IARIW 2025

IARIW 2025

Thursday, October 2 & Friday, October 3

Constructing Comparable Global Poverty Trends

Daniel Gerszon Mahler
(World Bank)
Elizabeth Foster
(World Bank)
Christoph Lakner
(World Bank)
Zander Prinsloo
(World Bank)
Rostand Tchouakam Mbouendeu
(World Bank)
Samuel K. Tetteh-Baah
(World Bank)

dmahler@worldbank.org

Paper prepared for the IARIW–World Bank–UEB/VNU Conference on "Improving Well-being Measurement in Data-challenged Environments in Developing Countries for Better Evidence-based Policies" October 2-3, 2025

Session 3A: Poverty Measurement and Global Trends
Time Slot: Friday, October 3, 2025, 10:30 AM–12:30 PM

Global Poverty Monitoring Technical Note 45

Constructing Comparable Global Poverty Trends

Daniel Gerszon Mahler, Elizabeth Foster, Christoph Lakner, Zander Prinsloo, Rostand Tchouakam Mbouendeu, and Samuel K. Tetteh-Baah

June 2025

Keywords: Poverty; inequality; comparability; survey design; household surveys

Development Data Group
Development Research Group
Poverty and Equity Global Department

Abstract

Countries frequently revise how they measure income or consumption due to changes in data collection and questionnaire design. These changes create comparability breaks in poverty trends over time. This paper develops three methods to create global, regional, and country-level poverty trends that are comparable within countries over time. It does so by using national accounts growth to bridge non-comparable sequences. Accounting for comparability breaks creates large differences in some country-level poverty trends, but the global extreme poverty trend built from these comparable poverty series remains largely unchanged.

All authors are with the World Bank. Corresponding author: Daniel Gerszon Mahler (dmahler@worldbank.org). The authors are grateful for feedback received from the Global Poverty Monitoring Working Group. The authors gratefully acknowledge financial support from the UK government through the Data and Evidence for Tackling Extreme Poverty (DEEP) Research Programme.

1. Introduction

In low- and middle-income countries, monetary poverty is measured using data from household surveys that elicit the income or consumption patterns of a representative sample of households in a country. If a household's income or consumption falls short of a given poverty line, the household is classified as being poor. Numerous studies have shown that the design of household surveys impacts the data collected, measured income or consumption, and therefore also measured poverty rates (Beegle et al. 2012, De Weerdt et al 2020, Jolliffe 2001, Friedman et al. 2017, Kilic et al. 2019). For example, when measuring food consumption, it matters whether households are asked to keep a diary of their consumption or to recall their consumption. In the latter case, it is of consequence whether they are asked to recall their consumption over the past seven days or 30 days. For both food and non-food consumption, it matters how disaggregated the elicited categories of consumption are, for example whether households are asked to report their rice consumption in aggregate, or brown and white rice separately. Hence, systematic differences introduced by alternative survey instruments imply that welfare and poverty indicators produced under one design are not directly comparable with those generated under another.

Guidelines for data collection and questionnaire design have changed over time as more information about the accuracy and bias of various measurement approaches are uncovered and as consumption patterns in countries change (Deaton and Grosh 2000, Deaton and Zaidi 2002, Mancini and Vecchi 2022). To ensure that household surveys are on the frontier of poverty measurement, the underlying data collection and questionnaire design occasionally need to be revisited. Though this creates more reliable poverty estimates, it also results in comparability breaks in estimated poverty rates. At times, measured consumption increases by over 50% because of changes to the survey design and construction of consumption aggregates, which biases poverty trends using fixed poverty lines if unaccounted for (Mahler et al. 2024).

For single countries, solutions to this problem are well known and frequently applied, and include survey-to-survey imputations that predict the distribution of consumption had there been no change to the questionnaire design (Dang et al. 2017, Dang et al. 2019, Mathiassen and Wold 2021, Roy and Van der Weide 2025) or bridge surveys that collect an old consumption aggregate on a subsample in a newer survey. Yet the problem extends beyond single countries. To understand how poverty is evolving globally and by regions, such methods cannot be applied consistently. As a result, the problem is most often ignored in the measurement of global poverty.

In this paper, we develop three methods to address the incomparability of how consumption is measured within countries over time. The methods use growth in national accounts data to bridge comparability breaks and use varying degrees of information from surveys of older design. We apply the methods to the World Bank's global poverty numbers in the Poverty and Inequality Platform (PIP) to shed light on whether the trends in poverty are robust to differences in how consumption is measured. We find large changes for many countries, but relatively muted changes at the regional and global levels. When changes are apparent at the regional and global level, they often reflect corrections of trends that very likely were unrealistic.

The rest of the paper is organized as follows. Section 2 briefly outlines the data we use, section 3 details the method, section 4 contains the results and section 5 concludes.

2. Data

We use data from the June 2025 vintage of the World Bank's Poverty and Inequality Platform (PIP), which contains more than 2,500 estimates of poverty from 172 countries. These welfare vectors are expressed in 2021 purchasing power parity (PPP) dollars. When looking at poverty trends, we will primarily be using the extreme poverty line of $3.00 (2021 PPP) per day (Foster ⓡ al. 2025).

PIP contains information on whether consumption and income aggregates within a country are comparable to each other. Poverty estimates over time are assumed to be comparable unless there is a known and sizeable change to the data collection, survey methodology, or welfare aggregate construction. The assessment of comparability is country-dependent and relies on the accumulated knowledge from past and current World Bank staff, as well as close dialogue with national data producers. The comparability indicator is binary, meaning that either two surveys are comparable or not. In practice, comparability is never binary but should be understood as a spectrum from fully comparable to not comparable.

We also use data from national accounts, particularly series of GDP per capita or Household Final Consumption Expenditure (HFCE) per capita. Both series are taken from the World Development Indicators.

3. Method

We develop three approaches to generate comparable series within countries over time that differ by how much information from past surveys they utilize. Before elaborating on how they are constructed, it is useful to summarize how PIP currently extrapolates and interpolates between welfare aggregates. This is done for the purpose of creating a global poverty estimate, which requires an estimate of poverty for every country every year, even if no household survey was conducted in a given year.

Beyond the latest welfare aggregate, PIP *extrapolates* the welfare vector forward, assuming welfare grows with real GDP per capita or real HFCE per capita. GDP per capita is used for low-income and lower-middle income countries (LICs and LMICs) while HFCE per capita is used for upper-middle and high-income countries (UMICs and HICs). For consumption vectors, it discounts the growth by 30% such that only 70% pass through to the growth rates observed in household surveys, reflecting that the remainder likely is saved. This growth is applied in a distribution-neutral manner, meaning that inequality stays unchanged in the extrapolation period. A similar approach is applied before the first data point to extrapolate backwards. More details and a justification for these choices are presented in Mahler and Newhouse (2024).

PIP *interpolates* between welfare aggregates from two surveys regardless of whether they are comparable or not. There are two kinds of interpolations depending on whether growth in national accounts (real GDP per capita for LICs and LMICs, real HFCE per capita for UMICs and HICs) is of the same sign as growth in welfare. Intuitively, both methods gradually give larger weight to the latter survey data point and use growth rates from national accounts to influence the pace with which welfare changes. The exact method is described in World Bank (2025).

The three approaches we develop all utilize these extrapolation and interpolation rules in various ways. We call them approach A, B, and C, and they are increasingly using more information available in past surveys as we detail below and summarize in Table 1. Throughout, we will use the case of Viet Nam as an example to illustrate and visualize how the three approaches work. For now, we neglect the fact that some countries have both income and consumption aggregates, which complicates matters, and describe cases where only income or consumption aggregates are available. The case where both are available is discussed in section 3.1.

For all three approaches, our starting point is the latest comparable sequence of welfare aggregates for each country. For certain countries this may be a single welfare aggregate, if this aggregate is not comparable to any prior estimate for the country. The latest comparable sequence is information we certainly want to use, as it reflects the most recent and generally highest quality standard used to measure income or consumption in each country. This sequence, extrapolated to 2024 using PIP's default extrapolation method, is common to all our approaches and common to the current numbers in PIP. Using this information implies that the comparable series we derive are consistent with the series in PIP at the country-level, regional level, and global level in the most recent years.[1]

Table 1: Information used in three approaches to construct comparable series

Sequence	*Changes*	*A*	*B*	*C*	*PIP*
Latest comparable	All	Use	Use	Use	Use
Old comparable	All	Don't use	Use	Use	Use
Non-comparable	Distributional	Don't use	Don't use	Use	Use
	Mean	Don't use	Don't use	Don't use	Use

Note: *All* changes refer to both changes to the distribution and the mean.

Approach A does not make use of any information before the latest comparable sequence. Before the start of this latest comparable sequence, it uses PIP's extrapolation rule to grow (or more likely, shrink) welfare vectors backwards in time. This means that inequality is assumed to remain unchanged prior to the first survey in the latest comparable sequence. The assumption behind this

[1] There are a couple of exceptions to this for countries that have both income and consumption aggregates. As we will outline in section 3.1, for a few cases, we will change the latest estimate from consumption to income, or vice versa, which means that there is not a complete overlap in global and regional poverty rates in recent years. This concerns Albania, Kosovo, and the Philippines.

approach is that any information collected using welfare aggregates different from the latest comparable sequence reflects a different notion of welfare, which is a weak signal of how welfare is measured currently. A justification for keeping inequality constant could be that predicting changes in inequality is quite challenging (Mahler et al. 2022).

Concretely, consider a country that has an old comparable sequence between t_0 and t_1 and its most recent comparable sequence between t_2 and t_3 (and consequently a non-comparable sequence between t_1 and t_2). Let $welf_t$ be the distribution of welfare from the survey (or from PIP's extrapolation and interpolation if no survey estimate is available) at time t and g_{NA,t_i,t_j} be the passthrough-adjusted growth in national accounts between two time periods. Then, the welfare aggregate estimated by approach A, $\widehat{welf}_{A,t_i}$, equals $welf_{t_i}$ for $i \geq 2$ and $\frac{welf_{t_2}}{1+g_{NA,t_i,t_2}}$ for $i < 2$.

Viet Nam provides an interesting case of how this approach works. Viet Nam is particularly illustrative due to the significant differences between PIP's estimates and the comparable series we derive. Viet Nam has the latest comparable consumption data for the period 2010-2022, meaning that Approach A gives the same mean, inequality and poverty rates as PIP from 2010 onwards (Figure 1). However, between 2008 and 2010 there was a change in how consumption was measured. The 2010-2022 series includes imputed rent and durables (as well as other changes) while the series ending in 2008 does not, leading to a mechanical increase in measured welfare from 2010 onwards (World Bank 2022). If ignored, this leads to a decline in poverty from 2008-2010 from 17.6% to 5.2%. Approach A ignores all survey data prior to 2010 and instead uses national accounts growth to project the welfare vector backwards. The national accounts backwards-extrapolated poverty rate for 2008 is 6.4%, much lower than the survey estimate, reflecting that part of the decline in poverty from 2008-2010 likely was caused by the break in comparability.

Approach B differs from approach A by utilizing growth and distributional changes from old comparable welfare vectors. These older welfare vectors are likely of worse quality, yet this approach holds that growth and distributional changes from older welfare vectors on expectation better capture what would have been measured with the most recent methodology than using PIP's extrapolation rule (i.e. better than assuming no distributional changes and growth from national accounts). To bridge non-comparable sequences, it uses PIP's extrapolation rule like Approach A.

Continuing the example from before, Approach B would be identical to Approach A from t_1 onwards, that is $\widehat{welf}_{B,t_i} = \widehat{welf}_{A,t_i}$ for $i \geq 1$. However, prior to t_1, it would utilize the survey-based growth incidence between t_0 and t_1. In contrast to Approach A, this means that Approach B allows for distributional changes prior to the latest comparable sequence. Concretely, denote the anonymized growth of the pth percentile between these two surveys as g_{welf,p,t_0,t_1}. Then, the welfare of the pth percentile in t_0 will be given by $\widehat{welf}_{B,p,t_0} = \frac{\widehat{welf}_{B,p,t_1}}{1+g_{welf,p,t_0,t_1,}}$.

Figure 1: Example of adjustments to Viet Nam

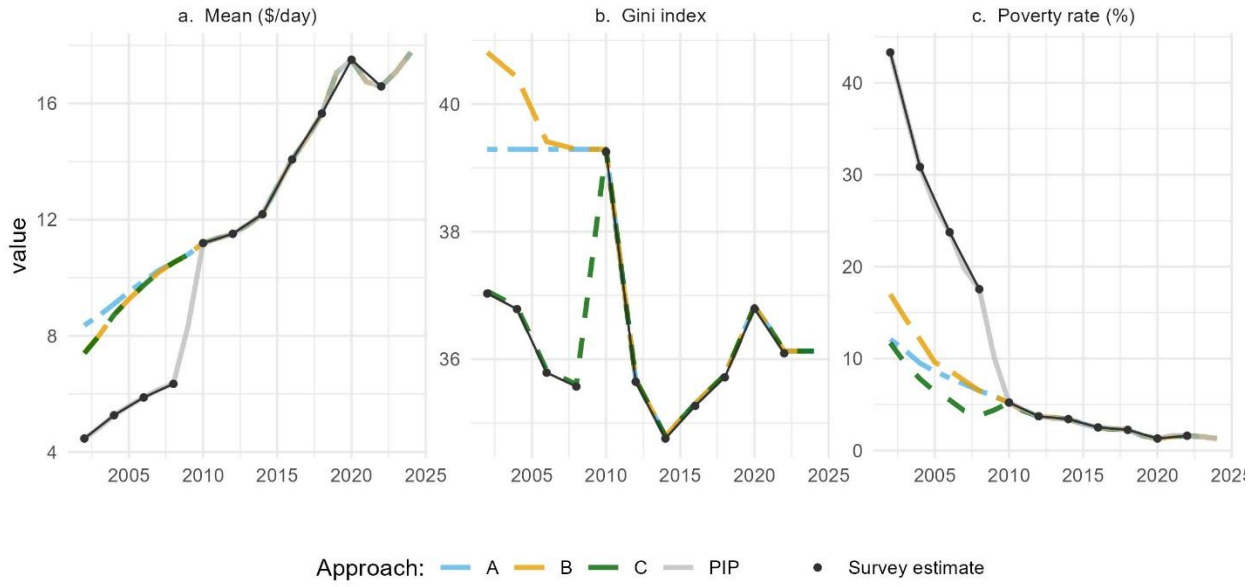

Note: Non-connected black lines indicate non-comparable survey estimates. The Gini index is available in PIP only for survey years.

In Viet Nam, Approach A and B are similar from 2008-10 where the trend is computed from national accounts growth while holding inequality constant. Before 2008 they diverge as comparable welfare data is available from 2002-2008, which Approach B uses. The growth in welfare suggested by the data from 2002-2008 is greater than the passthrough-adjusted growth from national accounts while the survey data suggests that inequality fell from 2002-2008 rather than remaining unchanged. This implies that Approach B suggests higher poverty in 2002 than Approach A (which has lower inequality and a higher mean in 2002).

Approach C is identical to approach B but utilizes one more aspect of past surveys. It assumes that distributional changes between *non-comparable sequences*, though in part reflecting methodological changes in data collection, welfare aggregate construction etc., are a better signal of changes to inequality than assuming no changes to inequality. When the non-comparable sequence is short and displays large changes to inequality – such as the Vietnamese case from 2008-2010 – this is unlikely to be a sound assumption. Yet in other cases, particularly for countries that do not have any comparable sequences and have many years between survey estimates – often the case in poorer countries – inequality changes between two non-comparable surveys might capture more signal than noise.

