Welfare & Poverty Measurement Using a Subset of Expenditures

Ethan Ligon

July 25, 2019

Some Problems in Poverty Measurement

1. Measuring *all* expenditures is expensive

If we could get by with a subset of expenditures we could survey more frequently.

2. Poverty lines depend on an arbitrary basket

Adjustments to prices or the poverty line based on the basket for a particular household gets the adjustment wrong for everyone else with a different basket.

3. Different household types are incomparable

Having more children affects the consumption basket, not just the size of the household.

A Question

Let

- x be the household's nominal expenditure budget
- p be a vector of all (nominal) prices
- z be household characteristics

Need

Suppose $\lambda(x, p, z)$ express the household's need as a function of budget, prices, and its characteristics.

Our Question

Does there exist such a function $\lambda(x, p, z)$ that solves these three problems?

Requirements

$\lambda(x, p, z)$ solves our problems if it satisfies the requirements:

1. Measuring *all* expenditures is expensive

Subsets valid: Our λ can be inferred from a subset of expenditures, and doesn't depend on *which* subset.

2. Poverty lines depend on an arbitrary basket

Independence: Our $\lambda(x, p, z)$ expresses need, so some fixed $\bar{\lambda}$ should express a fixed level of need, independent of prices or household composition.

3. Different household types are incomparable

Comparability: For any two household types z and z', and any prices, we can find budgets x and x' such that

$$\bar{\lambda} = \lambda(x, p, z) = \lambda(x', p, z').$$

The Household's Consumption Problem

A household with characteristics z facing prices p solves

$$V(x, p, z) = \max_{c \in \mathbb{R}^n} U(c, z) \qquad \text{such that } \sum_{i=1}^{n} p_i c_i \leq x.$$

This gives a set of n first order optimality conditions:

$$\begin{aligned} u_1(c, z) &= \lambda p_1 \\ u_2(c, z) &= \lambda p_2 \\ &\vdots \\ u_n(c, z) &= \lambda p_n \end{aligned}$$

where u_i denotes partial derivative of U w.r.t. c_i.

NB: λ the same across equations!

An Economist's Solution

Solution to the above is a set of demands $c_i(\lambda, p, z)$.

1. The system of n equations has the same λ in every one, so we can use any subset of equations we want to infer λ.
2. Using the budget constraint, λ must satisfy

$$\sum_{i=1}^{n} p_i c_i(\lambda, p, z) = x,$$

   which implicitly defines the function $\lambda(x, p, z)$.
3. Using the Envelope theorem,

$$\frac{\partial V}{\partial x}(x, p, z) = \lambda(x, p, z).$$

This λ satisfies all three of our requirements!

Can be interpreted as the marginal utility of expenditures.

Transition

What we've done so far

Established that a "neediness" function $\lambda(x, p, z)$ exists that

- Allows subsets of expenditures;
- Allows households with different consumption baskets to share the same poverty line, even when prices change;
- Allows comparison of need of households of different types (they can now share the same poverty line).

What we need to do next

Find a practical way of measuring λ for different households from subsets of expenditures.

Consumption Portfolio Approach

Rather than pretending that all households have the same "consumption portfolio", we *exploit* the fact that poor households don't just have smaller budgets, but that they have different consumption baskets than do wealthier households.

Example

Consider this expenditure data, for three households in the West Bank in 2016–17, with different household sizes.

Good	HH1	HH2	HH3
Long-grain rice	265.00	135.00	47.50
Frozen fish	82.00	30.00	50.00
Potato	50.00	80.00	20.00
White bread - Kmaj	84.00	195.00	53.50
Tomatoes	90.00	100.00	41.00
Tomato paste or solid (tinned)	10.00	35.00	8.00
Cucumbers	69.00	170.00	8.00
Featherless fresh chicken	190.00	340.00	155.00
Ground coffee	72.00	398.00	8.00
Imported chocolate	2.00	329.00	4.00

What can we infer about relative need?

NB: One HH from 1%, 50% and 99% of distribution.

