Technical Annex

Annex 1. The ILO nowcasting model and global updating model for unemployment rates

The estimates and projections presented in this Monitor provide an update of the projections published in the WESO Trends 2023, and on ILOSTAT as the ILO modelled estimates, November 2022. The update is made using two separate methods, depending on country data availability. A direct nowcast and forecast for each individual country is made for those 61 countries that publish quarterly time series of the unemployment rate. The methodology is described in Appendix B of WESO Trends 2023. The remaining countries are not projected at the country level but as regional groups.

The methodology for this indirect approach estimates by how much the projected annual change in the unemployment rate differs because of newly updated data that drives those changes, compared to data that were available at the time of the WESO Trends 2023 estimates. This ensures that the revisions to the projection do not derive from a change in the methodology, but rather only from a change in the data.

Two types of data driving the dynamics of unemployment rates are used. The first is the projection of GDP growth for all countries within a region, which is a standard indicator to use for labour market projections. The second includes projections of unemployment rates of those countries where a direct nowcast and projection is made. The logic here is that there is some global interdependency of labour markets, the extent of which varies by region.

The target variable is the annual change in the weighted average unemployment rate of a region, or of the world, for countries without a direct projection. Those series are available for the time period 1991 to 2021, meaning a relatively short time period which limits the number of explanatory variables that can be used in a regression. Therefore, the dimensionality of the country data on GDP growth and unemployment rates is reduced using principal component analysis. The weighted unemployment rate is then regressed on the components. Multiple specifications are possible in terms of the number of components of each indicator. Furthermore, including too many regressors with the short time series creates the risk of overidentification, which could reduce the forecasting performance. For that reason, a leave-one-out cross-validation procedure is used in conjunction with Jackknife model averaging (see Appendix B of ILO 2023 [WESO Trends 2023] for details).

The averaged model specification is applied to the updated data of GDP growth and direct unemployment projections, and also to those that were available for the production of the projections in WESO Trends 2023. The difference between these two predictions shows the revision to the regional and global unemployment rates that is due to updated explanatory variables. This allows the derivation of the revised unemployment rate projections.

The indirect method is used for the world, for each (sub)region and for the income groups. Revised estimates of employment by region and income group will imply a change in the distribution of unemployment within each income group across regions, and within each region across income groups – although those categories are not estimated. To limit those changes, unemployment is adjusted. For each region-income pairing, unemployment is computed in two ways: 1) the share of unemployment in that income group within the region, as derived from the previous estimates (November 2022), is applied to the revised estimate of unemployment in the region; and 2) the share of unemployment in that region within the income group, as derived from the previous estimates (November 2022), is applied to the revised estimate of unemployment in the income group. Those two estimates are then averaged, which in turn allows to compute the adjusted estimates by region and income group. Furthermore, the approaches will deliver slightly different estimates for the global unemployment rate, and hence need to be made consistent. The final estimate is based on an average of the implied global unemployment from the regional and income estimation,[1] and from the global estimation.

Annex 2. Jobs gap estimates

The ILO jobs gap estimates provide regional and income group estimates for the sum of the unemployed, the potential labour force, and willing non-jobseekers divided by the sum of the extended labour force and willing non-jobseekers for the population aged 15 and older. These aggregate estimates are built by aggregating country-level data, which includes both nationally reported observations and imputed data for countries with missing data. The gender-specific country-level data used for the models includes the unemployment rate, unemployment-to-population ratio, the share of the extended labour force that is unemployed or in the potential labour force (LU3), and the economically inactive rate. The country-level data also includes the percentage of people aged 65 and older, log GDP per capita, and categorical variables for geographic region and levels of economic development.

The imputations for missing country data are produced with the predictions of five separate econometric models. The five models were chosen from an array of candidate models based on cross-validation, which selects the models with the highest accuracy in predicting the jobs gap rate in pseudo out-of-sample simulations. First, a model produces estimates from 2004 to 2019 for countries with at least one yearly data point of the jobs gap rate by sex. Second, a model produces estimates from 2004 to 2019 for those countries with no data on the jobs gap rate during the entire period. The third and fourth models are used to produce estimates for 2020 and the period of 2021 and 2022, respectively. The final model produces projections for 2023 using a nowcasting methodology. Drawing on available real-time economic and labour market data, the nowcast model estimates the historical statistical relationship between these indicators and the jobs gap rate and uses the resulting coefficients to predict how the jobs gap rate will change in response to the most recent observed values of the nowcasting indicators in Q1 2023. An indirect approach is applied for the remaining countries with no real-time data: this involves extrapolating the change in the jobs gap rate from countries with direct nowcasts. Since all the models estimate the jobs gap rate separately for the total population, women, and men, the aggregated estimates for women and men may be incompatible with the total population estimates. The subcomponents for women and men are adjusted proportionally to match the total population estimates.

Annex 3. Policy simulation: the effects of universal basic pension coverage for older persons

The ILO has developed a policy simulation to quantify the effects of implementing a social protection floor for older persons in the developing world, in terms of economic growth and other key development and social justice indicators. This annex describes the methodology used and detailed results.

The structure is as follows. First, we present the data used. Second, the main empirical strategy to identify the causal impact of a set of historical pension expansions on fertility, non-farm employment and GDP per capita is discussed. Third, we present the main results and describe how we translate these into the expected impact of achieving universal old-age basic pension coverage in all countries. Finally, the framework we use for quantifying the impact of pension expansions on other outcomes is explained.

[1] Following the adjustment procedure, the implied global unemployment from the regions and income groups is the same.