With the hypothetical example from before, Approach C would be identical to Approach A and B from t_2 onwards. For t_1, it uses the survey distribution at the time but rescales the distribution to match the mean with what growth from national accounts would have suggested. That is,

$\widehat{welf}_{C,t_1} = welf_{t_1} * \frac{1+g_{welf,t_1,t_2}}{1+g_{NA,t_1,t_2}}$. Before t_1, it uses the distributional incidence of the survey growth, like Approach B, that is $\widehat{welf}_{C,p,t_0} = \frac{\widehat{welf}_{C,p,t_1}}{1+g_{welf,p,t_0,t_1}}$.

In Viet Nam, the Gini increased from 35.6 to 39.3 between 2008 and 2010. Approach C holds that this information is relevant and hence forces the distribution in 2008 to have a Gini of 35.6. This results in lower poverty in 2008 when compared to approach B, which has a higher Gini in 2008.

3.1. Switches between income and consumption

Some countries switch between income and consumption. PIP generally prefers consumption estimates to income estimates and selects the former when both are available for a given year. If only income estimates are available for a particular year, it is used. If neither income nor consumption estimates are available in a particular year, PIP interpolates between two consumption aggregates if such exist, and between two income aggregates otherwise (if such exist). If only income aggregates are available after and only consumption aggregates before (or vice versa), PIP interpolates between two aggregates of different type.

The approach above, next to ignoring that income and consumption aggregates are highly non-comparable, often causes annual switches from income to consumption aggregates. This happens, for example, for Serbia as shown in Figure 2, where PIP's poverty estimates change between income and consumption aggregates seven times between 2010 and 2020. There is nothing wrong with the separate income and consumption series, the challenge only arises when deciding on a preferred annual estimate from the two series.

Figure 2: Example of adjustments to Serbia

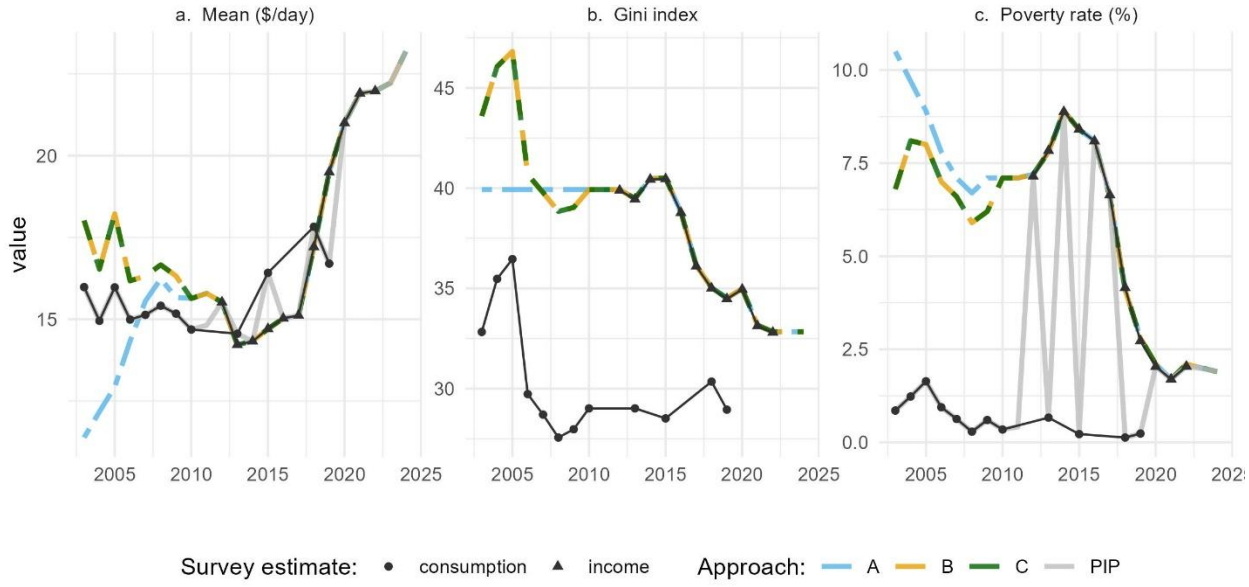

To create a comparable series for a country over time, unnecessary jumps between income and consumption should be avoided. As an alternative to PIP's current procedure, we select between income and consumption estimates based on which are deemed preferable by World Bank Poverty Economists – World Bank staff experts on poverty and inequality in a particular country. This often aligns with the metric used by countries for their own poverty monitoring. In Serbia, the income series is preferred throughout. This prevents a jump between income and consumption for countries that have both in a similar time period but still leaves occasional cases of switching when a country permanently changes which metric to use. For example, Serbia's first income estimate is from 2012 while there are many consumption estimates prior to that. Our three approaches deal with such past estimates differently.

Whenever a country switches between income and consumption estimates, we do not use the welfare growth and distributional changes associated with such a switch in any of the three approaches. Rather, we use the extrapolation rule to go from the end of the income (consumption) sequence to the start of the consumption (income) sequence. Not even approach C, which utilizes inequality changes between non-comparable welfare aggregates of the same type, uses this information. The reason is that there are systematic differences between income and consumption inequality and hence changes in distribution arising from such switches are highly unlikely to signal actual distributional changes.

For countries that have comparable sequences of the welfare type no longer used (for example, Serbia's consumption sequence from 2003-2010), Approach A ignores this information, while B and C utilize it. This mimics how old comparable sequences of the same welfare type were handled. The only difference is that Approach B and C discounts or inflates g_{welf,p,t_i,t_j} to reflect that income growth generally is larger (in absolute terms) than consumption growth. For countries that currently use consumption, but which have old comparable income sequences, the percentile-specific income growth, $g_{inc,p,t_i,t_{j,}}$ is discounted 30%. Conversely, countries that currently use income but have old comparable consumption sequences, such as Serbia, the percentile-specific consumption growth, $g_{con,p,t_i,t_{j,}}$, is inflated by 30%. This reflects the fact that only about 70% of growth in GDP is estimated to pass through to consumption vectors while 100% of growth in GDP is estimated to pass through to income vectors.

The growth and inequality changes from non-comparable sequences of the old welfare type are not used by approach A and B, but the distributional changes are used by approach C, in line with how non-comparable sequences of the same welfare type are used. A summary of the information used is given in Table 2.

We apply the three methods to all countries with survey data in PIP and generate three alternative annual distributions for all countries and all years. We next calculate poverty rates at various lines from these distributions, and aggregate these poverty rates to regional and global poverty rates using population-weighted averages. In line with PIP, for the purpose of calculating global poverty rates, we assign countries without any surveys the regional average poverty rate.

Table 2: Information used when switching between income and consumption

Sequence	*Changes*	*A*	*B*	*C*	*PIP*
Old comparable of old welfare type	All	Don't use	Use fraction	Use fraction	Use
Non-comparable of old welfare type	Distribution	Don't use	Don't use	Use	Use
	Mean	Don't use	Don't use	Don't use	Use
Switches between welfare types	All	Don't use	Don't use	Don't use	Use

4. Results

Figure 3 shows the global extreme poverty rates currently in PIP and the three series using the more comparable series we have recovered. The series are mechanically nearly identical in 2024 and differ only marginally from 2013-2024. In 2013, PIP's series declines 2.4 pct. points while Approach A and B decline only by 1.1 and 1.0 pct. points. This coincides with a comparability break in China that added included rent and made more changes to its measured consumption, which increased measured consumption and hence lowered poverty. It has been estimated that this change in measured consumption lowered China's extreme poverty rate by 31 million people (World Bank 2016). Our more comparable series attempt to mitigate this break.

Figure 3: Global extreme poverty rates, 1981-2024

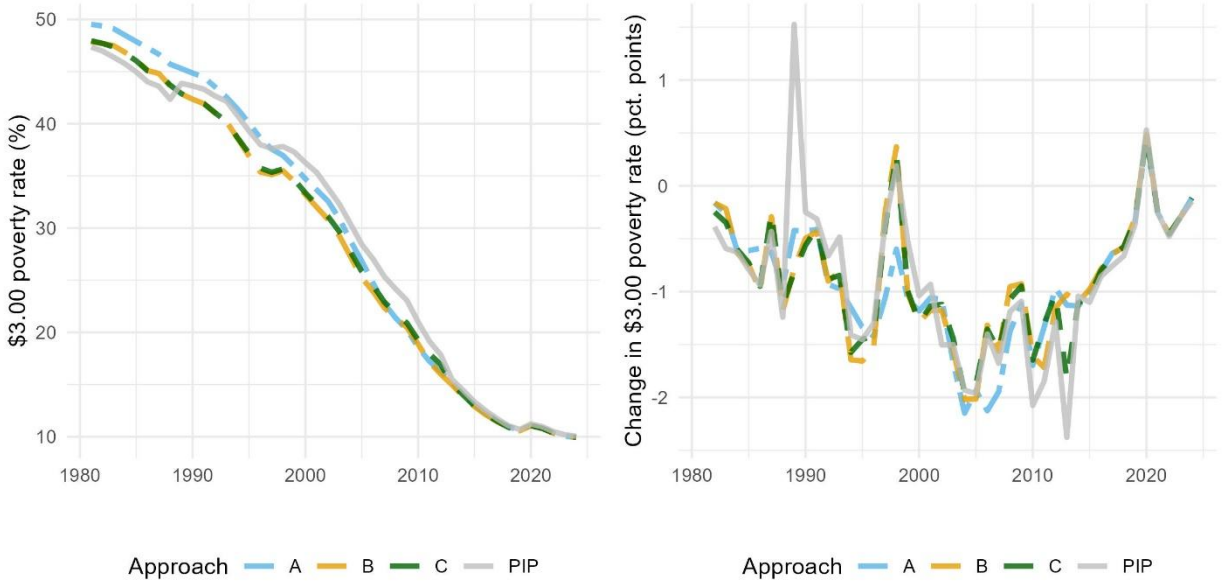

In the 1990s and 2000s, the trends are relatively parallel but systematically a little lower with Approach B and C than in PIP, largely reflecting the break in the comparability in China in 2013.

The next large change in trends happens from 1988 to 1989, where PIP has an increase in poverty of 1.5 pct. points, while the three comparable series decline in the range of 0.4 pct. points to 0.8 pct. points. This again reflects an incomparability in China, which went from an income to a consumption aggregate from 1987 to 1990. These two breaks together mean that the overall decline in poverty from 1981 to 2024 is broadly similar with approach B, C, and PIP. In general, systematic effects across countries are almost fully overwhelmed by sporadic changes in some of the high population countries. Approach A has a slightly larger decline in poverty.

This masks larger differences at the regional level (Figure 4). The divergences in the trend for East Asia & Pacific largely reflect the comparability breaks for China. In Europe & Central Asia, the trends are rather aligned except for a large difference in 2012 when Uzbekistan switches from a consumption to income aggregate and the comparable series suggests a lower poverty rate. The trends converge in the 1990s due to opposing impacts in Turkey, Romania, Kazakhstan, and Russia. In South Asia, Approach B and C align fairly well with the series in PIP in recent decades. This is in large part because a large comparability break in India when it switched from a uniform to a mixed modified reference period is solved at the country-level in PIP, and hence accounted for in all series. Still, Approach A suggests a much more rapid decline, reflecting that (passthrough-adjusted) growth in GDP in India was faster than growth in consumption from household surveys. In Sub-Saharan Africa, the trend from 2014-2019 differs noticeably, with PIP suggesting a small decline in poverty but the three comparable series suggesting a small increase. Together with the COVID-induced increase in poverty in 2020 and the stagnation in the years that followed, poverty is higher in Sub-Saharan Africa in 2024 than in 2014 with the comparable series.

Figure 4: Regional extreme poverty rates, 1981-2024

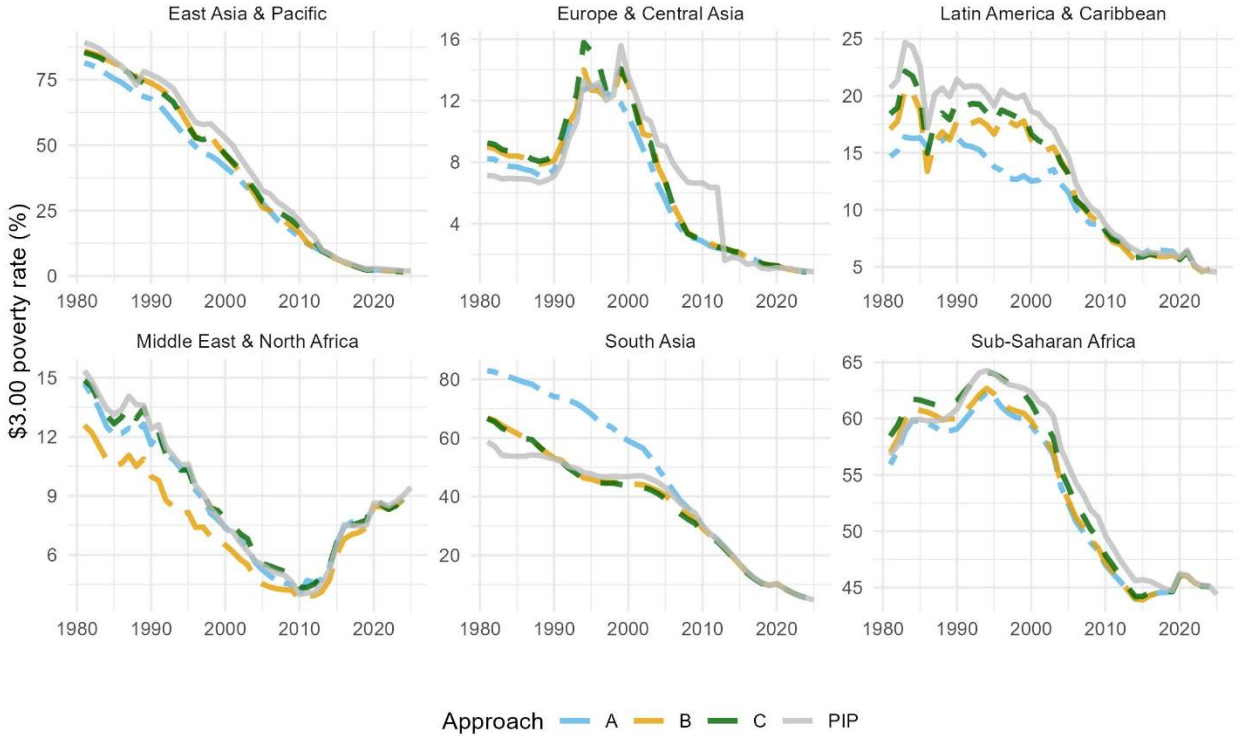

Figures A.1-A.15 in the Annex has similar figures for all 172 economies with survey data in PIP.

5. Conclusion

This note has suggested approaches to bridge non-comparable poverty trends using growth rates from national accounts. The approaches can be readily applied across countries given that national accounts growth rates are nearly universally available. National accounts growth has been shown to perform relatively well as predictors of welfare, when compared to more complex methods (Mahler et al. 2022).

For future work, it would be interesting to use metadata about consumption aggregates to account for comparability breaks. One could, for example, predict the impact on welfare distributions of changing the recall period, including durable goods, or other common changes when countries update their household surveys, and through such a prediction adjust older welfare distributions. Currently, such data does not exist with sufficient completeness. It would also be interesting for future work to deal with comparability issues between countries, such as differences in how the latest consumption aggregates are constructed and whether income or consumption is used.

