political Economy* (JPE) and the *Quarterly Journal of Economics* (QJE) are inadvertently swapped: the figures labeled as QJE correspond to JPE and vice versa. This discrepancy is purely cosmetic and does not affect any of the paper's analyses or conclusions. We thank Adam Gorajek for pointing out this labeling inconsistency.