1. Data

We use two sources of data on old-age social protection. First, we rely on current social protection coverage data, SDG indicator 1.3.1. The coverage of older persons by country is made available through the ILO World Social Protection Database and discussed in detail in the World Social Protection Report (International Labour Organization 2022).

Second, to analyse the effect of historical pension expansions, we rely on data from *PensionWatch*, a knowledge hub for old-age pensions administered by *HelpAge*, a global network of 171 organizations. Their *Social Pensions Database* contains an exhaustive list of basic old-age pensions, including the year when they were introduced and the share of the population above the age of 60 covered.[2] For each of the 112 programmes, a list of references is available from which the information was extracted.[3] This sample includes many countries that were in the low- or lower-middle-income group at the time of the pension expansion and thus share many characteristics with the countries that will be of main interest in the simulation.[4]

In addition to the pension data, we obtain information for the outcomes of interest. We list below each outcome and the corresponding data source:

- GDP per capita: Penn World Tables for historical data, World Bank and IMF for current data.
- Share of non-agricultural employment, labour income by gender: ILO modelled estimates.
- Total fertility rate, life expectancy: UN World Population Prospects.
- Poverty, Income distribution: World Bank, Poverty and Inequality Platform (PIP).
- Expected years of schooling: CEDLAS, World Bank, UNESCO and UNICEF; via UNDP's Human Development Indices[5].

Finally, we collect data to construct a proxy for country-level gender norms using the following indicators and data sources:

- Sex ratio at birth: UN World Population Prospects.
- Variables related to legal inequality: World Bank, Women, Business and the Law index.
- Proportion of seats in parliament, ministerial positions held by women: Inter-parliamentary Union via World Bank.

2. Methodology

The goal of this analysis is to quantify the effect of universal basic old-age pension coverage. These are major political reforms that require a substantial fiscal expenditure. They are thus unlikely to be random policy shocks but rather a function of other macroeconomic variables, such as GDP growth. Estimating the macroeconomic impact of pensions requires constructing a credible counterfactual path had the country not introduced the old-age pension.

To construct this counterfactual, we rely on the synthetic control method introduced by Abadie and Gardeazabal (2003) and Abadie, Diamond and Hainmueller (2010).[6] This method relies on the pool of so-called donor countries that have not yet implemented a pension expansion. It assigns a weight to each donor country such that the distance between the expansion country and the combination of donor countries is minimized.

---

[2] We use the version of the *Social Pensions Database* published on March 1, 2018.
[3] See http://www.pension-watch.net/ (last accessed in April 2023).
[4] Examples include Bangladesh's *Old-age Allowance Programme* and Kenya's *Older Persons Cash Transfer* (OPCT).
[5] See https://hdr.undp.org/sites/default/files/2021-22_HDR/hdr2021-22_technical_notes.pdf for details (last accessed in April 2023).
[6] See Abadie (2021) for a comprehensive review of the synthetic control literature.

The distance is computed using pre-expansion characteristics. We use the outcome variable itself (i.e. fertility, non-agricultural employment share, or GDP per capita), the growth in the outcome variable, and population.

Following the steps above, we obtain a counterfactual evolution of the outcome variable for each expansion. We can then compare this with reality. Computing the difference between real data and the counterfactual, averaging over all expansions gives us the average impact of the pension expansions in our sample.

More formally, following the notation in Abadie (2021), consider a sample of $J+1$ countries and suppose that the first country ($j=1$) is treated, i.e. implements a pension expansion, whereas the other countries ($j=2,\dots,J$) have not. Let there be K predictor variables and X_{kj} the value of variable k for country j. The predictors are the average in the pre-expansion period of the outcome itself, growth in the outcome and population ($K=3$). Let w_j be the non-negative weight for country j and let $\sum_{j=2}^{J+1} w_j = 1$. Then, the weights are chosen to minimize:

$$\sqrt{\sum_{k=1}^{K} v_k (X_{k1} - w_2 X_{kh} - \dots - w_{J+1} X_{kJ+1})}$$

where $v_1,\dots,v_K$ are non-negative and capture the relative importance of the different predictors (see Abadie (2021) for details). Denoting Y_{jet} the outcome for country j and expansion $e=1,\dots,E$ at time t, and w^* being the weights that minimize the distance metric above, the estimated treatment effect at time t and expansion e is:

$$\hat{\tau}_{et} = Y_{1et} - \sum_{j=2}^{J+1} w_{je}^{*} Y_{jet}$$

We repeat the above for each pension expansion and average over the individual estimated treatment effects to obtain the average treatment effect at time t:

$$\hat{\tau}_t = \frac{1}{E} \sum_{e=1}^{E} \hat{\tau}_{et}$$

Note that the synthetic control method was originally applied to cases with a single treated unit, but it has been increasingly applied to multiple treated units. A recent example in the empirical macroeconomics literature is Funke, Schularick and Trebesch (2022) who use this to study the impact of electing populist leaders on GDP and other economic outcomes. To interpret our estimates as causal, we assume that country implementation of a pension expansion conditional on the three target variables is due to factors unrelated to economic development and not in anticipation of future shocks to the economy.

Since using the synthetic control method in cases with multiple treated units is still a relatively novel approach in the literature, work on how to conduct statistical inference is limited. For the baseline results, we follow Funke, Schularick and Trebesch (2022) and plot the average pre-expansion standard deviation, which gives a sense of how closely aligned the counterfactual and the actual data are prior to the expansion. Intuitively, if the counterfactual evolves similarly to the actual data before the expansion, a large divergence after the expansion is suggestive of a significant treatment effect.