Statistical approach

Recall first order equations. Stack for J households:

$$\begin{bmatrix} u_1^1(c^1) & u_1^2(c^2) & \cdots & u_1^J(c^J) \\ u_2^1(c^1) & \ddots & & u_2^J(c^J) \\ \vdots & & \ddots & \vdots \\ u_n^1(c^1) & \cdots & \cdots & u_n^J(c^) \end{bmatrix} = \begin{bmatrix} p_1 \\ p_2 \\ \vdots \\ p_n \end{bmatrix} \begin{bmatrix} \lambda^1 & \lambda^2 & \cdots & \lambda^J \end{bmatrix} = \mathbf{pL}^\top$$

If we observed LHS, we could recover $(\mathbf{p}, \mathbf{L})$ via matrix decomposition (e.g., first principal components of LHS matrix).

Methods for measuring neediness

We're after a latent index of wealth (compare Filmer and Pritchett (2001)), constructed from data on log item expenditures, after controlling for prices and household size and composition. Obtain this by estimating demand functions in two steps:

1. Use regression to relate demand to prices & household characteristics.
2. Obtain latent measure of wealth via principal components on residuals from regression.

Step 1: Controlling for prices and characteristics

Let l index a time and place (e.g., the West Bank in 2011). Assume all households at a given time and place face the same prices.

Notation

- y^j_{il} log expenditures on good i for household j who lives in region/period l.
- z^j vector of household characteristics, such as the number of boys, girls, men, and women.

Step 1: Controlling for prices and characteristics

The regression

The first stage estimates a system of partial demand equations:

$$y_{il}^{j} = \alpha_{il} + \delta_i' z^j + r_{il}^{j}. \tag{1}$$

- α_{il} Controls for prices at time and place l
- δ_i Good-specific influence of characteristics z^j.
- r_{il}^{j} Residuals which include all other influences on demand *except* prices and characteristics.

Step 2: Obtaining latent measure of wealth

Concatenate step 1 residuals

$$\mathbf{R} = [r^j_{1l}, r^j_{2l}, \dots, r^j_{nl}]$$

Residuals contain all information about (relative) wealth.

- Wealthier households have larger r^j_{il}, other things equal; and
- By controlling for prices and observed characteristics we've held other things equal.

Step 2: Obtaining latent measure of wealth

Calculate first principal components

Use principal components to estimate both the latent components and weights which best explain the variation in residuals:

$$\mathbf{R} = \mathbf{LB},$$

L First column are the wealth index we want; call these $\log \lambda_l^j$, identified up to location and scale.

B First row are weights; call these β_i.

Scale We impose the standard normalization which sets the sample variance of $\mathbf{L}$ to one.

Interpretation of $\log \lambda$

From estimation above we have a neediness index $\log \lambda$ that vary across households, and a vector of weights β that vary across goods. There are several ways to interpret these.

Informal Household neediness—value household places on an additional shekel of expenditures.

Economic Log of marginal utility of expenditures—household will try to keep this constant over time.

Linear algebra Left singular vector of matrix **R** corresponding to largest singular value.

Statistical First principal component of log expenditures, after controlling for time-place effects and observable household characteristics.

Interpretation of β_i

Informal How responsive consumption of a particular good is to changes in wealth.

Economic Proportional to income elasticities of different goods.

Linear Algebra Right singular vector of matrix $\mathbf{R}$ multiplied by largest singular value.

Statistical Square root of diagonal of Fisher information matrix.

Subsetting Expenditures

We can do estimation for many goods (the PECS has more than 700). Not all are equally informative:

- Some goods aren't observed often enough estimate demand;
- Some goods won't have expenditures that vary much with wealth; (e.g., salt). If an estimated weight β is close to zero that is evidence that expenditures on that good aren't vary informative.

Criterion for subsetting

Goods with smallest β_i are least valuable. Drop sequentially until overall fit of demand system falls (trade-off between loss of information and cost of items in expenditure module).

Alternative criterion

Use Analysis of Variance to identify goods that have most information. In our application produces similar results.

Experiments & Validation

We're done with main discussion of methods. Next, we discuss some Monte Carlo experiments to validate the consumption portfolio approach, and also to try some other ideas for using subsets of expenditures.

Data

We work with a particular dataset on household expenditures from the Palestinian Expenditure and Consumption Survey (PECS) collected in 2011.

- Notable for large number of expenditure items—over 700.
- Interesting question about welfare & poverty in West Bank vs. Gaza Strip.
- Out of over 700 goods, Khawaja (1998) uses 155 in a "basic basket" to calculate poverty in Palestine in 1997.

