Extended Abstract

Please do not add your name or affiliation

Paper/Poster Title	Measuring output on intercropped plots: Evidence from crop cuts in Malawi

Abstract prepared for presentation at the 97th Annual Conference of the Agricultural Economics Society, The University of Warwick, United Kingdom

27th – 29th March 2023

Abstract	***200 words max***
We use randomized crop cuts to measure output on intercropped maize plots in Malawi. We document several empirical regularities. First, crop cuts on subplots leads to higher estimated yield than crop cuts on full plots. Second, intercropped plots are farmed more intensively than purestand plots across many inputs. Finally, compared to a "gold standard" adjustment strategy, we show that adjusting area for intercropping leads to vastly different estimates of yield for maize. These findings are important when considering the effects of intervention on the productivity of single crops on intercropped plots, which is a common occurrence in many parts of the world.	

Keywords	agricultural productivity, crop cuts, intercropping
JEL Code	Q10 see: www.aeaweb.org/jel/guide/jel.php?class=Q)

Introduction	***100 – 250 words***
We use randomized crop cuts to measure output on intercropped maize plots in Malawi. We document several empirical regularities. First, crop cuts on subplots leads to higher estimated yield than crop cuts on full plots. Second, intercropped plots are farmed more intensively than purestand plots across many inputs. Finally, compared to a "gold standard" adjustment strategy, we show that adjusting area for intercropping leads to vastly different estimates of yield for maize. These findings are important when considering the effects of intervention on the productivity of single crops on intercropped plots, which is a common occurrence in many parts of the world.	

Methodology	***100 – 250 words***
The crop cuts used in this paper were incorporated into the Malawi Fifth Integrated Household Survey (IHS5), implemented as a collaboration between the Malawi National Statistical Office (NSO) and the World Bank Living Standards Measurement Study (LSMS). The field work for the survey took place in 2019 and 2020, with the listing exercises taking place in late 2019 and the bulk of the enumeration taking place in 2020.	

The crop-cut enumeration areas (EAs) were a subset of all EAs included in the survey. A total of 72 rural EAs were randomly selected to be included in the sample, covering the 2019/2020 cropping season. Of these 72 EAs, half were assigned to a full-plot crop cut and half were assigned to a sub-plot crop cut. Within all 72 EAs, enumerators identified the universe of households with purestand (PS) maize plots or intercropped (IC) maize plots.

At the time of harvest, enumerators returned to the plot with the farmer and processed the crop cuts according to World Bank guidelines. On all plots, enumerators first harvested maize plants in each quadrant, collecting the following information:

1. total biomass (i.e. not discarding any portion of the plants)
2. number of harvested plants and cobs
3. weight of cobs prior to shelling
4. weight of grain resulting from shelled cobs
5. on intercropped plots only, the number of plants of each secondary crop within the quadrant at the time of harvest

For full-plot crop cuts, maize was then harvested on the remaining sections of the plot, following the same steps.

We consider the following adjustments to area:

- $A_{maize} = A_{full}$

  This uses the entirety of the plot/subplot/quadrant area to calculate yield. This amounts to no adjustment at all to the total area. We refer to this measure as the **unadjusted** measure**.**
- $A_{maize} = A_{full} \times P_{maize}$
  The survey asked respondents what proportion of the total area was dedicated to maize. For the second measure, we use this proportion, P_{maize}, to scale the area for maize. We refer to this measure as the **self-reported percentage** measure**.**
- $A_{maize} = A_{full} \times \frac{P_{maize}}{P_{maize} + P_{legumes} + P_{others}}$
  When responding about proportion, the total proportions reported by respondents sometimes total to greater than one. One possible explanation for this is that cropping areas can overlap within a plot. This option rescales the maize proportion based on the total proportion of each reported crop on the plot. We refer to this measure as the **self-reported ratio** measure**.**

Results	100 – 250 words

The overall rates (kg/ha) for many of the inputs are surprisingly similar across intercropping status, even though area is defined as the unadjusted (larger) measure for the intercropped plots. In fact, for some of the inputs, input intensity is actually higher on intercropped plots, even when we focus on maize-specific inputs, like seeds and plants. The maize plant rate on intercropped plots is around 20 percent higher than on purestand plots, and this difference is statistically significant.

The other inputs are not maize specific, but also show some differences. The use of inorganic fertilizer is slightly higher on intercropped plots, while organic fertilizer and pesticide use are relatively similar. Hired labor is slightly higher on intercropped plots, but not significantly so. Household labor, on the other hand, shows large, significant differences. Households apparently apply around 30 percent more person days of labor per hectare on intercropped plots than purestand plots. One possible explanation for this is that households with more intercropped maize plots are systematically different from households with more purestand maize plots. For example, if poorer households are more likely to intercrop and also face more labor market failures, they may allocate more household labor to their own plots than do richer households, who have more purestand plots. This would be consistent with previous research on market failures in Malawi (Brummund & Merfeld, 2022; Dillon et al., 2019), as well as in developing countries more generally (LaFave & Thomas, 2016).

Unadjusted total plot area is not significantly different from zero in the yield regressions. This is different from previous research which found a negative relationship between yield and plot size (e.g. (Barrett et al., 2010)) but is consistent with more recent research arguing that previous estimates were biased by reporting error (Gourlay et al., 2019). We use crop cuts in this paper, which is similar to the methodology in Gourlay, Kilic, & Lobell (2019) and may explain why we are unable to reject constant returns to scale in the yield regressions.

The estimated difference in yield on intercropped plots relative to purestand plots is dependent on the yield measure used. When not adjusting area (column one) or when using self-reported maize percentage (column five), intercropped plots have significantly lower yield than purestand plots. Apparently, the similarities between the unadjusted measure and the self-reported measure continue here. The correlation is also negative on the other three measures but is smaller in magnitude, especially for the self-reported ratio, which has a coefficient less than one-third as large as the coefficient in column one or column five. Again consistent with other tables, the two self-reported measures lead to vastly different conclusions.

Discussion and Conclusion	100 – 250 words

Having established that different area adjustments can lead to drastically different yield estimates, we next turn to a more practical issue: which alternative measure should we prefer? We believe that the adjusted plant ratio is the best measure to use when adjusting yield measures on intercropped plots. However, the number of plants is generally not asked on household surveys, like the LSMS. The other measures, on the other hand, are generally available on the LSMS and similar surveys.

The adjusted seed ratio is generally quite similar to a “gold standard” measure, indicating correspondence with the plant ratio, at yields below around eight on the graph, while the self-reported ratio is closest at higher yields. The unadjusted measure and the self-reported percentage are quite similar to one another; in fact, they are so close that it is often difficult to make out the separate lines. From the graph, it is not immediately apparent whether the self-reported ratio or the self-reported percentage/unadjusted measure is a better estimate when the plant ratio is not available, while it seems clear that seed ratio is the best choice for much of the distribution.

While both the self-reported ratio and the seed ratio significantly outperform the other two alternatives, the self-reported ratio clearly outperforms the adjusted seed ratio, and the difference is significant (p=0.075). Based on these statistics, it seems that the self-reported ratio is the preferred alternative if we assume the adjusted plant ratio is the gold standard by which to measure the alternatives.