

## [Measuring Productivity Dispersion: Lessons from counting one-hundred million ballots](#)

CFM-DP2017-25

Ethan Ilzetzki<sup>1,2,3</sup>, Saverio Simonelli<sup>4,5</sup>

<sup>1</sup>London School of Economics, <sup>2</sup> Centre for Macroeconomics, <sup>3</sup> Centre For Economic Policy Research, <sup>4</sup>University of Naples Federico II, <sup>5</sup> Center for Studies in Economics and Finance

Measuring output per worker is an important, yet challenging, task in economics. Workers' productivity is the main driver of income differences across countries, and its growth is the main proximate cause for economic growth over time (Caselli 2005). Yet measuring workers' output is less simple than it may seem at a first glance. While it is occasionally possible to observe workers' product directly in small scale settings, e.g. an individual production line, it is more difficult to isolate their value added from that of other factors in large scale settings, e.g. across an entire country. It isn't straightforward to separate workers' contribution from that of capital or other productive factors. Further, productivity in firms is often measured as revenue per worker, which may confound productivity with market power or other market imperfections (Syverson 2011)

We measure the productivity of electoral workers who counted ballots in a general election and two referenda in Italy. The objective isn't to understand productivity in the electoral process, but rather because this task is particularly useful in isolating workers' productivity in naturally-occurring data spanning an entire country. Using data on ballot counting times from the 2013 Italian general election and two referenda in 2016, we measure electoral volunteers' productivity in close to 8,000 municipalities. Combined, volunteers counted more than one-hundred million ballots. Each polling station had a fixed number of vote counters and polling stations were designed to minimize variation in eligible voters per station. Using observed turnout, we calculate the number of votes counted per person-hour: a direct, output-based, measure of workers' productivity. The task is managed at the national level and is uniform across the country, with identical guidelines in all polling stations, allowing a direct comparison of workers' productivity across municipalities. The task is simple, manual, and repetitive. There is virtually no physical capital or technology involved, and it would seem that only minimal education is required to count votes productively. Direct pecuniary incentives are identical for all volunteers and involve a lump-sum payment that is independent of performance or time-on-task.

We document large productivity dispersion across provinces in this very uniform task. In fact, the dispersion is slightly larger in vote counting than in firms. The vote counting productivity gap between northern and southern Italy is 28%, compared to a 20% north-south labour productivity difference in firms. That a regional productivity divide exists even in an identical task in a setting with rudimentary technology, virtually no physical capital, and similar pecuniary and institutional incentives, suggests large productivity differences “embedded”. Using a development accounting framework, this measure explains up to half of the firm-level productivity dispersion across Italian provinces and more than half the north-south productivity gap in Italy. We explore potential drivers of our measure of labour efficiency and find that its association with measures of work ethic and trust is particularly robust.