Summary Statistics

	Gaza Strip	West Bank	Pooled
Boys	1.75	1.32	1.46
	(1.42)	(1.32)	(1.37)
Girls	1.68	1.25	1.39
	(1.52)	(1.31)	(1.40)
Men	1.62	1.57	1.59
	(1.12)	(1.13)	(1.13)
Women	1.62	1.55	1.58
	(0.98)	(0.91)	(0.93)
Rural	0.21	0.29	0.27
	(0.41)	(0.46)	(0.44)
Camp	0.23	0.20	0.21
	(0.42)	(0.40)	(0.41)
HSize	6.67	5.69	6.01
	(2.79)	(2.66)	(2.74)
N	1408	2909	4317

Selected Expenditure Data

Expenditure Category	Agg. shares	Mean shares
Transport & Communication	0.303	0.140
Furniture and Utensils	0.126	0.125
Education	0.126	0.129
Other Non-Food Consumption Expenditure	0.101	0.058
Medical Care	0.061	0.040
Housing	0.042	0.094
Meat and Poultry	0.032	0.058
Recreation	0.028	0.027
Clothing & Footwear	0.023	0.034
Bread & Cereals	0.021	0.046
Tobacco	0.017	0.030
Other Cash Non-Consumption Expenditure	0.014	0.017
Vegetables, Legumes, & Tubers	0.014	0.032
Fruits & Nuts	0.012	0.026

Mean vs. Aggregate Shares

Where mean and aggregate weights differ it's a sign baskets differ across rich and poor; thus CPI bias, *especially* for poor people.

Evidence of CPI bias

A statistic ρ_{it} measures this bias: for good i at time t,

$$\rho_{it} = \log\left(\frac{\sum_{j=1}^{N} x_{it}^{j}}{\sum_{j=1}^{N}\sum_{k=1}^{n} x_{kt}^{j}}\right) - \log\left(\sum_{j=1}^{N}\frac{x_{it}^{j}}{\sum_{k=1}^{n} x_{kt}^{j}}\right).$$

Aggregate vs. Mean Shares

Using *all* expenditure items, aggregated into categories

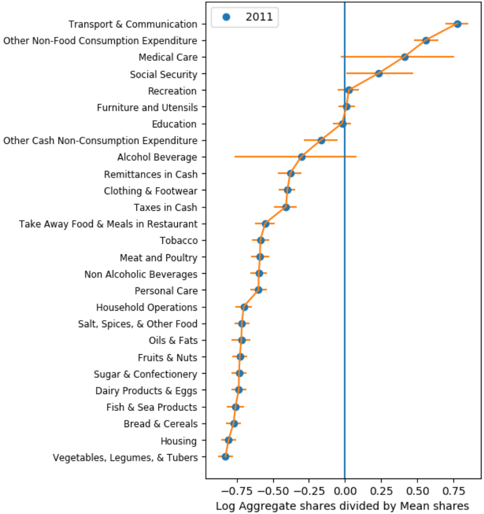

Plan

- Estimate demand system using full set of 155 "basic basket" goods; validate by showing that *estimated* demands reproduce the observed cross-sectional distribution of total expenditures.
- Monte Carlo experiments to characterize sampling error & measure reliability; reliability here implies we can also reliably construct standard measures of poverty and inequality.
- Reduce set of goods and re-estimate; validate using this *reduced* set of goods. We consider three ways of reducing the set of goods:
  1. Using demand system approach outlined above;
  2. Use only goods with a large expenditure share;
  3. Us a random subset from different households; overall we observe expenditures from some households on all items.

Estimates of β and δ (more elastic goods)

Goods	β_i	Camp	Rural	log H	R^2
Imported chocolate	0.59	0.04	-0.07	0.16	0.29
Oriental deserts	0.56	-0.14^{*}	-0.05	0.31^{*}	0.33
Fresh goat & sheep meat	0.55	-0.18	0.02	0.16	0.26
Featherless fresh chicken	0.51	-0.17^{**}	-0.08	0.22	0.53
Cake & Cookies	0.50	0.12	-0.23^{***}	-0.34	0.19
Fresh or pasteurized milk	0.48	-0.31^{***}	-0.34^{***}	0.13	0.25
Olive oil	0.47	-0.11	0.07	0.47^{*}	0.34
Bonbon, citrus products	0.45	-0.35^{***}	-0.19^{*}	-0.75^{**}	0.18
Fresh beef meat	0.44	-0.32^{***}	-0.08	0.22^{*}	0.24
Soft drinks, family size	0.42	-0.14^{***}	0.07	0.17	0.27
Soup (cubes)	0.40	-0.07	-0.04	0.57^{***}	0.20
Soft drinks, can	0.40	-0.07	-0.12^{*}	-0.08	0.16
Cardamom	0.40	-0.19^{*}	0.07	-0.73^{*}	0.34