References

Beegle, Kathleen, Joachim De Weerdt, Jed Friedman, and John Gibson. 2012. "Methods of household consumption measurement through surveys: Experimental results from Tanzania." *Journal of Development Economics* 98 (1): 3-18.

Dang, Hai-Anh H., Peter F. Lanjouw, and Umar Serajuddin. 2017. "Updating poverty estimates in the absence of regular and comparable consumption data: methods and illustration with reference to a middle-income country." *Oxford Economic Papers* 69 (4): 939-962.

Dang, Hai-Anh, Dean Jolliffe, and Calogero Carletto. 2019. "Data gaps, data incomparability, and data imputation: A review of poverty measurement methods for data-scarce environments." *Journal of Economic Surveys* 33 (3): 757-797.

Deaton, Angus, and Margaret Grosh. 2000. "Consumption.' In Grosh and Glewee (eds.) *Designing household survey questionnaires for developing countries lessons from ten years of LSMS experience*". Chapter 17: 91-133

Deaton, Angus, and Salman Zaidi. 2002. *Guidelines for constructing consumption aggregates for welfare analysis*. Living Standards Measurement Study Working Paper 135, World Bank, Washington DC.

De Weerdt, Joachim, John Gibson, Kathleen Beegle. 2020. "What can we learn from experimenting with survey methods?" *Annual Review of Resource Economics 12*(1): 431-447.

Foster, Elizabeth ⓡ Dean Jolliffe ⓡ Gabriel Lara Ibarra ⓡ Christoph Lakner ⓡ, and Samuel Kofi Tetteh-Baah. 2025. "Global Poverty Revisited Using 2021 PPPs and New Data on Consumption." World Bank Policy Research Working Paper Series. World Bank, Washington, DC.

Gibson, John, Kathleen Beegle, Joachim De Weerdt, and Jed Friedman. 2015. "What does variation in survey design reveal about the nature of measurement errors in household consumption?" *Oxford Bulletin of Economics and Statistics* 77 (3): 466-474.

Jolliffe, Dean. 2001.m "Measuring absolute and relative poverty: the sensitivity of estimated household consumption to survey design." *Journal of Economic and Social Measurement* 27 (1-2): 1-23.

Friedman, Jed, Kathleen Beegle, Joachim De Weerdt, and John Gibson. 2017. "Decomposing response error in food consumption measurement: Implications for survey design from a randomized survey experiment in Tanzania." *Food Policy* 7): 94-111.

Kilic, Talip, and Thomas Pave Sohnesen. 2019. "Same question but different answer: experimental evidence on questionnaire design's impact on poverty measured by proxies." *Review of Income and Wealth* 65 (1): 144-165.

Mahler, Daniel Gerszon, Raul Andres Castañeda Aguilar, and David Newhouse. 2022. "Nowcasting global poverty." *The World Bank Economic Review* 36(4): 835-856.

Mahler, Daniel Gerszon, and David Newhouse. 2024. "*Changes to the Extrapolation Method for Global Poverty Estimation*." Global Poverty Monitoring Technical Note 35. World Bank, Washington DC.

Mahler, Daniel Gerszon, Elizabeth Foster, Samuel Kofi Tetteh-Baah. 2024. "How Improved Household Surveys Influence National and International Poverty Rates." Global Poverty Monitoring Technical Note 40. Washington, D.C.: World Bank.

Mancini, Giulia, and Giovanni Vecchi. 2022. "On the construction of a consumption aggregate for inequality and poverty analysis." Washington, D.C.: World Bank.

Mathiassen, Astrid, and Bjørn K. Getz Wold. 2021. "Predicting poverty trends by survey-to-survey imputation: the challenge of comparability." *Oxford Economic Papers* 73 (3): 1153-1174.

Roy, Sutirtha Sinha, and Roy Van Der Weide. 2025. "Estimating poverty for India after 2011 using private-sector survey data." *Journal of Development Economics* 172: 103386.

World Bank. 2016. "Poverty and Shared Prosperity 2016: Taking on Inequality." Washington,
D.C.: World Bank.

World Bank. 2022. "From the Last Mile to the Next Mile: 2022 Vietnam Poverty and Equity Assessment." Washington, D.C.: World Bank.

World Bank 2025. Poverty and Inequality Platform Methodology Handbook. Edition 2025-06. Available at https://datanalytics.worldbank.org/PIP-Methodology/.

Appendix

Figure A.1: Impact on country-level poverty trends (1/15)

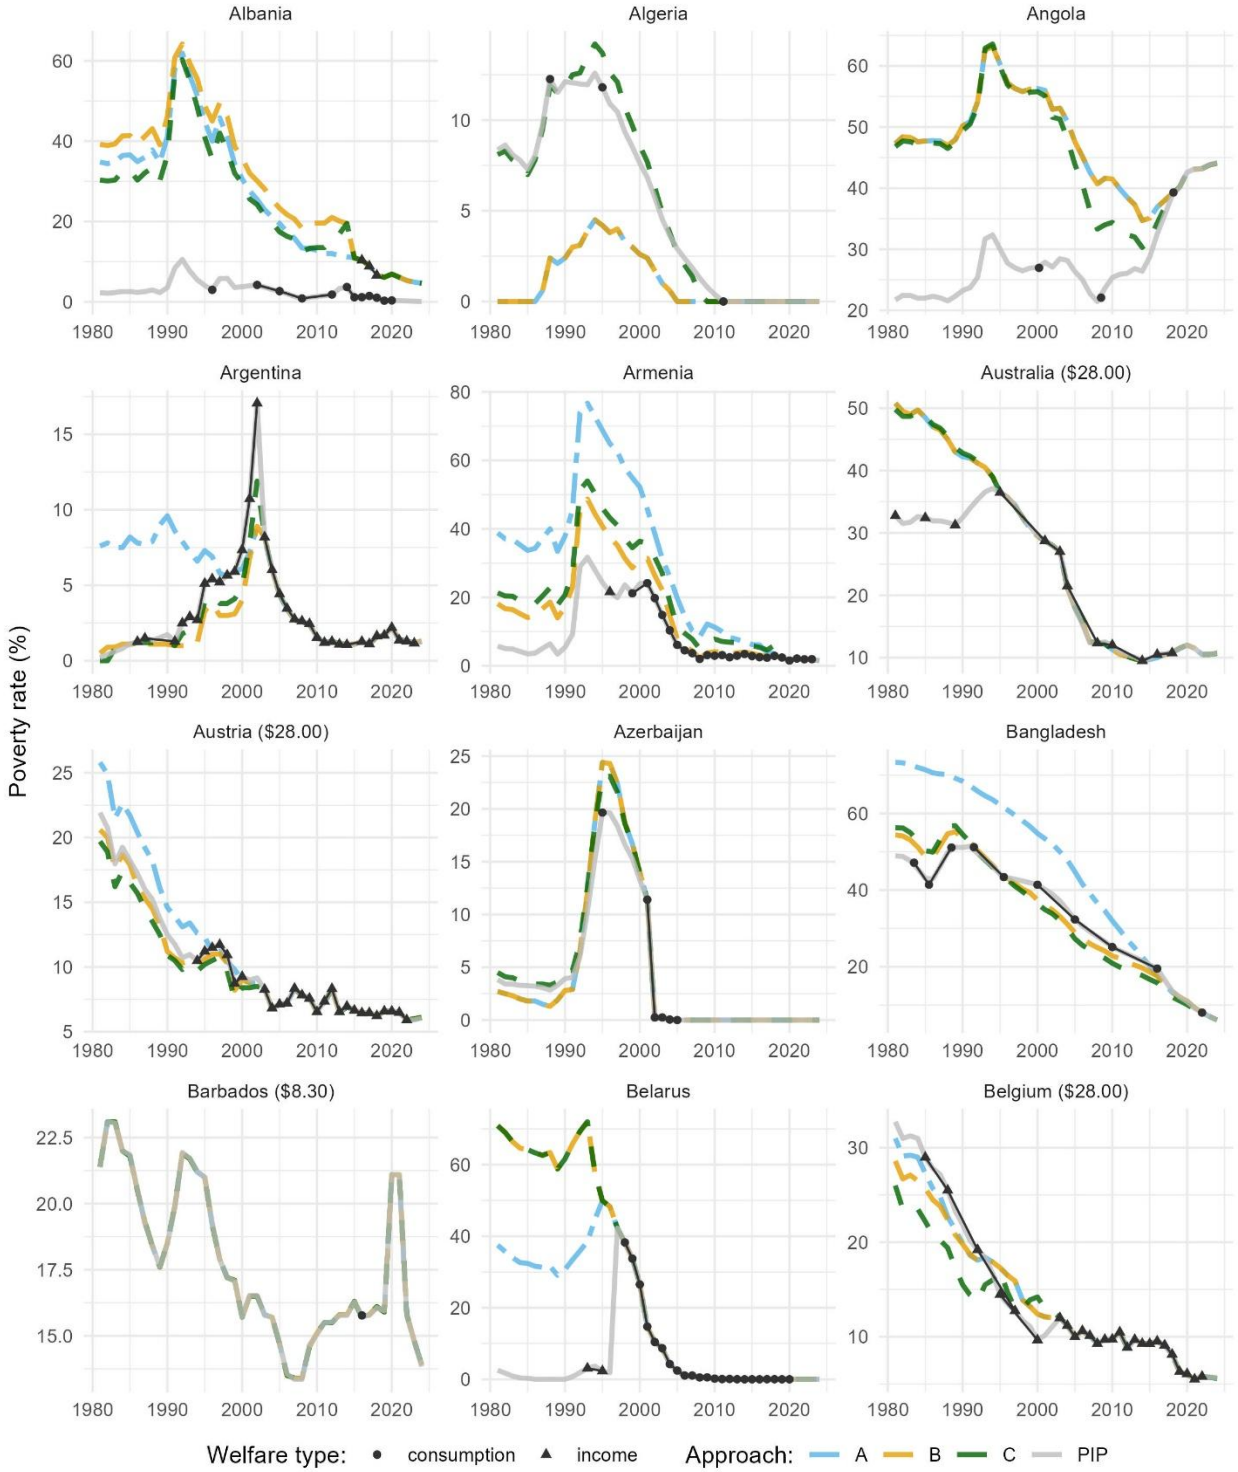

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.2: Impact on country-level poverty trends (2/15)

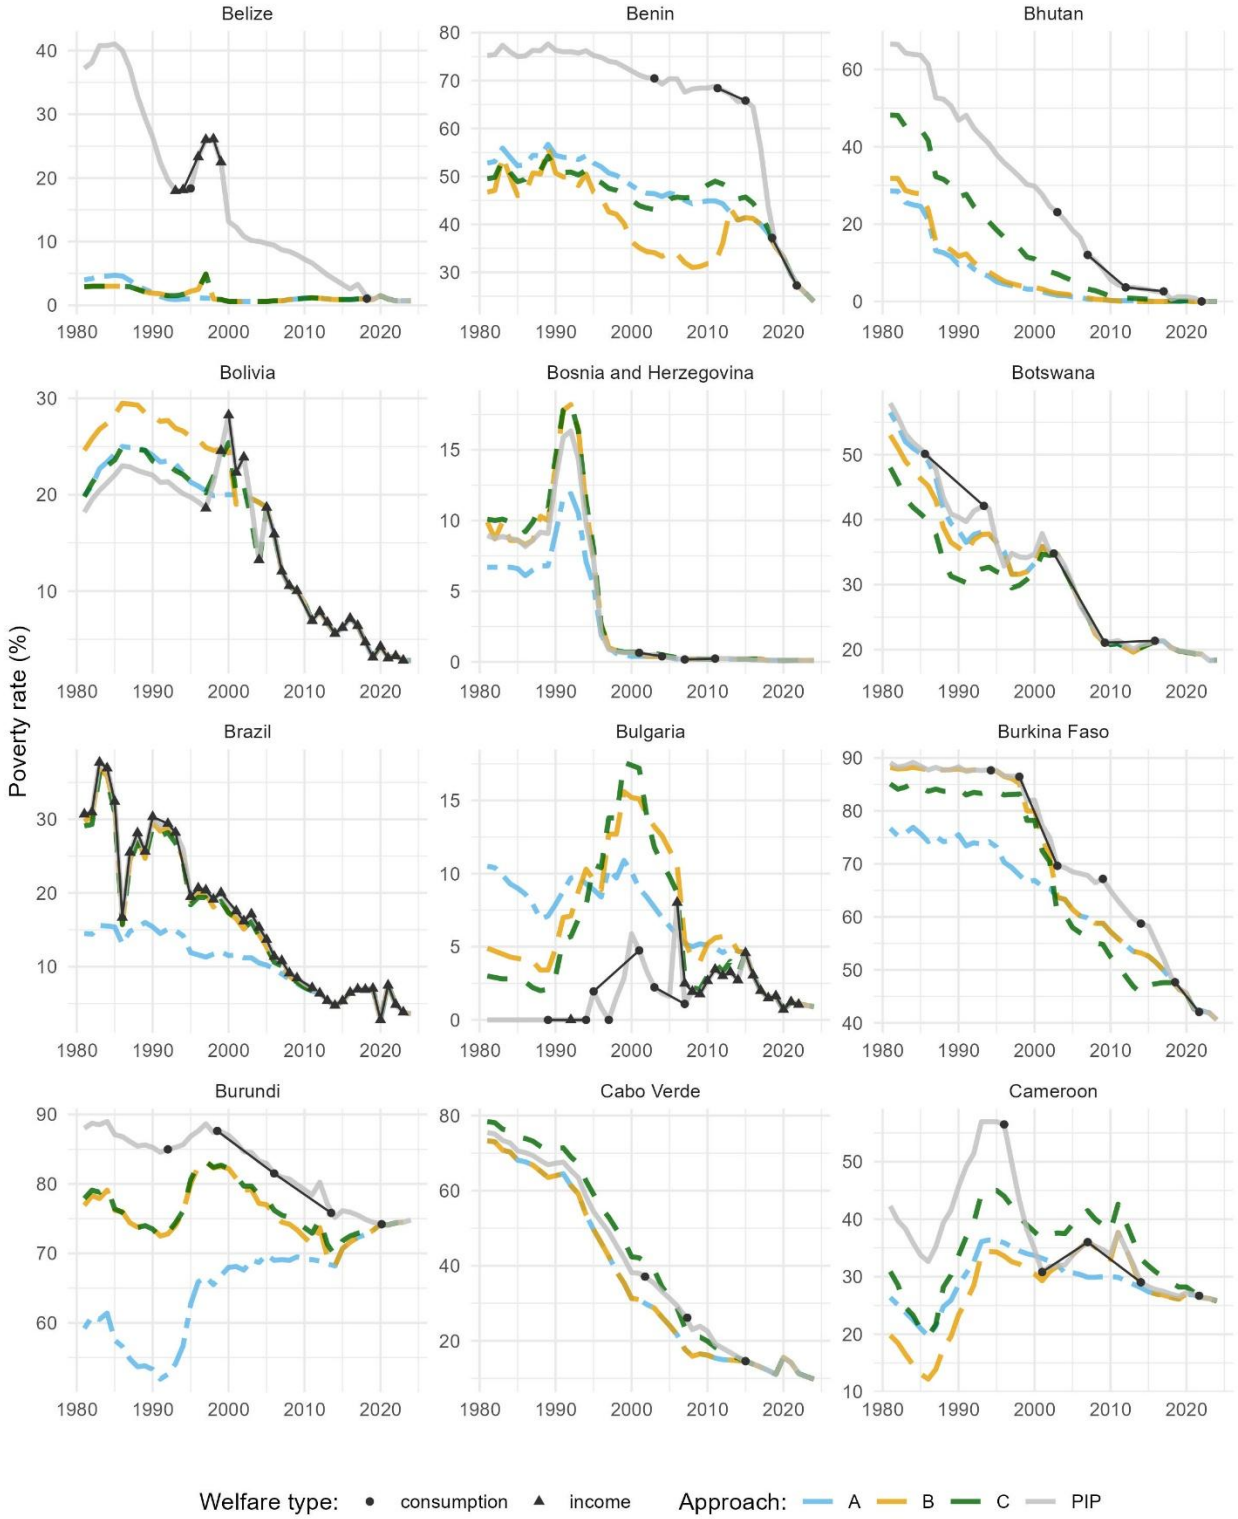

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.3: Impact on country-level poverty trends (3/15)