2.1 Sample restrictions

We now discuss restrictions to the sample of pension expansions that are necessary to apply the synthetic control method. If a country appears multiple times in the pension expansion data, we only keep the earliest expansion. For some observations the year of the expansion is unknown. We classify these countries as having an expansion throughout the whole period of analysis such that they do not enter the control group of non-expansion countries. We end up with 84 cases for which we observe the expansion year.

In our main empirical strategy we use a window of 10 years of data before and after each expansion. This further restricts the sample to expansions between 1960 and 2009, which finally leads to a set of 62 newly introduced pension schemes.

As discussed in Abadie (2021), in cases where a closely matching counterfactual cannot be constructed from the donor pool, the synthetic control method should not be applied. We thus drop badly matched cases from our sample.[7]

3. Results of historical pension expansions

To report results of the historical pension expansions we normalize the outcome variables to the period before the expansion $t-1$. For each of the outcomes, we plot the average of the pension expansion countries and the corresponding synthetic countries. In addition, we show the difference, i.e. the average treatment effect.

3.1 Fertility

We find that fertility decreased by 8.5 per cent in expansion countries compared to synthetic countries within 10 years on average. Figure 3.1 shows that the synthetic country closely matches reality before the expansion. After the expansion there is an impact on fertility that steadily grows over time. In Figure 3.2 we plot the 10-year impact of each individual expansion against initial fertility. We find the absolute treatment effect to be roughly linear. Countries with high initial fertility saw TFR decrease by .5, whereas countries with a fertility rate around the replacement rate of 2.1 saw very little if any fertility decline.

Figure 3.1. The effect of pension expansions on the total fertility rate

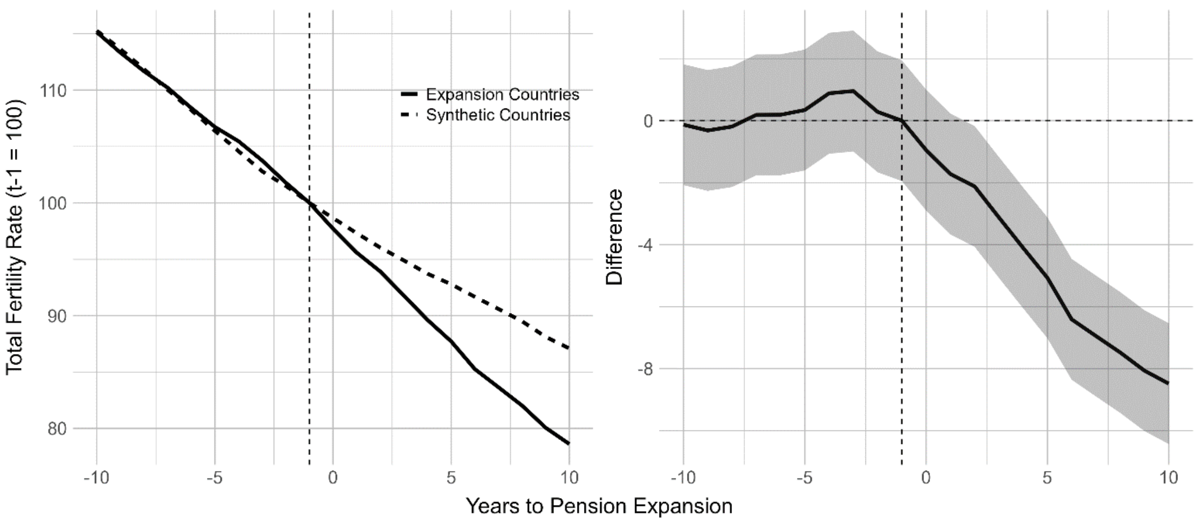

[7] A pension expansion is dropped from the sample when one of the yearly normalized pre-expansion differences between actual and counterfactual data is larger than a specific threshold. For GDP per capita and fertility, this value is set to 0.2 and for the non-farm employment share to 0.1, such that we drop roughly the 5 per cent of worst matched cases for each of the outcomes. In practice, varying these thresholds does not change the overall results substantially, but improves the fit of the counterfactual in the pre-expansion period.

Figure 3.2. Change in the total fertility rate depending on initial level

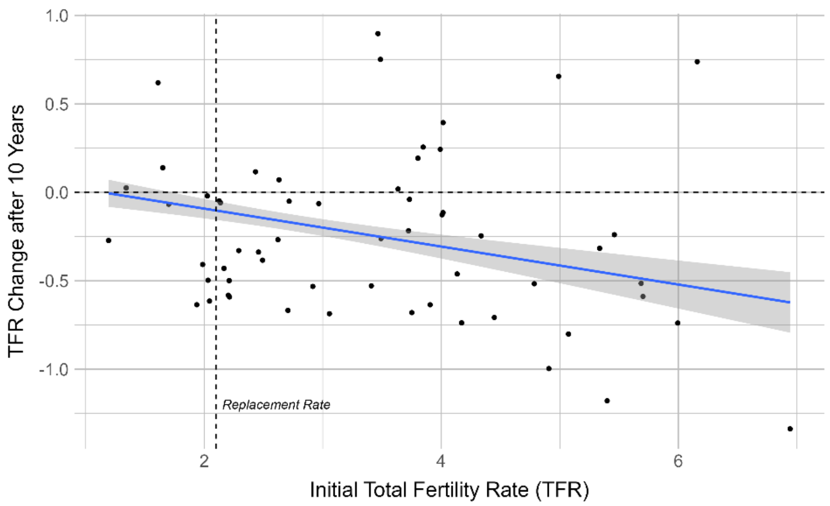

Our findings are broadly in line with the existing literature. Rossi and Godard (2022) study the effect of Namibia's Pension Act (adopted in 1992) on fertility. They rely on pre-reform variation in coverage across both ethnic groups and regions that is eventually eliminated by the reform. They divide the population into multiple groups depending on how far their expected pension benefits are from the poverty line. Fertility declines estimated in this manner point to large effects of 0.45 or more.