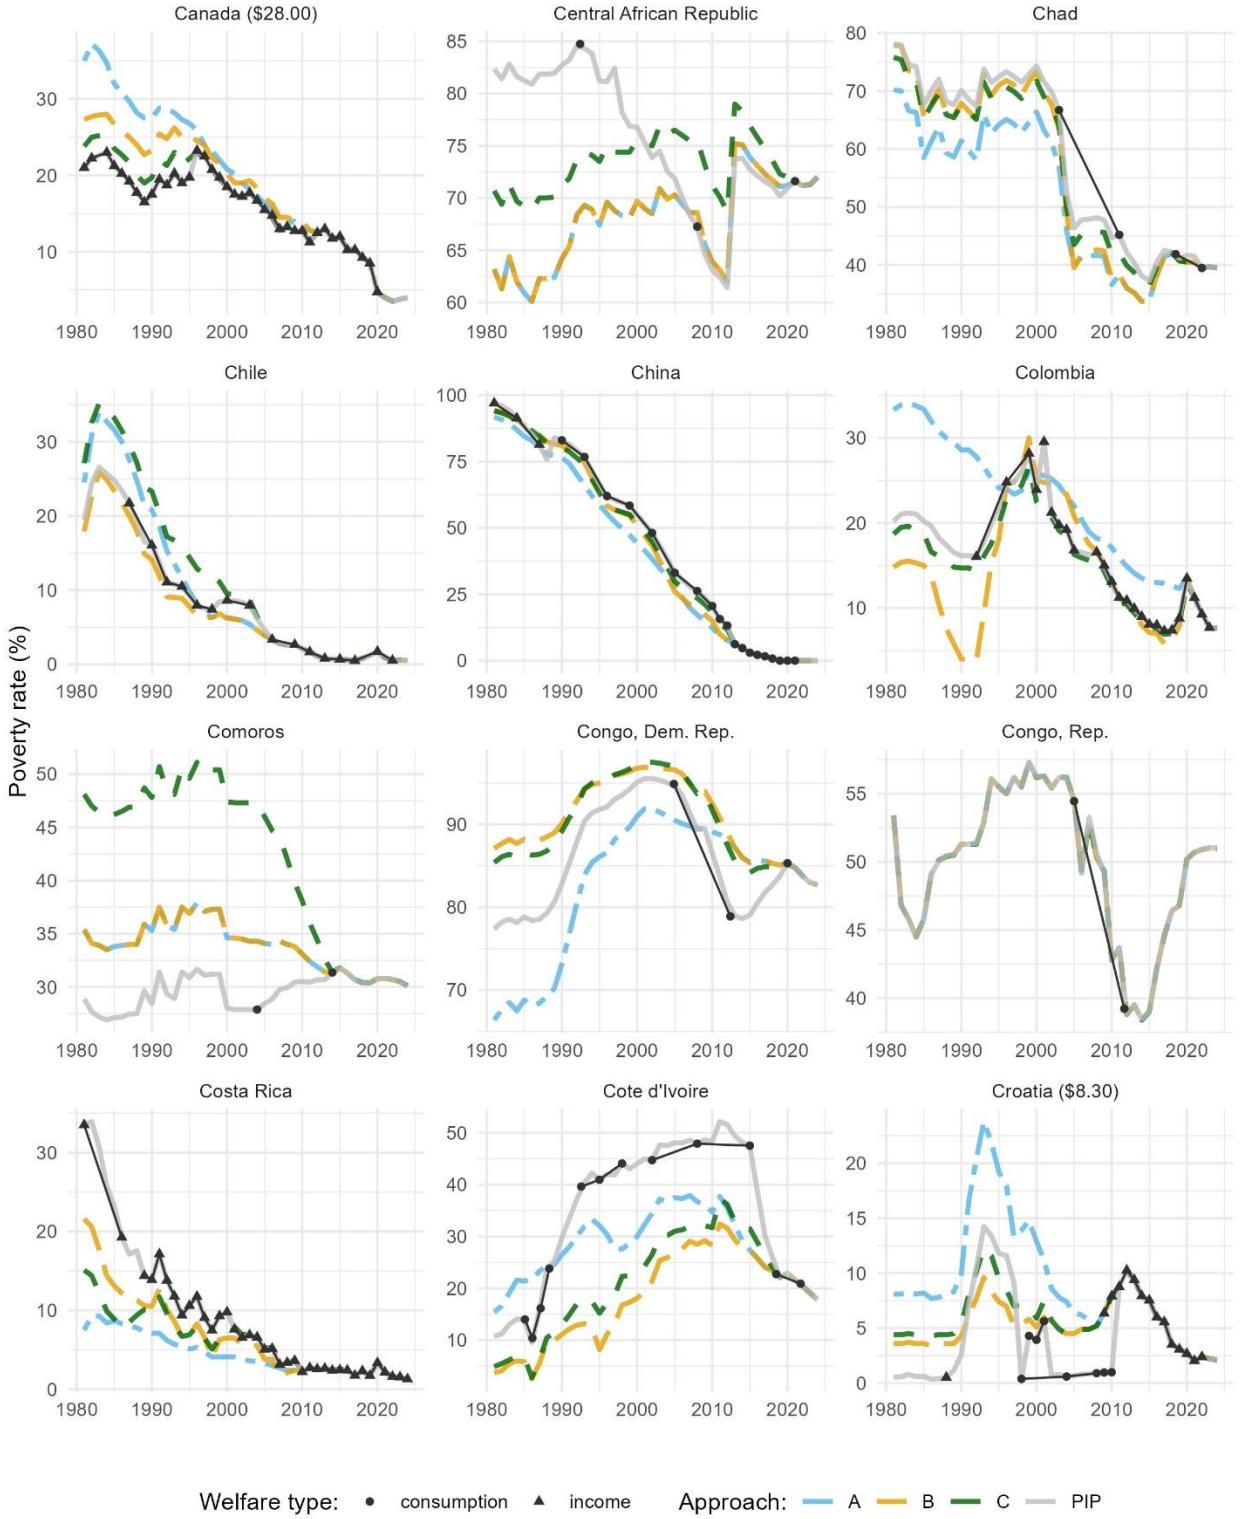

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.4: Impact on country-level poverty trends (4/15)

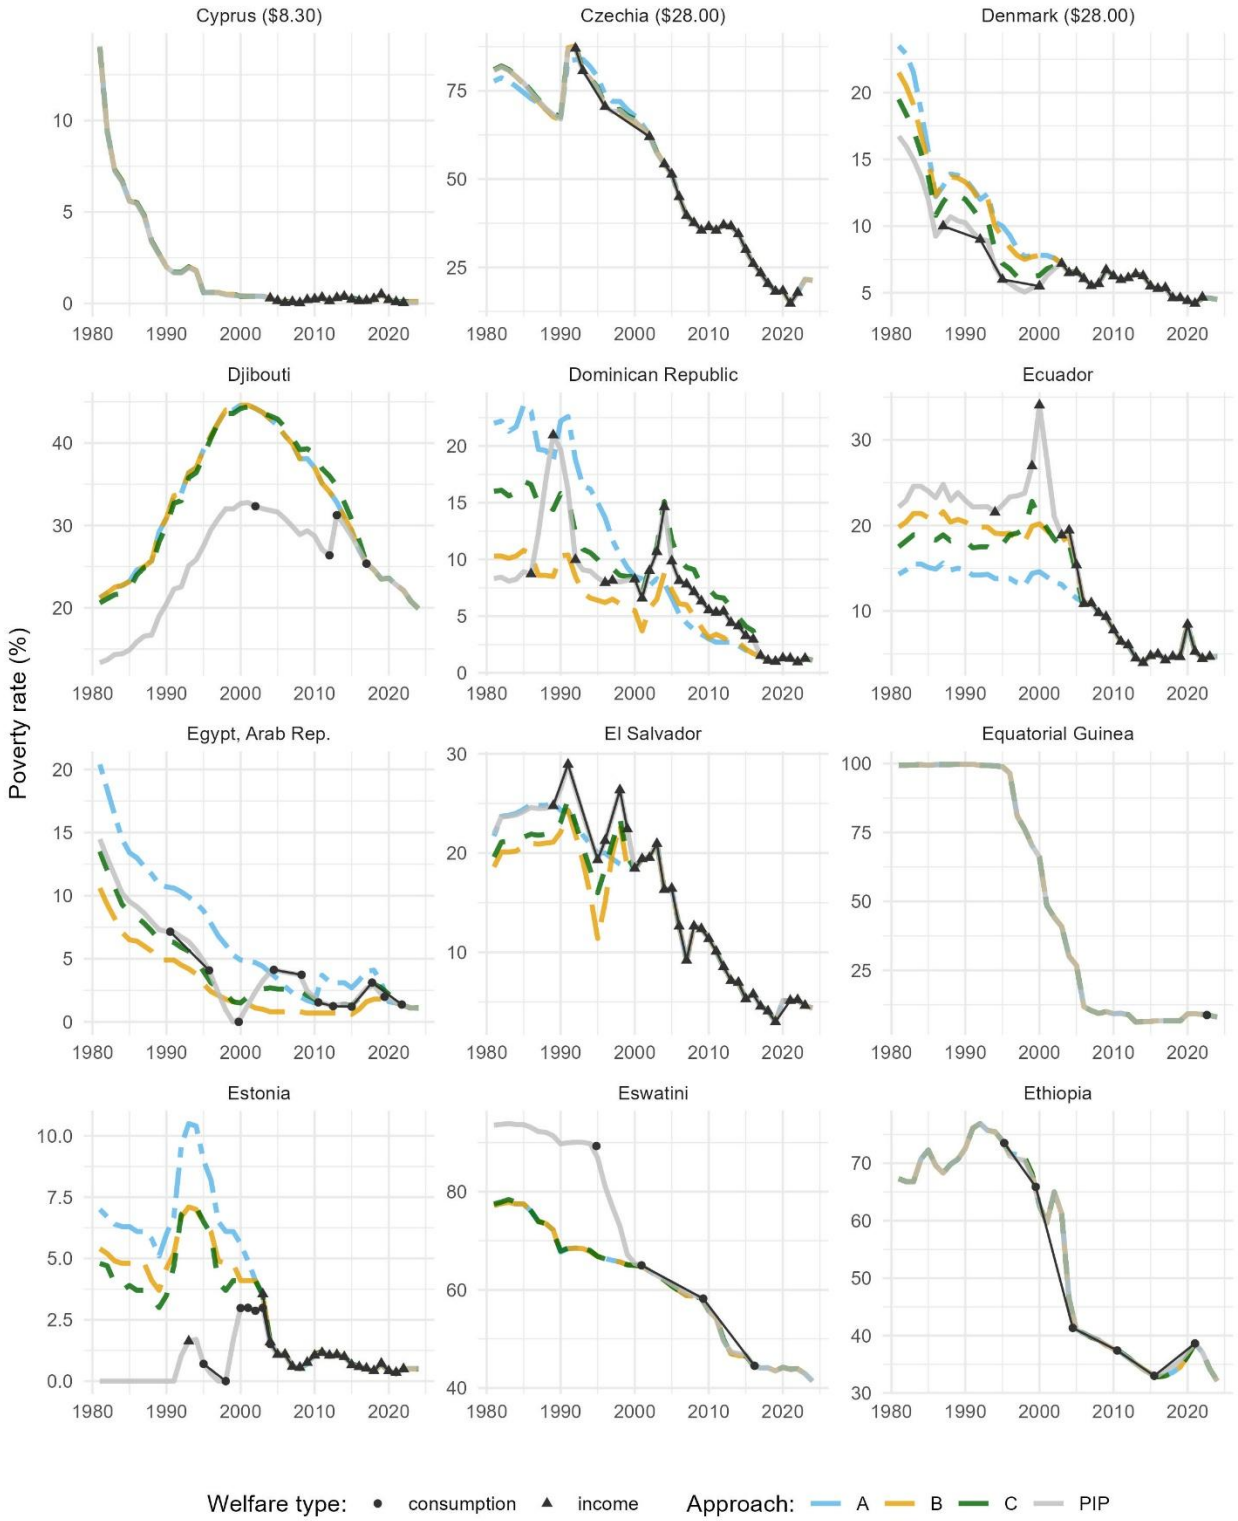

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.5: Impact on country-level poverty trends (5/15)

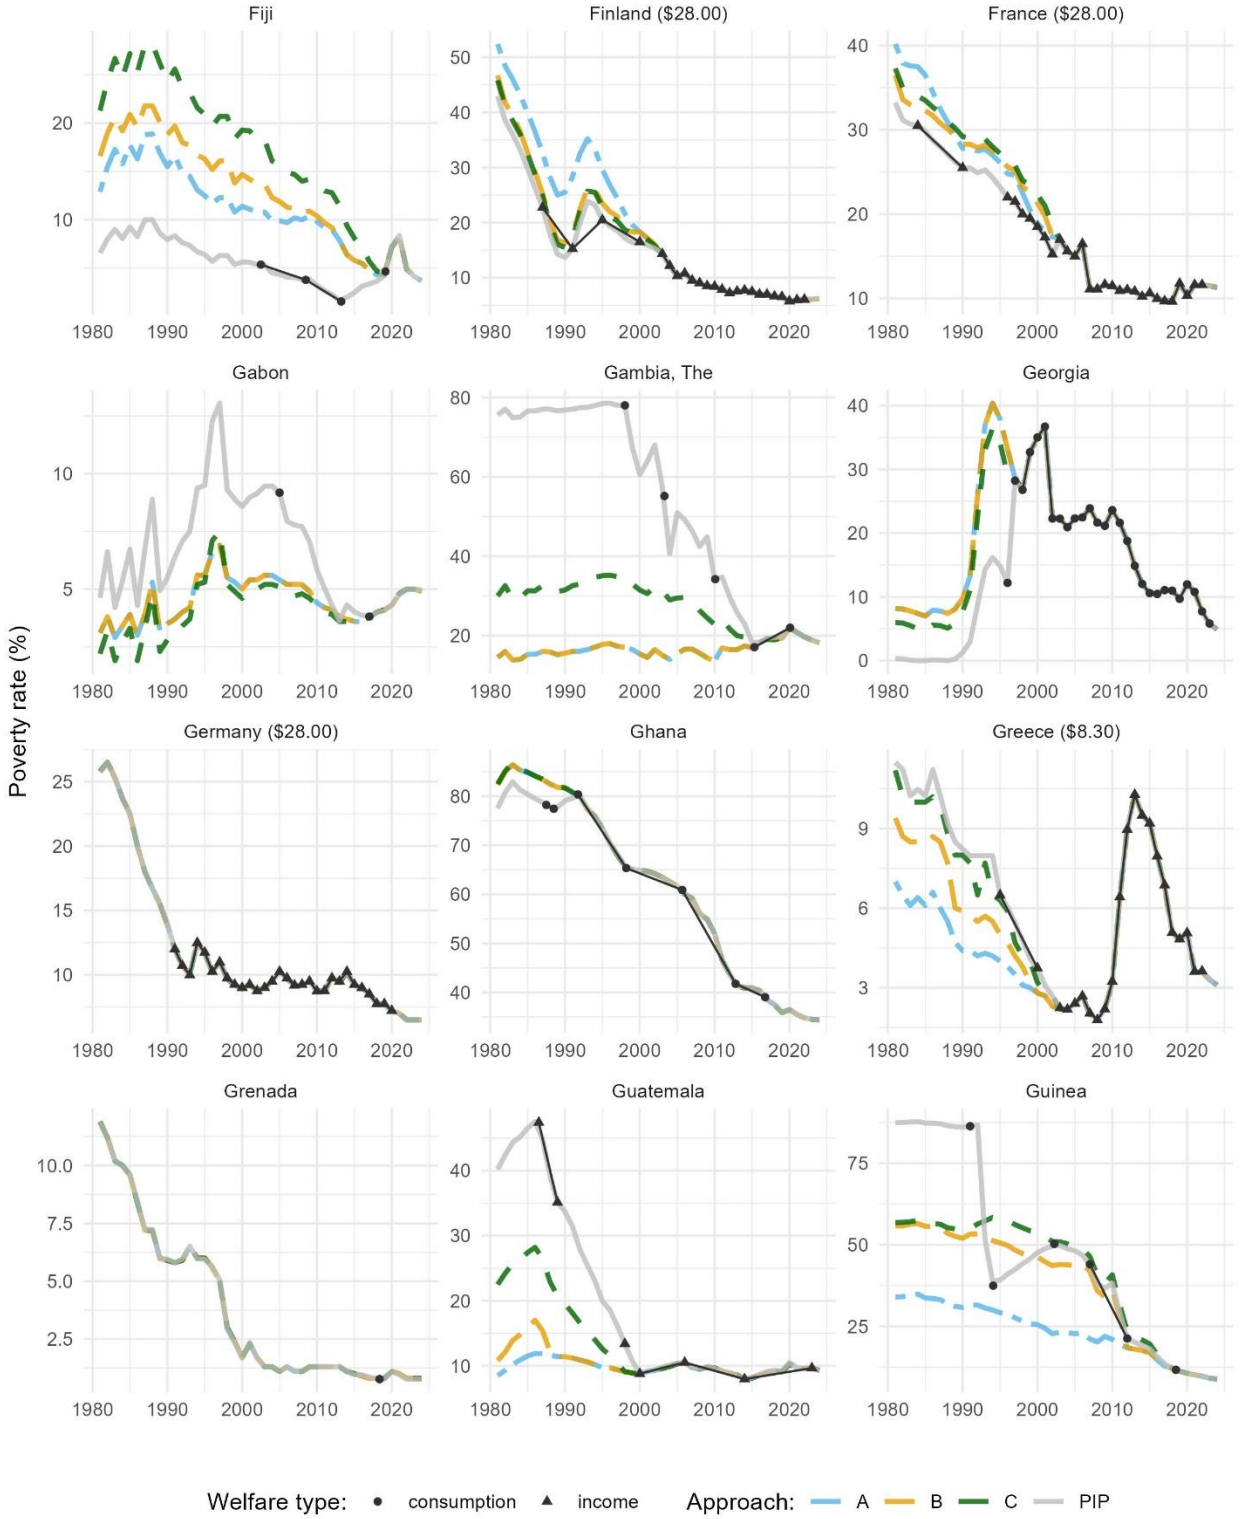

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.6: Impact on country-level poverty trends (6/15)

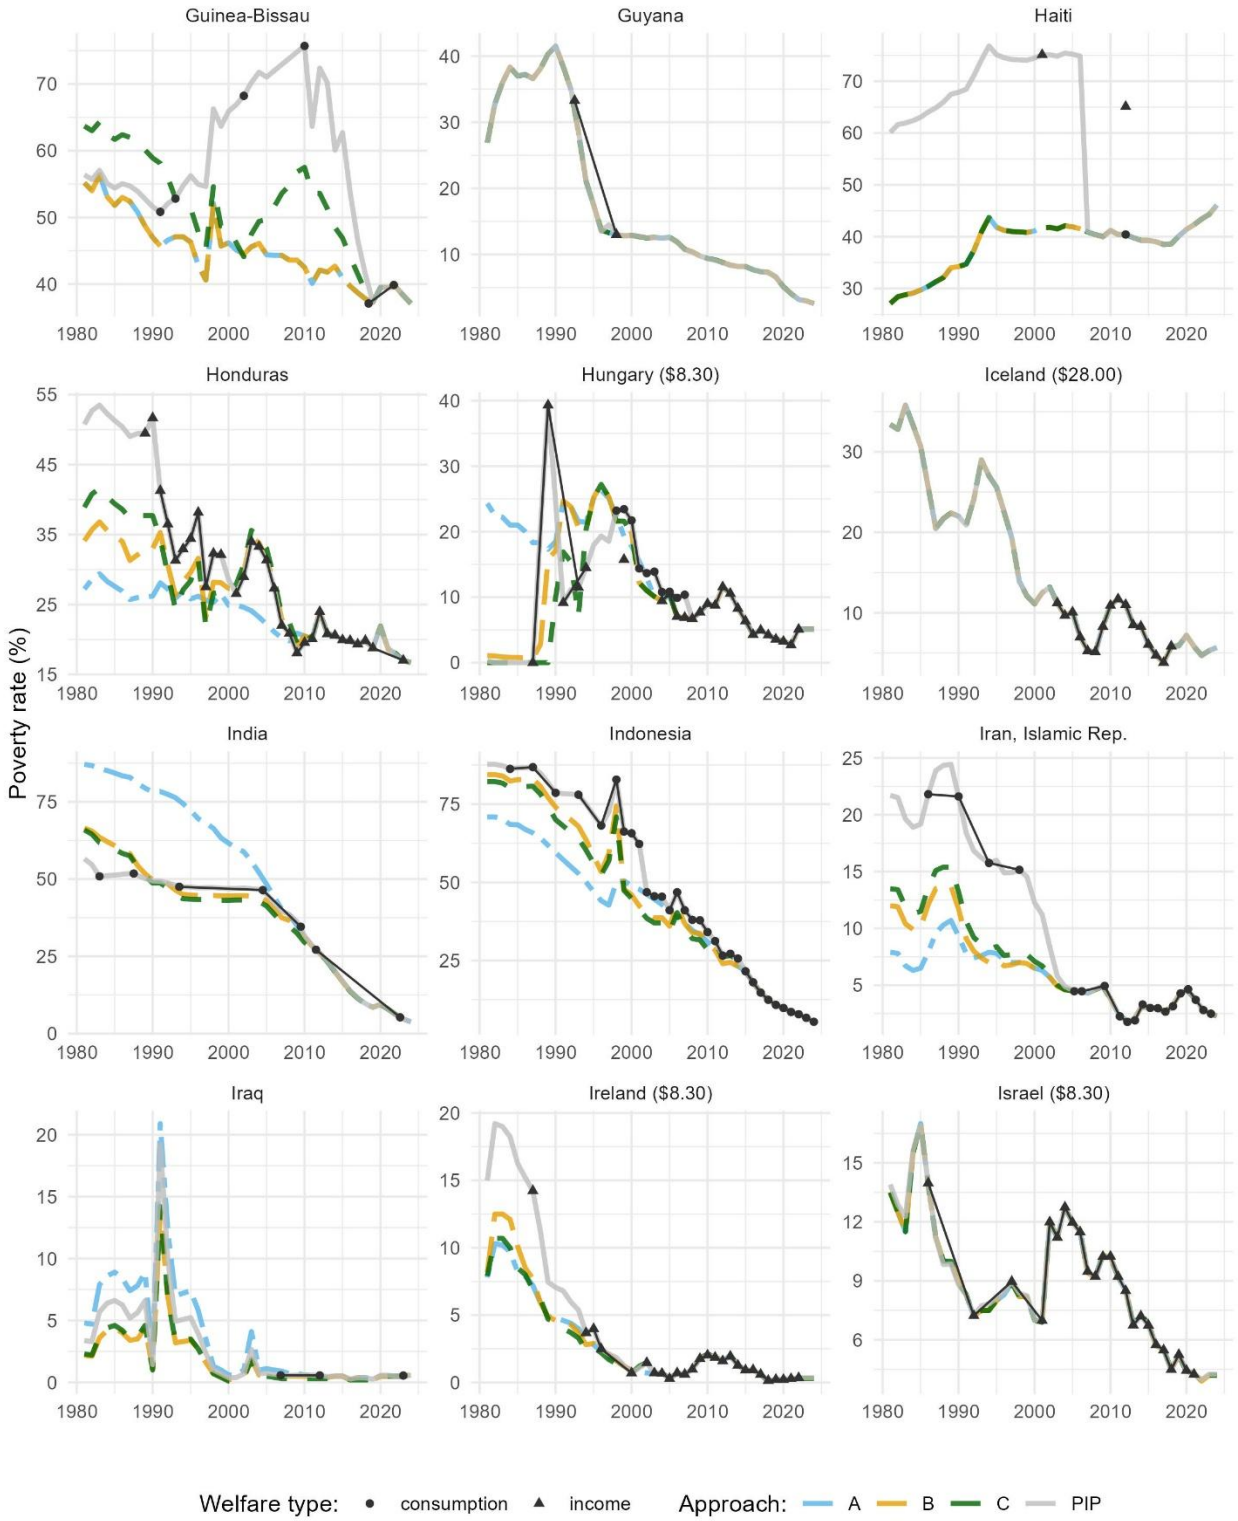

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.7: Impact on country-level poverty trends (7/15)

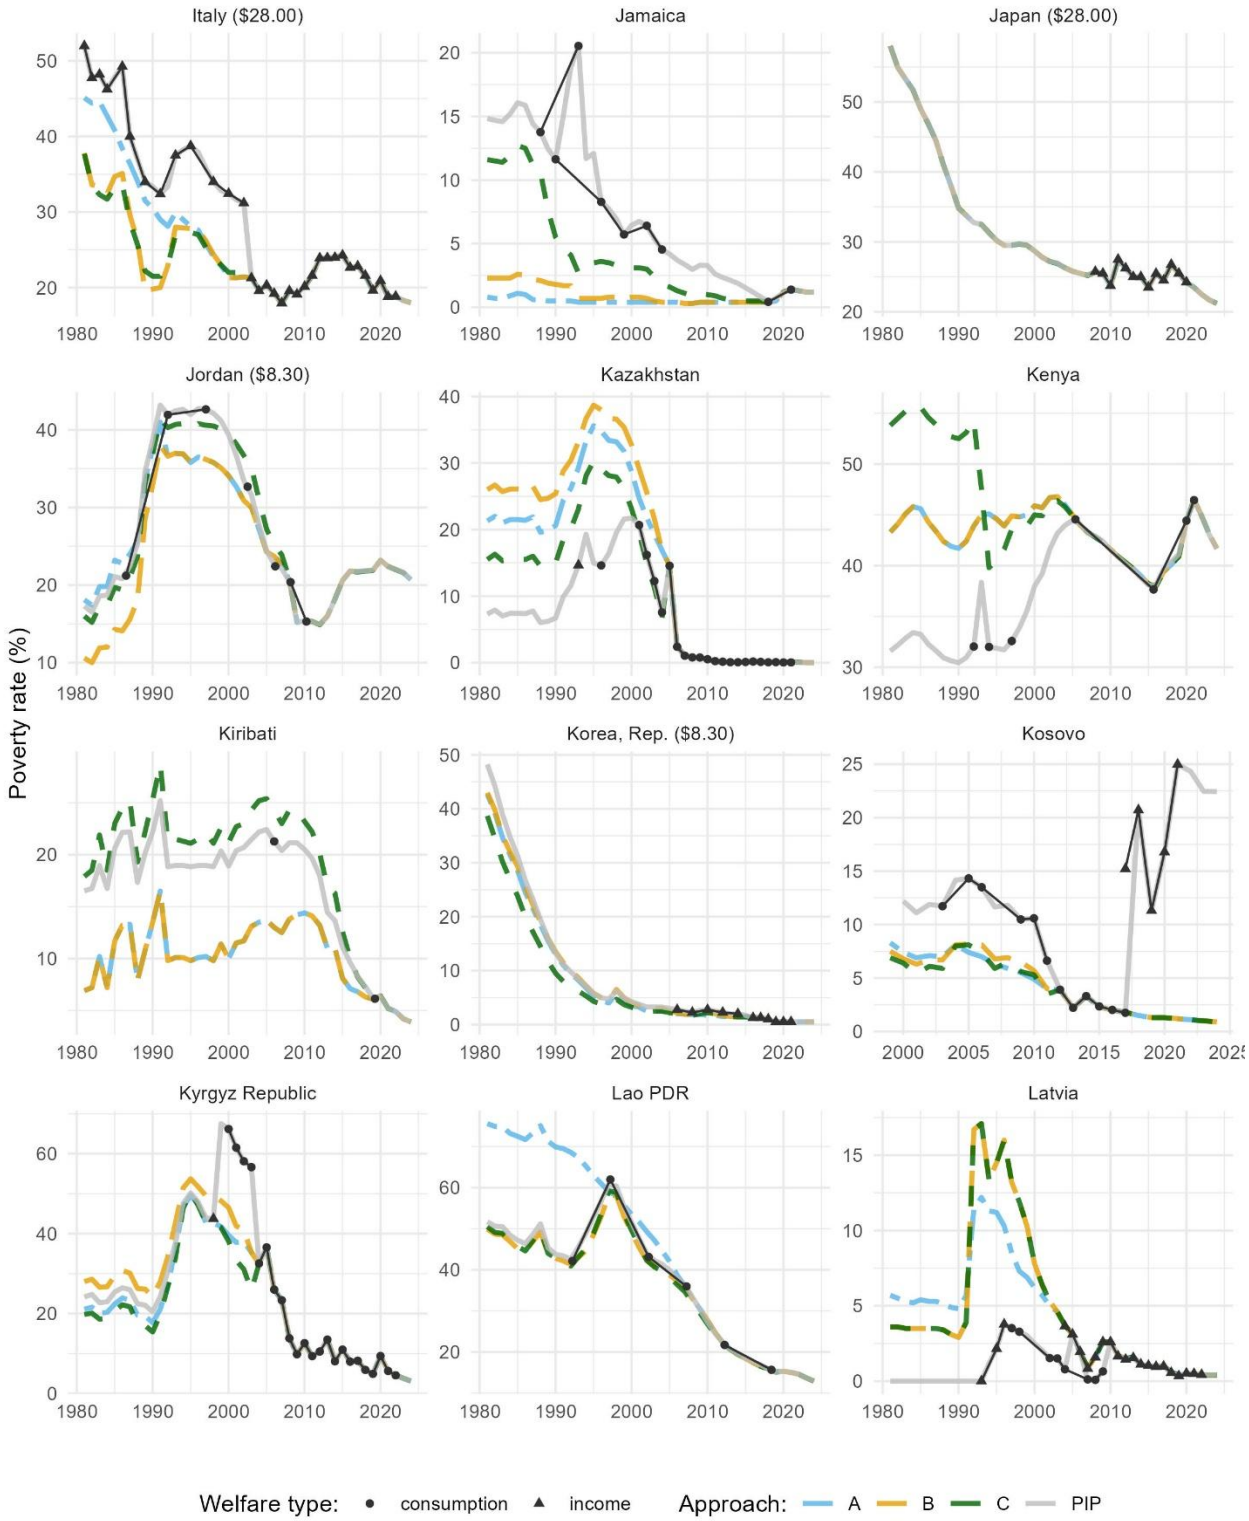

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.8: Impact on country-level poverty trends (8/15)