Another pension scheme that has recently been studied in the literature is China's rural pension scheme that has been found to decrease the number of children by 0.08-0.17 (Shen, Zheng and Yang 2020). This shows that even when policies that heavily restrict family size are in place and initial fertility is low, pensions can have an effect. Finally, Danzer and Zyska (forthcoming) find a decrease in the number of children per woman by 1.3 within 20 years after the introduction of Brazil's rural pension scheme.

Overall, our estimate that fertility decreased by around 0.5 within 10 years in high-fertility settings is thus in line with the literature. Our contribution lies in confirming this result using a large set of pension expansions from a variety of countries. This also allows us to trace out the relationship between pre-expansion fertility and the subsequent fertility decline which has not been possible previously.

3.2 Non-agricultural employment share

It is important to note that for the non-agricultural employment share, we partially rely on imputed data. These estimates should thus be interpreted with caution, and we complement them with a country-level case study. As Figure 3.3 illustrates, we find a 10-year decrease in the agricultural employment share by 7 per cent in expansion countries compared to synthetic countries. The pre-expansion period shows that it is not always possible to construct a reasonable counterfactual for non-agricultural employment. Instead, there are slight differences between treatment and control prior to the expansion. In Figure 6.1 we show that if we drop three cases where the pre-expansion counterfactual is relatively far from the actual data, the pre-expansion fit improves while the post-expansion effect remains very similar.

Figure 3.3. The effect of pension expansions on the non-agricultural employment share

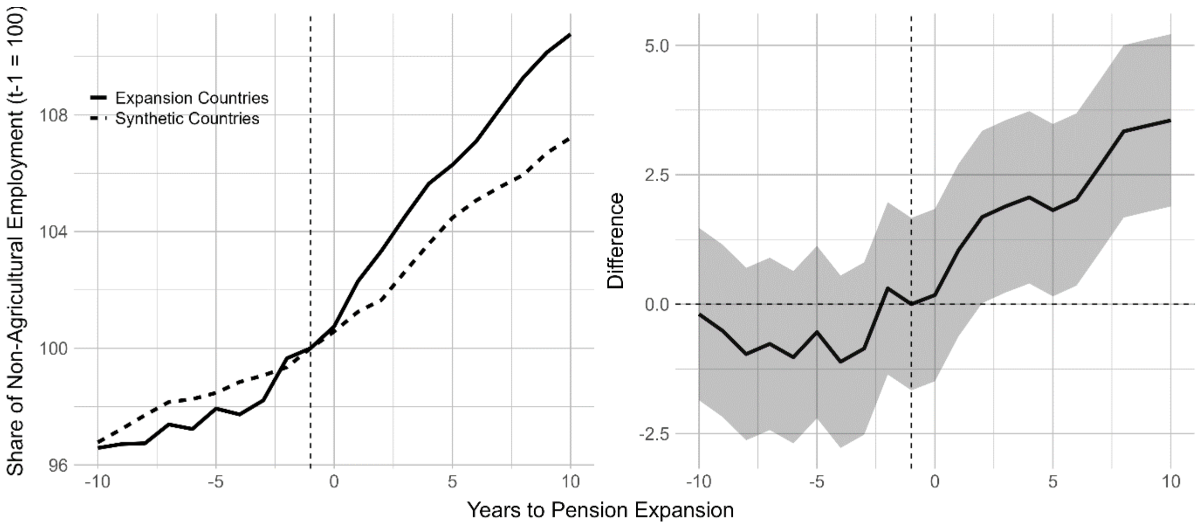

The existing literature on how social protection affects sectoral employment is limited. A notable exception is a recent study of China's rural pension scheme (Huang and Zhang 2021). The authors find that adults younger than age 60 are 3.3 percentage points more likely to engage in nonfarm work after the pension expansion. They hypothesize that this is driven by an increase in labour supply to pay pension premiums. However, we find evidence for this effect using our database of non-contributory schemes, suggesting that other mechanisms are at play.

To further investigate whether pension schemes affect sectoral employment shares, we study Namibia's pension scheme. Relying on the same identification strategy as Rossi and Godard (2022), we find that participation in non-agricultural employment increased following expansions of pension coverage. In the control group, among social groups that experienced low increases in pension coverage (below the median), the increase in non-agricultural employment between 1994 and 2010 was 7 per cent. In contrast, the treatment group that faced the largest expansion of pension coverage (at or above the median) saw an increase of 24 per cent. Before the reform the incidence of non-farm work across the groups was quite close, 53 and 47 per cent respectively.[8]

Hence, while the literature is still nascent, the different pieces of available evidence show that pension expansions strongly accelerate the transition from farm to non-farm work.

3.3 Economic growth

We find large effects of pensions on fertility and non-agricultural employment. It is therefore reasonable to expect a large effect on economic growth and living standards. As Figure 3.4 shows, we find that GDP per capita increased by 10.7 per cent in expansion countries compared to synthetic countries within 10 years on average.