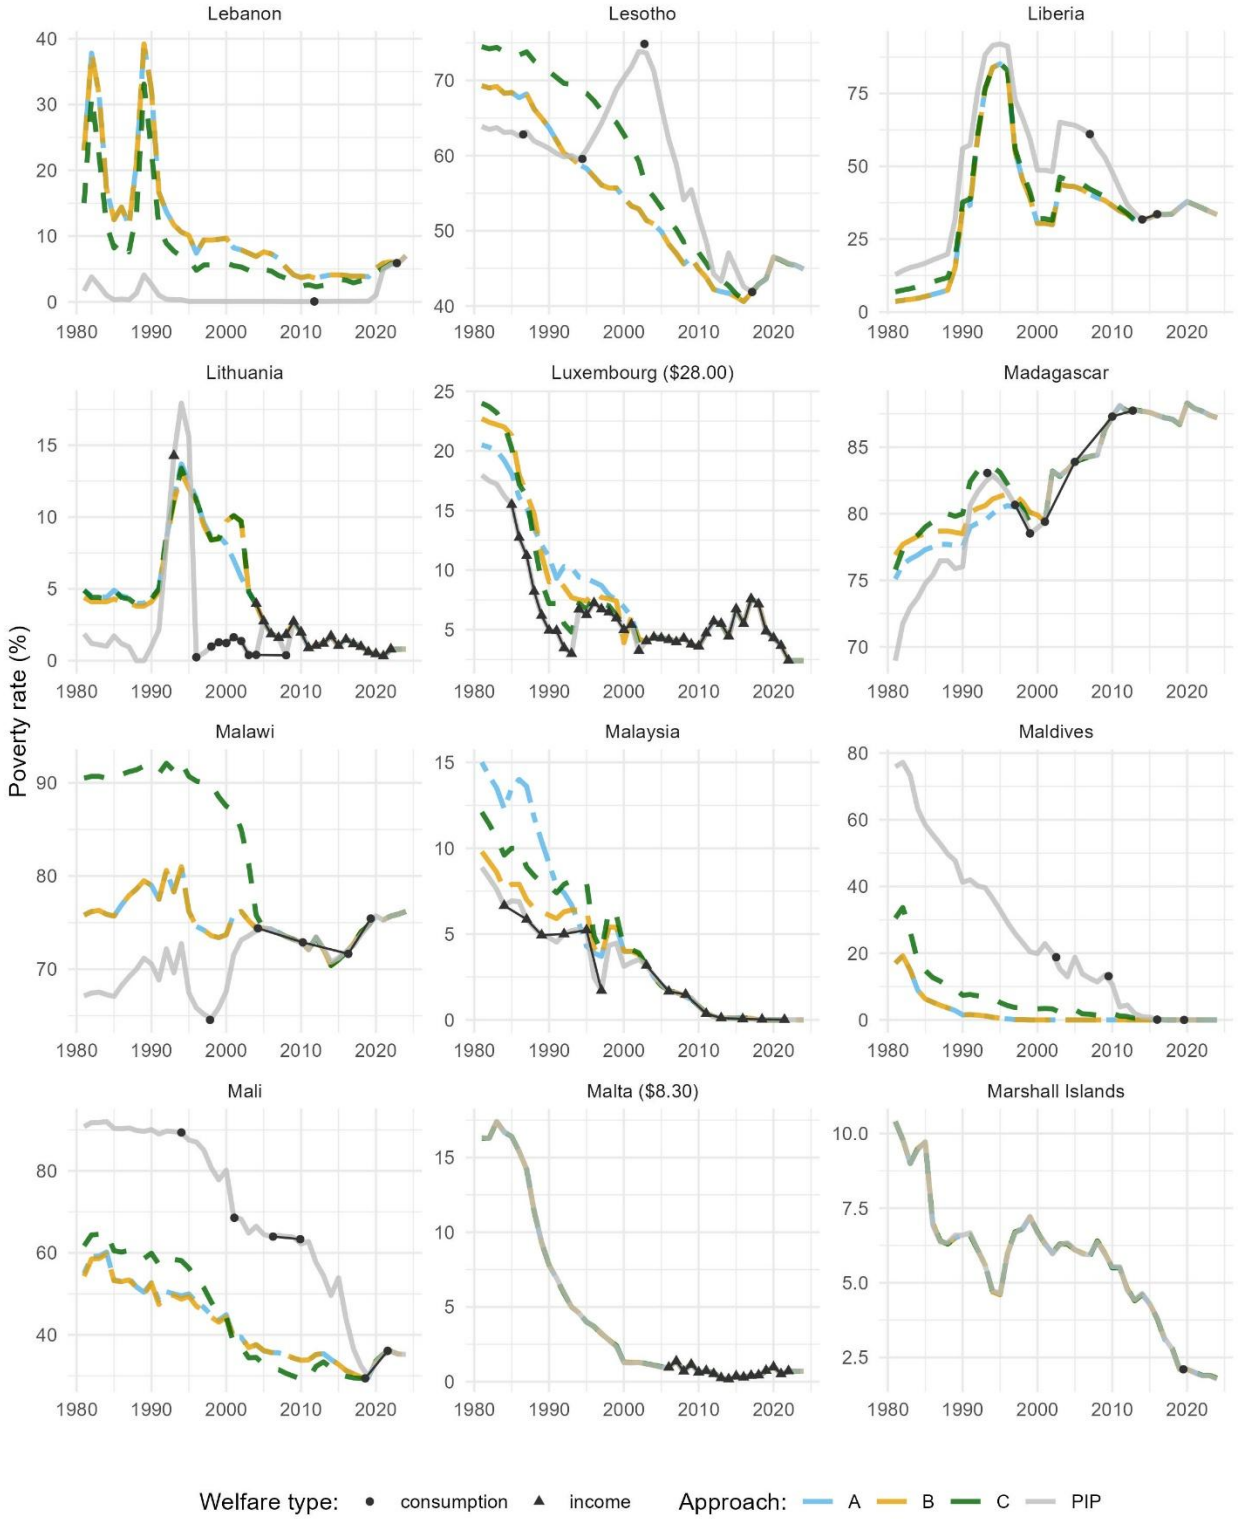

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.9: Impact on country-level poverty trends (9/15)

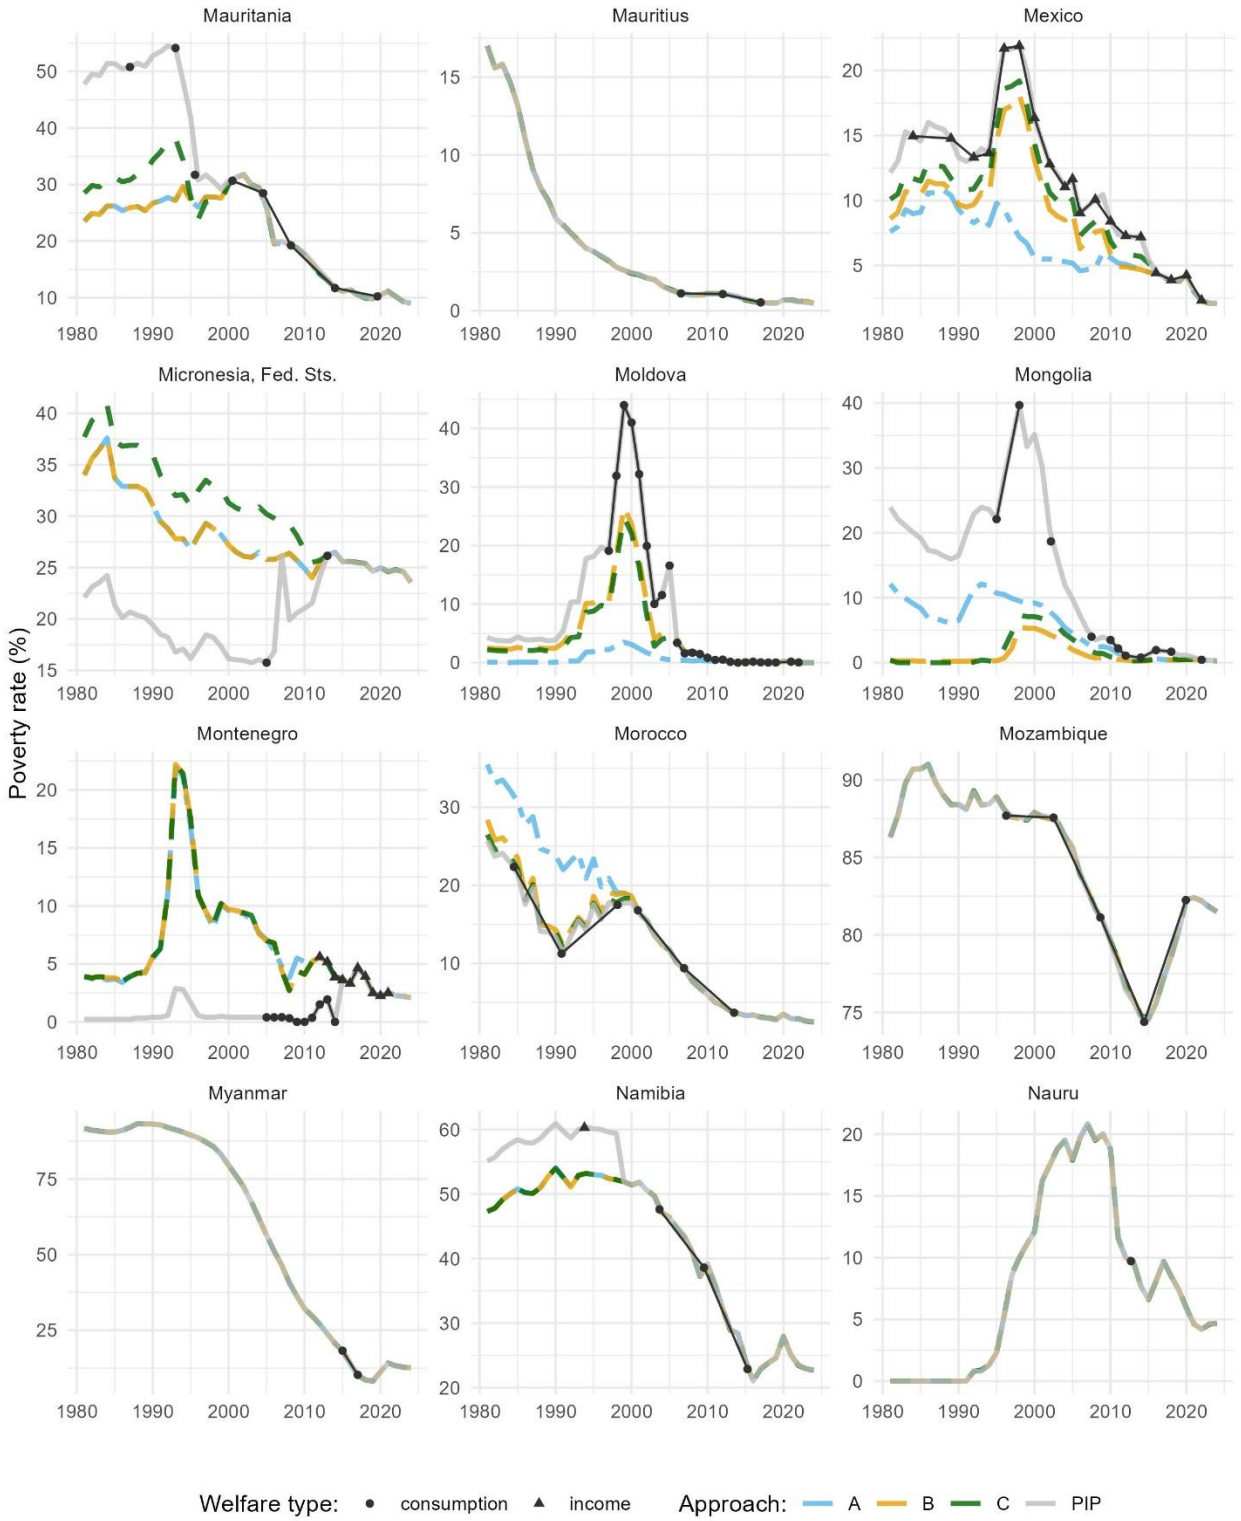

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.10: Impact on country-level poverty trends (10/15)