[8] In a regression of non-agricultural work with year, cluster, age group, and education controls, the interaction term of the coverage expansion by reform dummy is significant at the 5 percent confidence level. The ILO Harmonized Microdata repository has been utilized for this exercise. Only surveys for 1994 and 2010 are used due to data availability.

Figure 3.4. The effect of pension expansions on GDP per capita

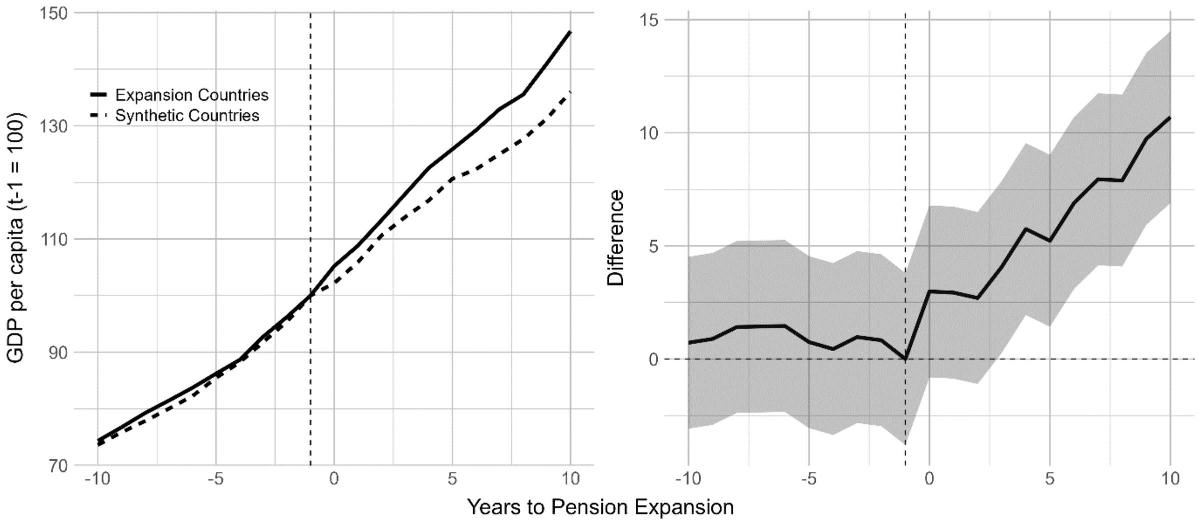

The empirical evidence of how pensions can affect GDP per capita is limited. Purely empirical studies that focus on one specific country typically struggle to quantify effects on GDP, as there can be various types of spillovers and general equilibrium effects that comparisons across different regions or population groups fail to capture. Our identification strategy is not subject to this issue as we can look at macroeconomic outcomes at the country level. To back up our synthetic control estimates with further evidence, we use the estimated changes in fertility and sectoral employment shares to calculate the implied GDP increase considering only these two factors in section 4.5.

3.4 Estimating the impact of universal coverage

Thus far, we have calculated the average effect of past pension expansions. To quantify the benefits of increasing coverage by a certain percentage, the average treatment effect needs to be scaled by the actual coverage of the different pension expansions. To do so, we compute the average coverage of the pension expansions in the estimation sample using the information contained in the dataset of historical pension expansions.[9] The sample varies slightly across outcomes due to data availability, but coverage tends to be around one-third. Dividing the average treatment effect by the average coverage gives us the effect of going from no to full coverage. Finally, we multiply this scaled effect by the current coverage gap obtained from the SDG indicator for each country and outcome.

For fertility, we proceed slightly differently since the effect of the pension expansion depends strongly on the initial level. There is an approximately linear relationship between initial fertility and the effect of the expansion. We thus fit a linear regression model to the individual effects (which is shown in Figure 3.2). For countries that are already below the replacement rate of 2.1, the change in fertility is assumed to be zero. In addition, we impose a lower bound for our estimated fertility rate at 1.42, the 5th percentile of the historical fertility distribution.

For the change in sectoral employment shares, we use the estimate to compute a constant relative decline in agricultural employment.

[9] For around 12 per cent of pension expansions, data on coverage is not available. They are dropped when computing the average.

Since we think of the demographic transition and sectoral change as the two main mechanisms through which pensions affect economic growth, we adjust the direct GDP per capita estimates in a second step. We regress the original GDP estimate on the employment and fertility estimates and use the fitted value from this regression as the adjusted estimate (see Table 3.1). There is no constant in the regression, which implies a zero effect on GDP in the absence of a change in fertility or the sectoral employment shares. However, this also implies that the mean of the fitted values is no longer equal to that of the original data. Here, we find an estimated increase in GDP per capita with an unweighted average that is around one third lower than the average original effect.

Table 3.1. Adjusting GDP per capita effects

Dependent variable:	GDP per capita effect
Agri. employment share effect	-2.271*** (0.3498)
TFR effect	-0.1005** (0.0436)
Observations	145
R^2	0.11441
Dep. var. mean	0.18928

3.5 Robustness

Time placebos

One placebo experiment commonly conducted in the literature is to shift the treatment back in time (see e.g. Funke, Schularick and Trebesch (2022)). We assume that all pension expansions happened 5 years earlier than they did and use the years 10 to 6 before the actual expansion to construct the synthetic control. If the pension expansions had a causal impact, we expect to see no effect prior to the expansion.

Figures 6.2, 6.3 and 6.4 show that, while using only 5 years of pre-expansion outcomes to construct the counterfactual negatively affects its fit, countries start to diverge from their synthetic counterpart only after implementing the expansion. This lends support to the causal interpretation of our results.