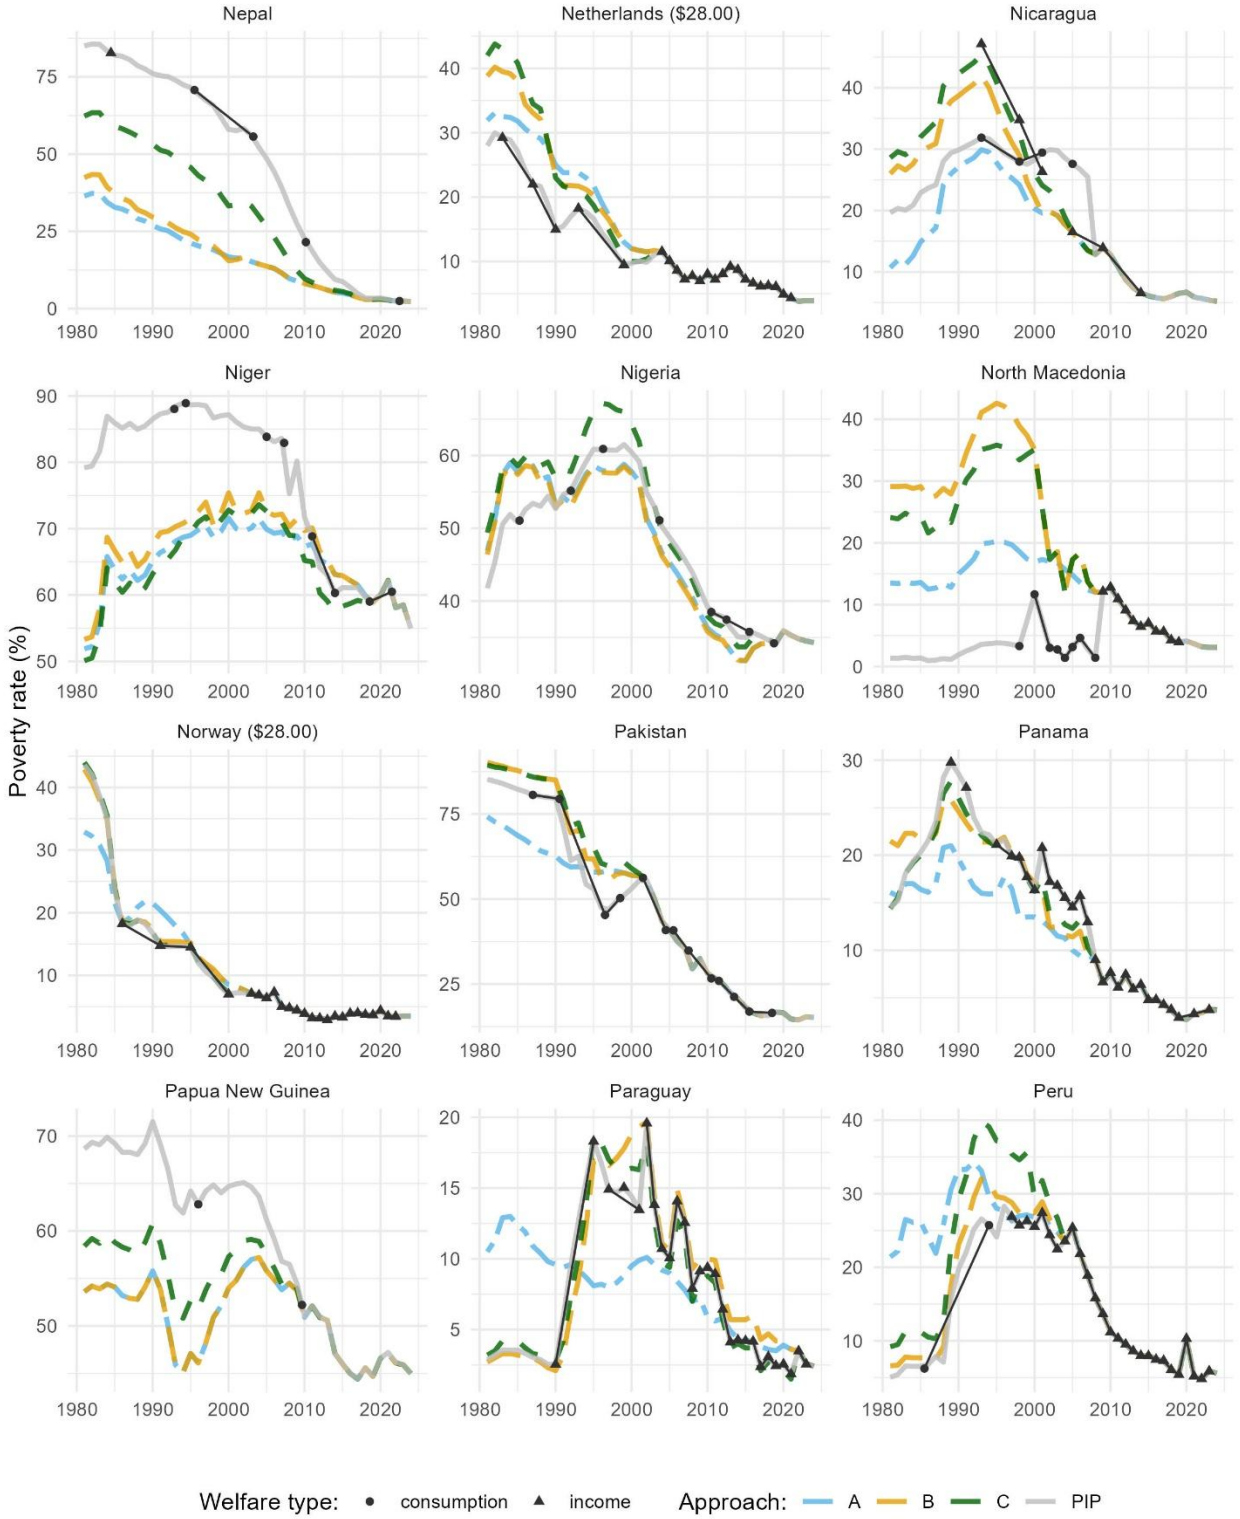

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.11: Impact on country-level poverty trends (11/15)

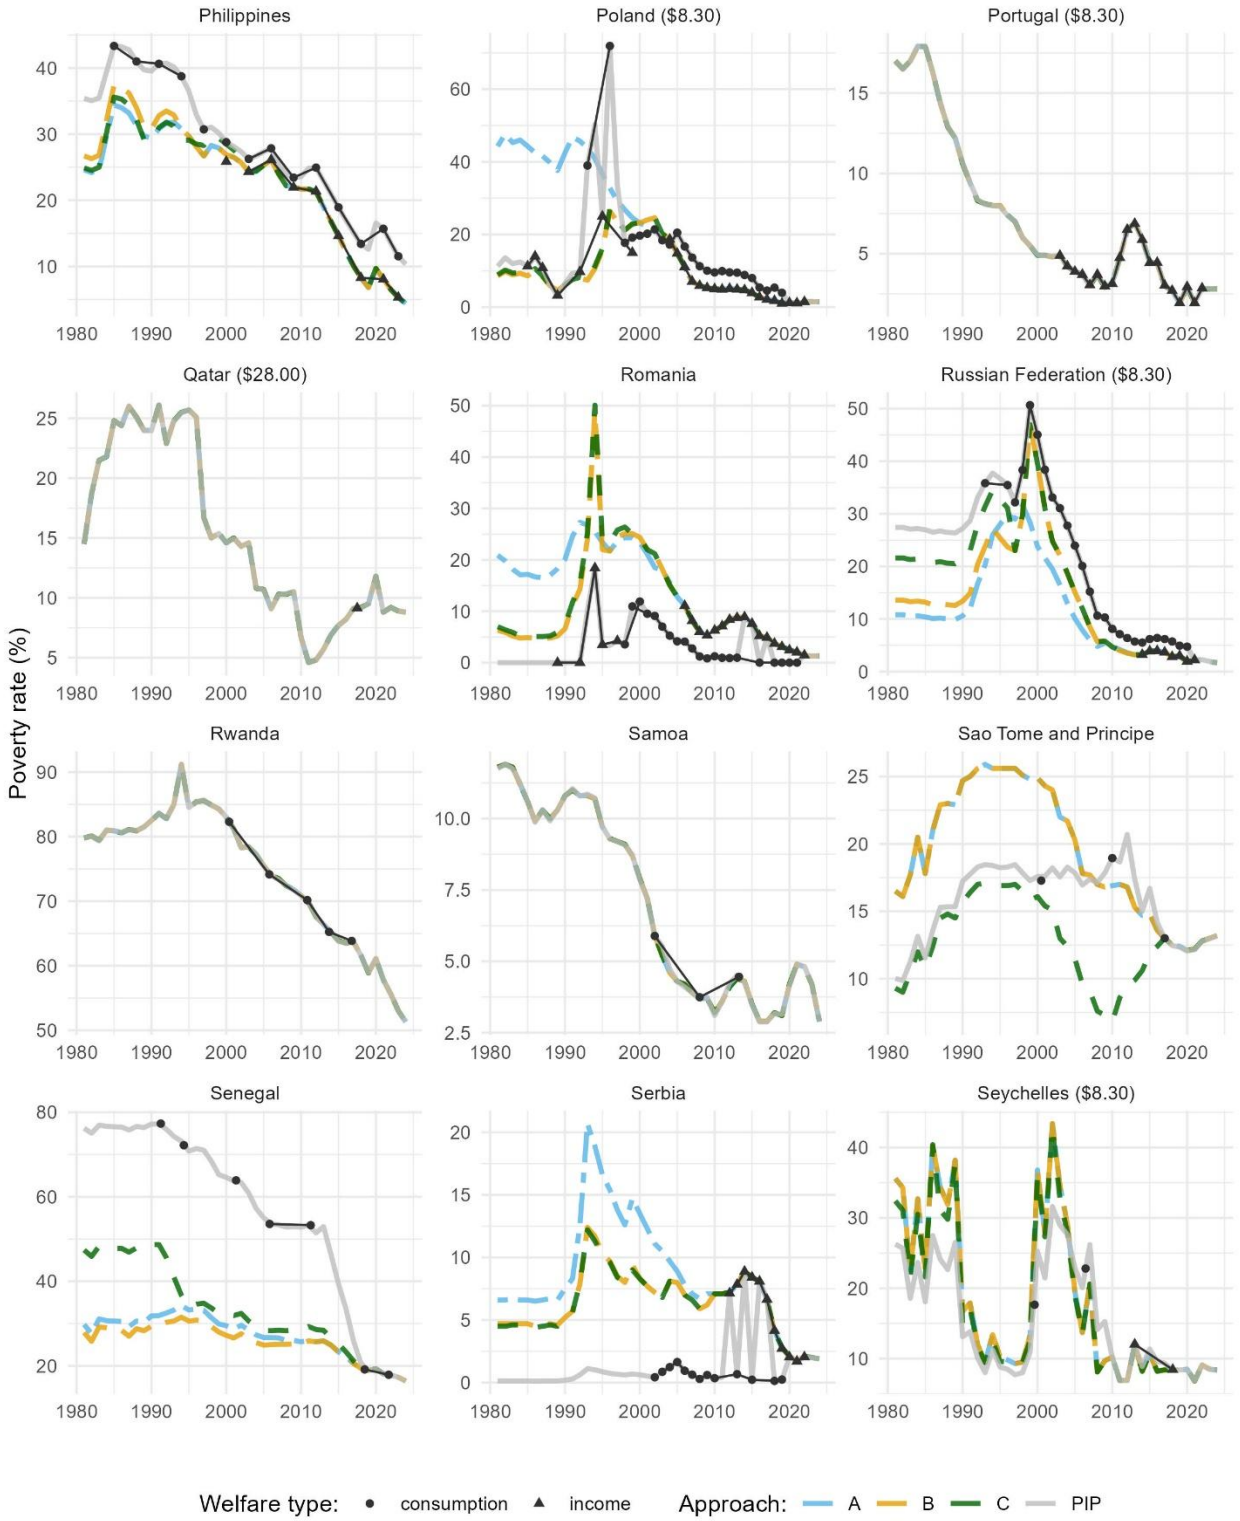

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.12: Impact on country-level poverty trends (12/15)

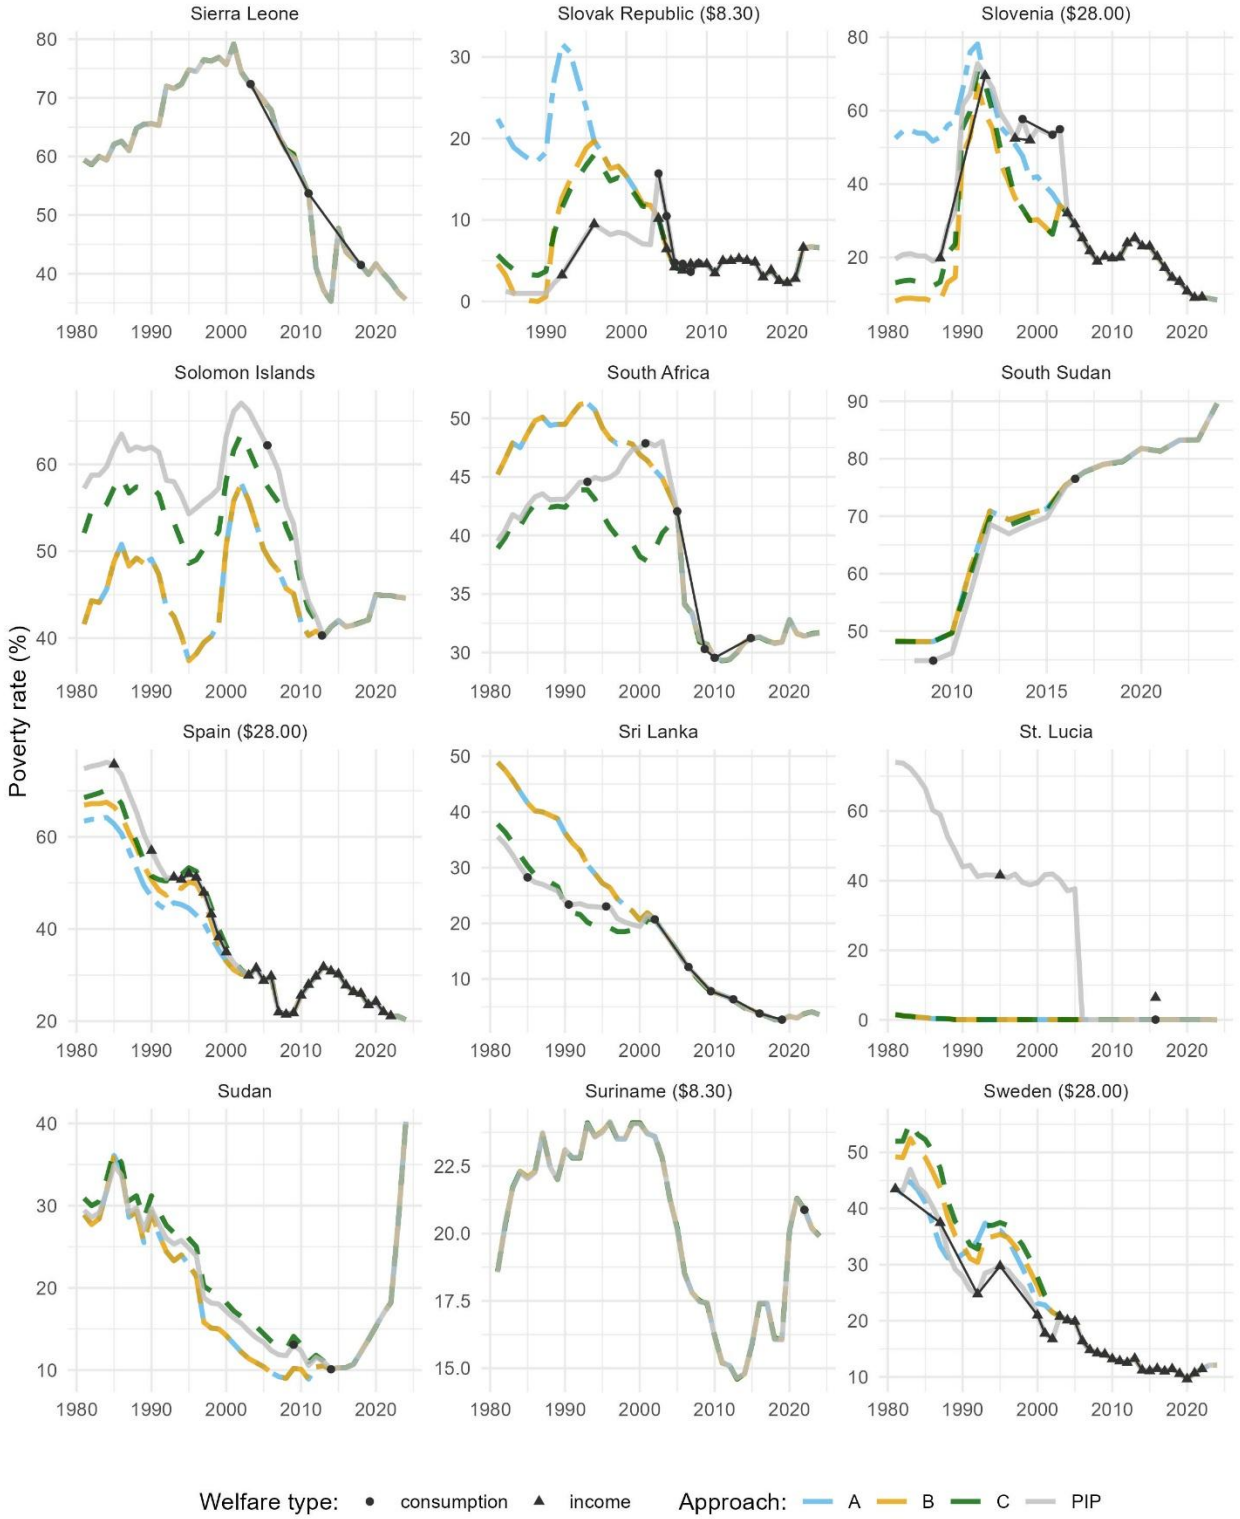

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.13: Impact on country-level poverty trends (13/15)