Linking fertility and non-agricultural employment with economic growth

While the synthetic control approach yields a direct estimate of pension expansions on GDP, considerable uncertainty remains, in part because empirical studies on this topic are scarce. This section quantifies the effect of demographic and sectoral change on GDP using various techniques from the economic literature. This allows us to argue that the estimated change in these two key outcomes alone would result in a large increase in GDP that is similar to our direct estimate.

One important strand of literature studies the effects of fertility declines on living standards, often referred to as the *demographic dividend*.[10] This literature typically highlights the change in the age dependency ratio as the primary direct effect of lower fertility on growth. As the share of the population of working-age increases, the relative size of employment grows and output per capita is higher.

[10] See Bloom et al. (2003) or Aiyar and Mody (2011).

To understand how the change in fertility affects economic outcomes, we use a local projection approach, as popularized by Jordà (2005), to estimate impulse response functions for GDP per capita and TFR for a TFR decline today. As commonly done in the empirical macro literature on Structural Vector Autoregressive (SVAR) models, we impose a short-run restriction for identification. We allow fertility to have a contemporaneous impact on GDP per capita but not vice versa. We use this restriction as fertility can be reasonably assumed to be pre-determined one year in advance. We find a 10-year fertility elasticity of GDP per capita of around -0.23, meaning that a 1 per cent reduction in fertility today will increase GDP per capita by 0.23 per cent in 10 years' time. This effect increases to around 0.63 after 20 years, indicating that the fertility effects take longer to materialize. Multiplying our estimated relative fertility decline with this number and averaging across countries shows that the predicted effect due to fertility changes amounts to around 10 per cent of the baseline GDP effect after 10 years, rising to almost one third after 20 years.

We then turn to reallocation of labour across sectors as the other driving mechanism of GDP increases. We compute the ratio of value added in the agricultural and non-agricultural sectors, respectively. This serves as an approximation of the relative increase in GDP per capita of one worker moving from agriculture to another sector. As for some countries this ratio reaches very high values, we instead use the median, which is around 2.1, and apply it to all countries. Finally, we multiply this with the share of workers we predict to move from agriculture to non-agricultural jobs. This yields an estimated GDP per capita increase that accounts for around 78 per cent of the original estimate.

One concern might be that when a large share of workers shifts away from agriculture, average labour productivity in the non-agricultural sector declines – as workers might not have the right skills, education or training. However, Gollin, Lagakos and Waugh (2014) find that large differences in labour productivity across the sectors remain even after accounting for education, literacy, and other potential drivers. Moreover, Bustos, Caprettini and Ponticelli (2016) show that a labour-saving productivity shock in Brazilian agriculture led to large increases in manufacturing employment, while wages in that sector declined only modestly. Assuming that wages are roughly proportional to productivity, this implies that average labour productivity was only marginally affected. Hence, the evidence suggests that shifts outside of agriculture are expected to have a strong positive impact on aggregate labour productivity.

Table 3.2 summarizes the results of the multiple robustness exercises. We conclude that changes in non-agricultural employment can explain a sizeable share of the GDP effects we estimate. While the effects of fertility declines tend to be smaller in the short run, we expect them to grow in importance at longer horizons. Combining the effects of fertility and sectoral employment change on GDP per capita, we find an estimated increase that is similar to our baseline.

Table 3.2. Explaining GDP effects with demographic and sectoral change, developing countries

	Estimated change in GDP per capita	Per cent of the baseline estimate
Baseline estimate	946	
Sectoral change (country-level value-added)	1285	135.8
Sectoral change (median value-added)	746	78.8
Local projection analysis (10 years)	95	10
Local projection analysis (20 years)	260	27.4

4. *Additional effects*

Based on the estimates for our three targeted outcomes from above, we want to derive estimates for the effects of pension expansions on other relevant economic and social outcomes. Since we generally lack long time series data for these outcomes, the synthetic control method cannot be applied, and we instead use tailored estimation approaches for each set of indicators as outlined below.

4.1 Poverty and inequality

While time series data for poverty and inequality exists, it often relies heavily on imputation. Focusing only on real data points, we run panel regressions with different measures for poverty and inequality as the dependent variable. First, to capture the well-established direct effect of pension expansion on poverty and inequality, we include a pension coverage measure. Second, we want to capture the more indirect effects of pension expansions that come from the general improvement in economic development. Introducing GDP, fertility, and employment shares together as regressors results in multicollinearity. Instead, we extract the first principal component and use it as the variable summarizing the main indirect effects of the pension expansion, leading to the following specification:

$$y_{ct} = \beta * PC1_{ct} + \gamma * Coverage_{ct} + FE_c + FE_t + \varepsilon_{ct}$$

Where y_{ct} is the outcome variable for country c at time t, which is either the share of the population below the poverty line or the income share of the bottom 40 per cent, the top 10 per cent, or the remainder. $PC1_{ct}$ is the first principal component of GDP per capita, fertility, and the agricultural employment share. $Coverage_{ct}$ is a variable capturing social pension coverage. Prior to the expansion year this is set to 0 and is equal to the coverage of the social pension thereafter. FE_c and FE_t are country and time fixed effects, respectively. Standard errors are clustered at the country level. This specification boils down to a difference-in-difference design with a continuous treatment variable, $Coverage_{ct}$.

Tables 4.1 and 4.2 show the results for poverty and income inequality, respectively. Pensions have both a strong direct effect and indirect effect by improving overall economic development as captured by $PC1_{ct}$. Coefficients associated with pension coverage are statistically significant for two widely used international poverty lines and the income share of the bottom 40 per cent, suggesting a strong direct redistributive effect of pension expansions. We use the estimated changes in the variables contained in $PC1_{ct}$ and the coverage gap to predict how poverty and inequality evolve with full old-age pension coverage.