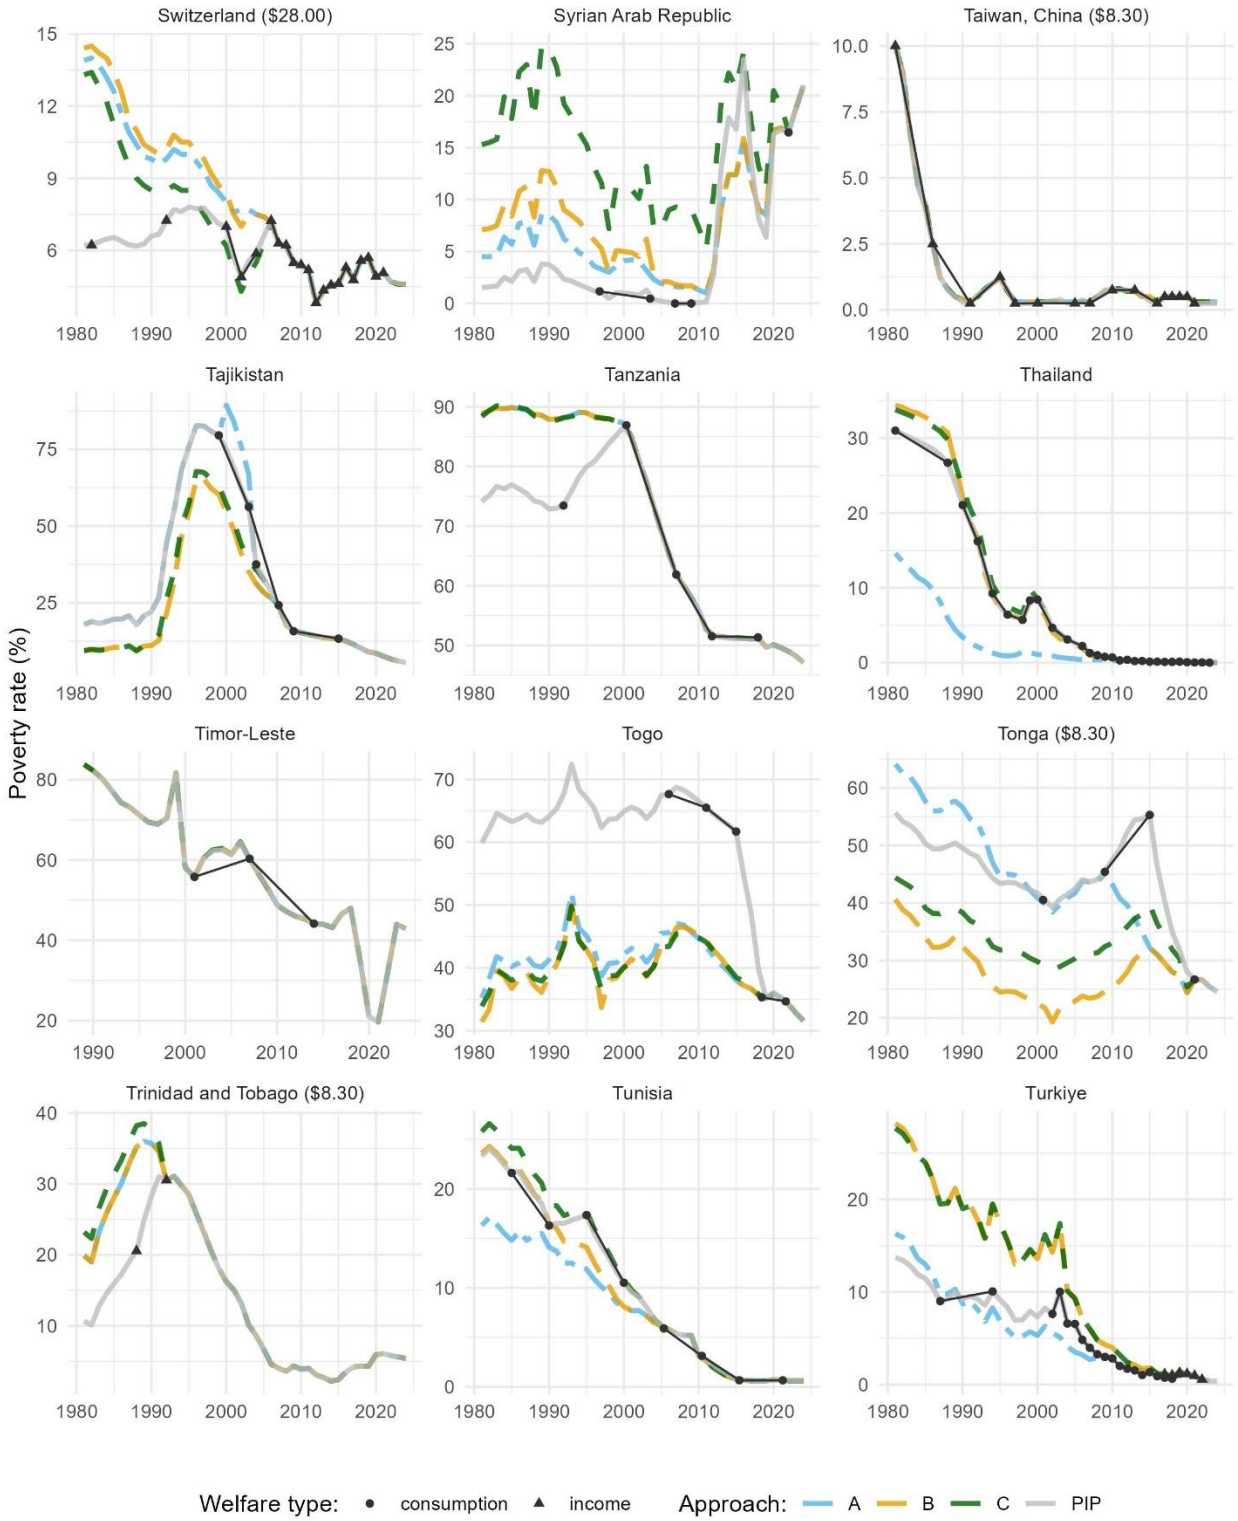

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.14: Impact on country-level poverty trends (14/15)

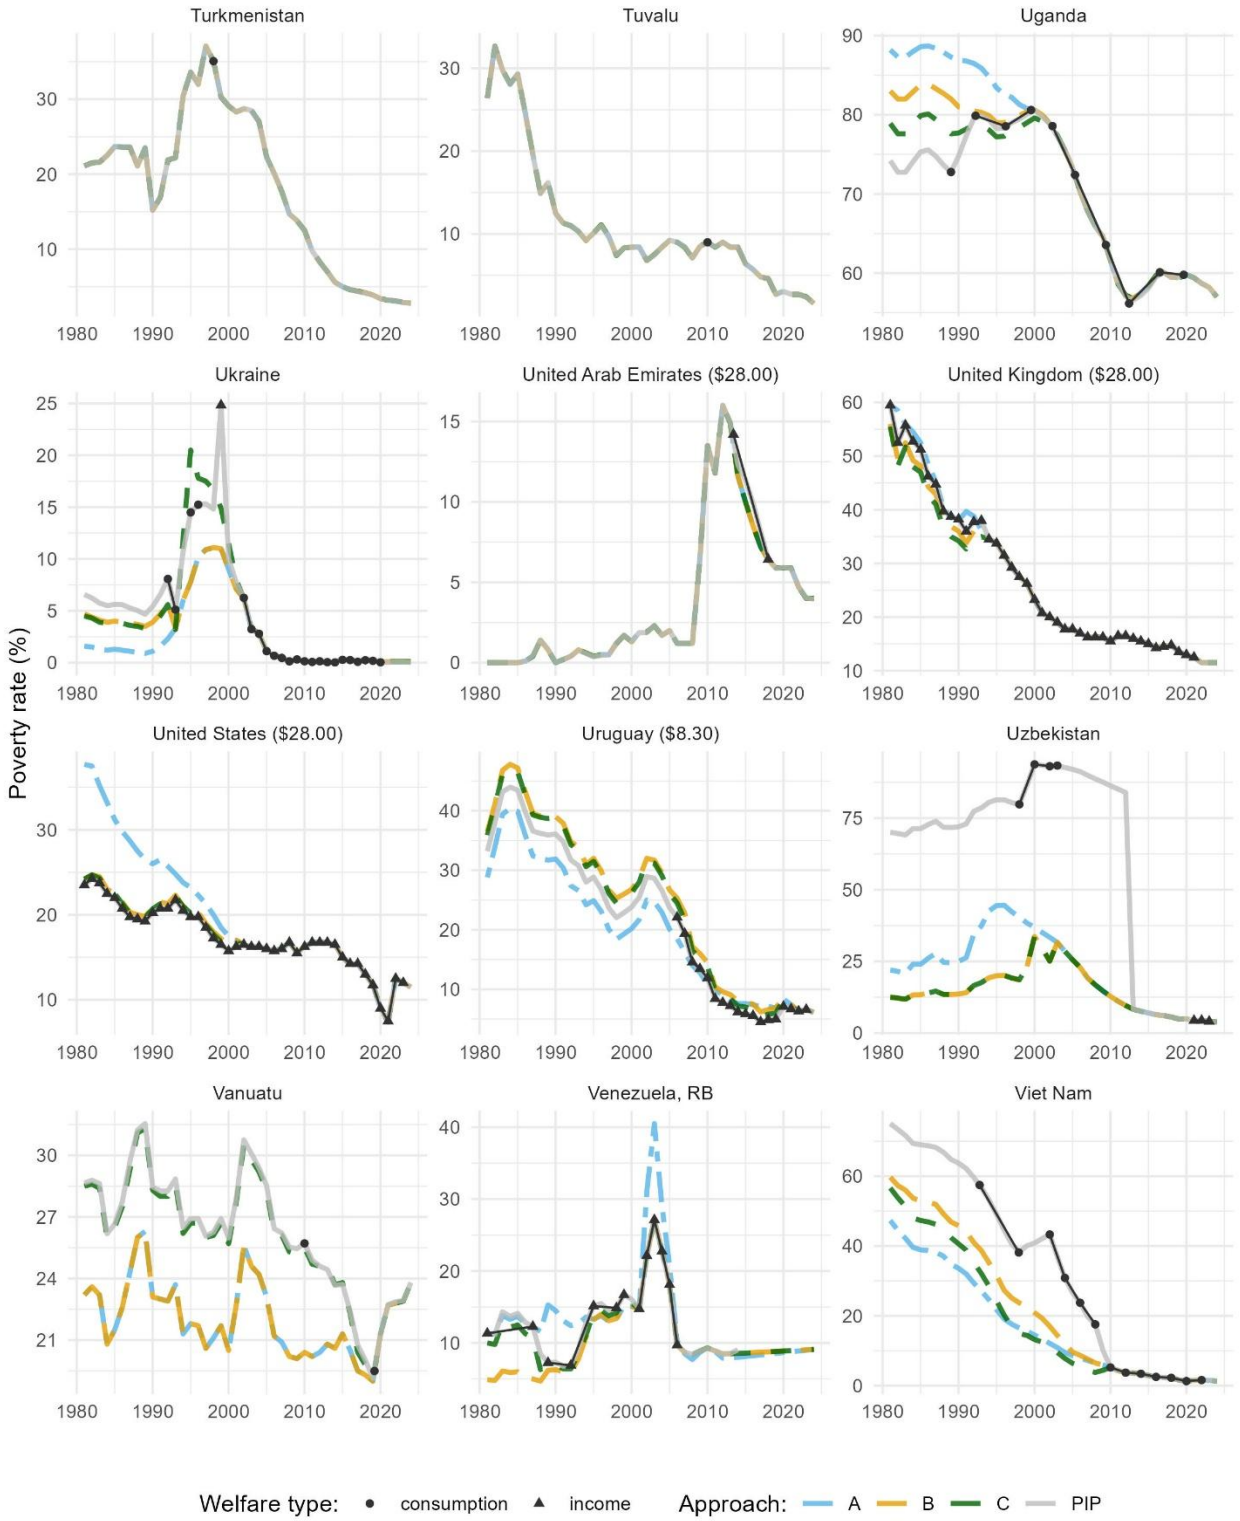

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.

Figure A.15: Impact on country-level poverty trends (15/15)

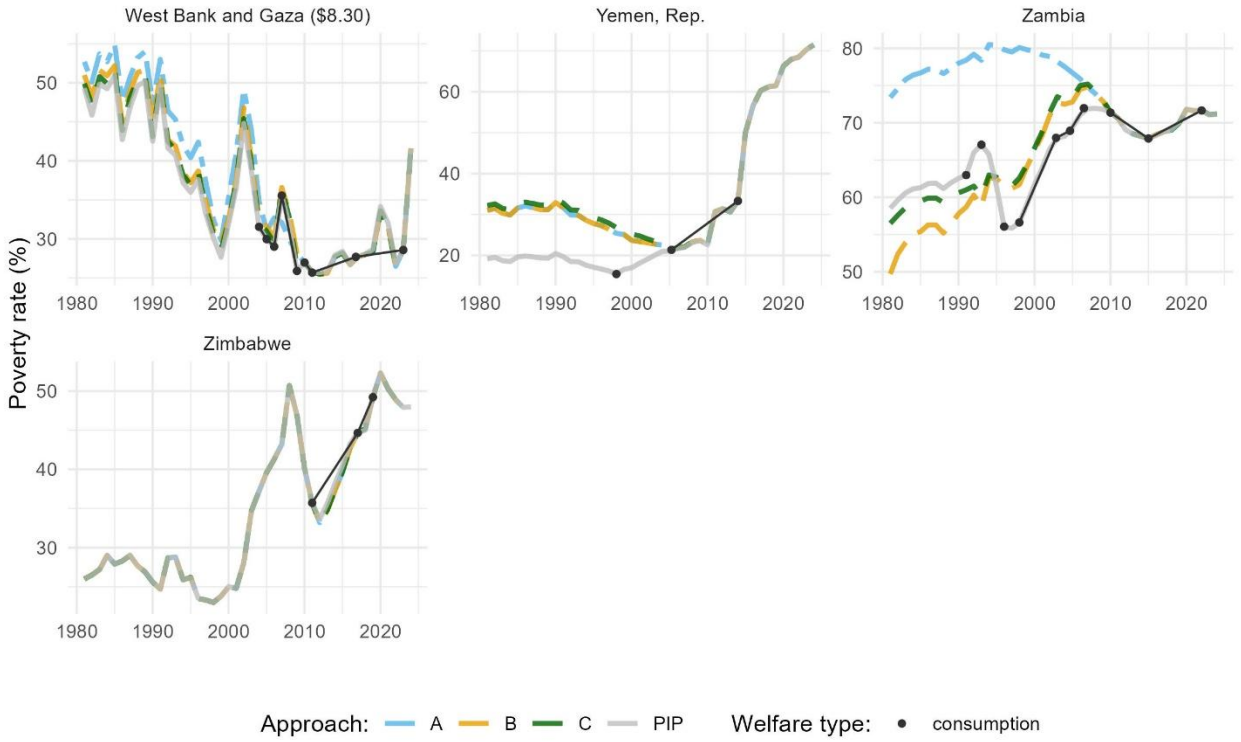

Note: Connected survey estimates indicate comparable data. The extreme poverty line is used unless poverty rates never exceed 10%, in which case the upper-middle-income poverty line ($8.30) is used, as mentioned next to the country name. If the poverty rates still never exceed 10%, the Prosperity Gap threshold ($28) is used.