Table 4.1. The effect of pension expansions on poverty

Dependent variable:	Share below 2.15 USD	Share below 3.65 USD
PC1	-0.0894*** (0.0119)	-0.1221*** (0.0166)
Social pension coverage	-0.0528*** (0.0089)	-0.0979*** (0.0184)
Fixed effects:	-------------------	-------------------
Country-Spell	Yes	Yes
Year	Yes	Yes
S.E.: Clustered	by: Country-Spell	by: Country-Spell
Observations	1,022	1,021
R2	0.96898	0.97473
Within R2	0.27213	0.23585
Dep. var. mean	0.05655	0.12936

Table 4.2. The effect of pension expansions on income inequality

Dependent variable:	Income share bottom 40 per cent	Income share 41th-90th percentile	Income share top 10 per cent
PC1	0.0293*** (0.0062)	0.0287*** (0.0057)	-0.0579*** (0.0083)
Social pension coverage	0.0276** (0.0121)	-0.0055 (0.0115)	-0.0221 (0.0167)
Fixed effects:	----------------------	-----------------------------	-------------------
Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
S.E.: Clustered	by: Country	by: Country	by: Country
Observations	1,046	1,046	1,046
R2	0.96859	0.92693	0.96803
Within R2	0.28923	0.24255	0.36275
Dep. var. mean	0.18036	0.52744	0.2922

4.2 Education and health

We can expect the transformational impact of pensions on the economy to spill over to education and health. Lower fertility rates can incentivise women to invest more in education and can improve their health outcomes. However, these effects are unlikely to materialize immediately and hence require a methodology different from the synthetic control method.[11] To simulate the effects on education and health, we rely on the following specification:

$$y_{ct} = \beta * PC1_{ct} + FE_c + FE_t + \varepsilon_{ct}$$

where the variables are defined as in section 0. The only difference is that we assume that there is no direct effect of pension coverage on health and education outcomes, but rather all the effect is going through the variables subsumed in $PC1_{ct}$.[12] For education, y_{ct} are the expected years of schooling for women and men, either separately or the ratio of the two.[13] To measure health outcomes, we rely on life expectancy. Tables 4.3 and 4.4 show that both schooling and life expectancy increase for women, as pensions affect fertility, non-farm employment, and GDP. The estimated coefficients are statistically significant. Life expectancy also increases for men. For both health and education, the ratio of women's and men's outcomes increases, meaning that women benefit more.

[11] Consistent with this, when applying the synthetic control method, we found no significant effects for health and education outcomes.

[12] If we include the coverage variable, it is positive but not statistically significant at conventional levels and we thus exclude it. Intuitively, while old-age pensions can have direct redistributive consequences, there is uncertainty regarding potential direct effects on education and health.

[13] Expected years of schooling are measured as the sum of the age-specific enrolment ratios in all levels of education.

Table 4.3. The effect of pension expansions on schooling

Dependent variable:	Expected years of schooling ratio	Women's expected years of schooling	Men's expected years of schooling
PC1	0.0322*** (0.0112)	0.4641** (0.2119)	0.2726 (0.2096)
Fixed effects:	----------------------------	-----------------------------	----------------------------
Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
S.E.: Clustered	by: Country	by: Country	by: Country
Observations	2,316	2,316	2,316
R2	0.8806	0.94105	0.92644
Within R2	0.0639	0.02111	0.00844
Dep. var. mean	1.0257	14.066	13.647

Table 4.4. The effect of pension expansions on life expectancy

Dependent variable:	Life expectancy ratio	Women's life expectancy	Men's life expectancy
PC1	0.0149*** (0.0027)	1.600*** (0.2877)	0.5289* (0.2806)
Fixed effects:	----------------------------	---------------------------	-------------------------
Country	Yes	Yes	Yes
Year	Yes	Yes	Yes
S.E.: Clustered	by: Country	by: Country	by: Country
Observations	2,419	2,419	2,419
R2	0.91745	0.97582	0.97589
Within R2	0.11973	0.12399	0.01545
Dep. var. mean	1.0845	76.791	70.876

4.3 Labour income by gender

As pensions have transformative effects for economic development, they will likely lead to different labour market outcomes by gender. For example, fertility and women's labour market decisions are typically closely intertwined. Contributing family work is prevalent among women and this form of work is common in agriculture. We thus look at the ratio of labour incomes of women and men, respectively.[14] The labour income gap for country c is defined as:

$$Labour\ Income\ Gap_c = \frac{Women's\ Total\ Labour\ Income_c}{Men's\ Total\ Labour\ Income_c}$$

[14] See ILO (2023) for a discussion of labour income by gender.

Time series data for the labour income gap are scarce and present a high level of statistical noise. We thus cannot apply any of the approaches that require time series data. Instead, we rely on the following cross-sectional specification:

$$Labour\ Income\ Gap_c = \beta * PC1_c + \theta * Gender\ Norms_c + FE_{r(c)} + \varepsilon_c$$

where $PC1_c$ is defined (as previously) as the first principal component of our three primary outcomes in country c and $Gender\ Norms_c$ is the first principal component of a variety of measures for progressive gender norms and laws.[15] This captures long-standing differences between countries in the extent and conditions of women's participation in the labour market, which are not controlled for in the absence of country fixed effects. We do control for region fixed effects ($FE_{r(c)}$).

Table 4.5 shows that our main drivers are positively associated with higher relative earnings by women. Again, we take the estimates obtained previously to predict how $PC1_c$ and subsequently $Labour\ Income\ Gap_c$ would evolve.

Table 4.5. The effect of pension expansions on the labour income gap

Dependent variable:	Gender gap in labour income
PC1	0.0696*** (0.0182)
Gender norms (principal component)	0.0006** (0.0002)
Fixed effects:	--------------------------
Region	Yes
S.E.: Clustered	by: Country
Observations	111
R2	0.6037
Within R2	0.28085
Dep. var. mean	0.53937

[15] The variables are the proportion of seats in parliament held by women, the proportion of ministerial positions held by women, the sex ratio at birth, length of paid maternity leave, indicator if a woman can be head of household in the same way as a man, the pay indicator score, and the marriage indicator score from the World Bank's Women, Business and the Law indicators.

5. References

Abadie, Alberto. 2021. "Using synthetic controls: Feasibility, data requirements, and methodological aspects." *Journal of Economic Literature* 59 (2): 391-425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." *Journal of the American Statistical Association* 105 (490): 493-505.

Abadie, Alberto, and Javier Gardeazabal. 2003. "The economic costs of conflict: A case study of the Basque Country." *American Economic Review* 93 (1): 113-132.

Aiyar, Shekhar, and Ashoka Mody. 2011. "The demographic dividend: Evidence from the Indian states." *IMF Working Paper Series.*

Bloom, David E, and David Canning. 2008. "Global Demographic Change: Dimensions and Economic Significance." *Population and Development Review* 34: 17-51.

Bloom, David, David Canning, and Jaypee Sevilla. 2003. *The demographic dividend: A new perspective on the economic consequences of population change.* Rand Corporation.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli. 2016. "Agricultural productivity and structural transformation: Evidence from Brazil." *American Economic Review* 106 (6): 1320-1365.

Danzer, Alexander M, and Lennard Zyska. forthcoming. "Pensions and Fertility: Micro-Economic Evidence." *American Economic Journal: Economic Policy.*

Funke, Manuel, Moritz Schularick, and Christoph Trebesch. 2022. "Populist leaders and the economy." *Working Paper.*

Huang, Wei, and Chuanchuan Zhang. 2021. "The Power of Social Pensions: Evidence from China's New Rural Pension Scheme." *American Economic Journal: Applied Economics* 13 (2): 179-205.

International Labour Organization. 2023. "Spotlight on Work Statistics n°12."

International Labour Organization. 2022. "World Social Protection Report 2020-22."

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review* 95 (1): 161-182.

Mason, Andrew, Ronald Lee, and NTANetwork. 2022. "Six Ways Population Change Will Affect the Global Economy." *Population and Development Review* 48 (1): 51-73.

Rossi, Pauline, and Mathilde Godard. 2022. "The Old-Age Security Motive for Fertility: Evidence from the Extension of Social Pensions in Namibia." *American Economic Journal: Economic Policy* 14 (4): 1-33.

Shen, Zheng, Xiadong Zheng, and Hualei Yang. 2020. "The fertility effects of public pension: Evidence from the new rural pension scheme in China." *PLoS ONE* 15 (6).

UNDESA, Population Division. 2022. "World Population Prospects 2022, Methodology Report."

6. Additional figures

Figure 6.1. Dropping badly matched cases, share of non-agricultural employment

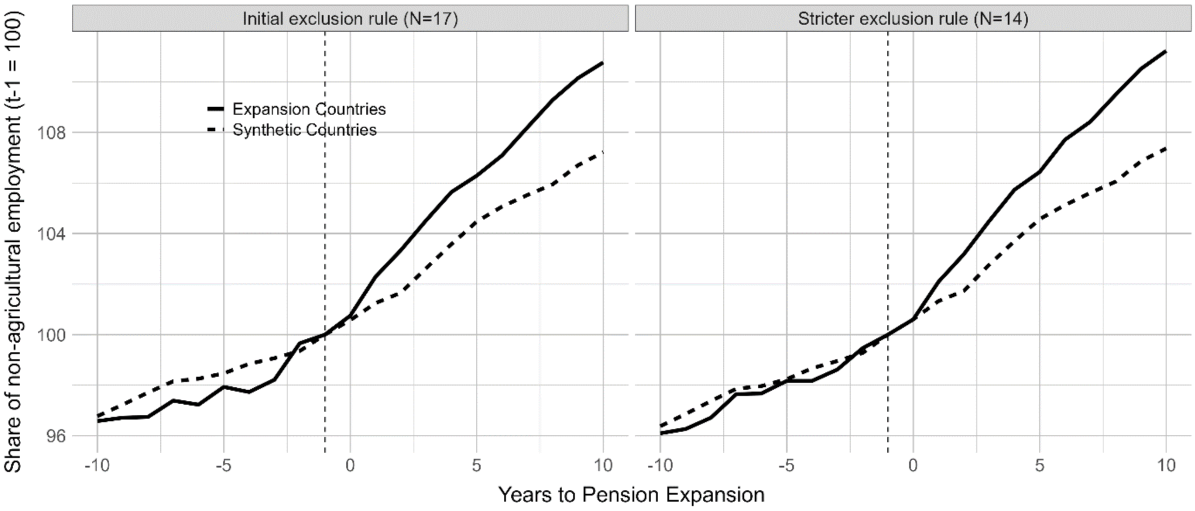

Figure 6.2. Time placebo, fertility

Figure 6.3. Time placebo, non-agricultural employment

Figure 6.4. Time placebo, GDP per capita