Estimates of β and δ (less elastic goods)

	β_i	Camp	Rural	log H	R^2
Soft white cheese	0.15	-0.06	0.04	0.08	0.15
Jam	0.14	-0.12^{**}	-0.03	0.02	0.13
Dry beans	0.14	-0.13^{***}	0.11^{**}	-0.04	0.14
Yogurt paste	0.13	-0.22^{***}	-0.23^{***}	0.35^{***}	0.10
Lentils	0.13	-0.10^{**}	0.12^{***}	-0.08	0.13
Frozen chicken	0.12	0.03	-0.01	-0.14	0.07
Gazelle fat	0.11	-0.09^{**}	-0.03	0.07	0.29
Yeast	0.10	-0.02	0.03	0.06	0.09
Sunflower oil	0.10	-0.09^{**}	-0.15^{***}	0.25^{**}	0.20
Crushed lentils	0.09	-0.05	0.10^{***}	-0.18^{**}	0.11
Concentrated juice	0.06	0.16^{**}	0.16^{***}	0.09	0.06
Garbage disposal	0.03	-0.23^{***}	-0.03^{**}	0.09^{**}	0.04
Sewage fees	0.02	-0.24^{***}	-0.25^{***}	-0.08	0.57
Imported white flour	0.02	-0.08	0.42^{***}	0.08	0.27

Predicting total expenditures (full basic basket)

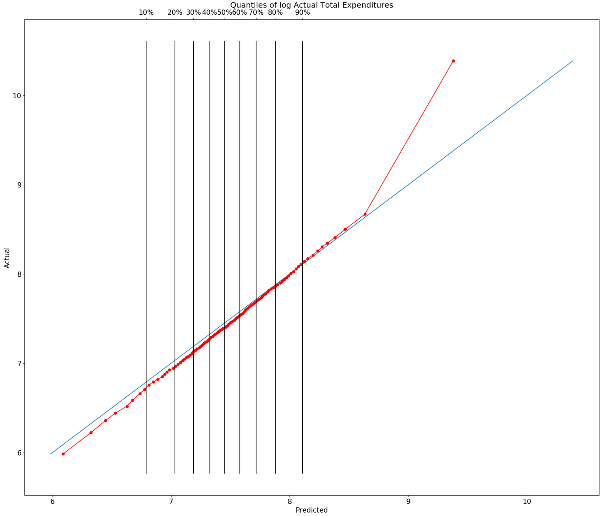

Using goods with larger income elasticities

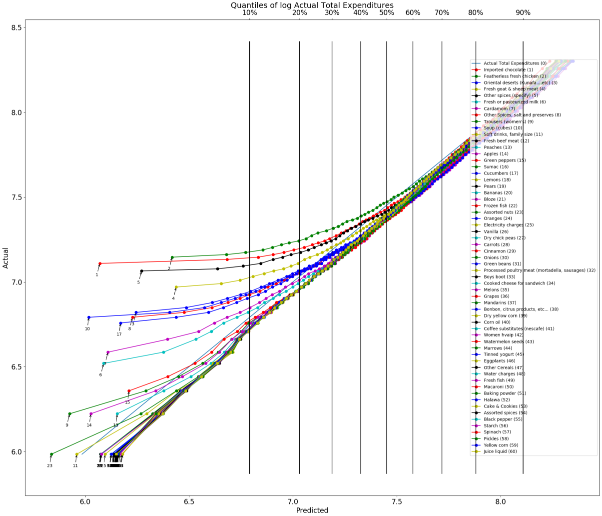

Takeaways (using more income elastic goods)

The consumption portfolio approach ranks goods according to how income elastic they are.

- By just using the 25 most income elastic goods we can predict the distribution of total expenditures just as well as one can by measuring 155 (or more) goods!

Using goods with larger expenditure shares

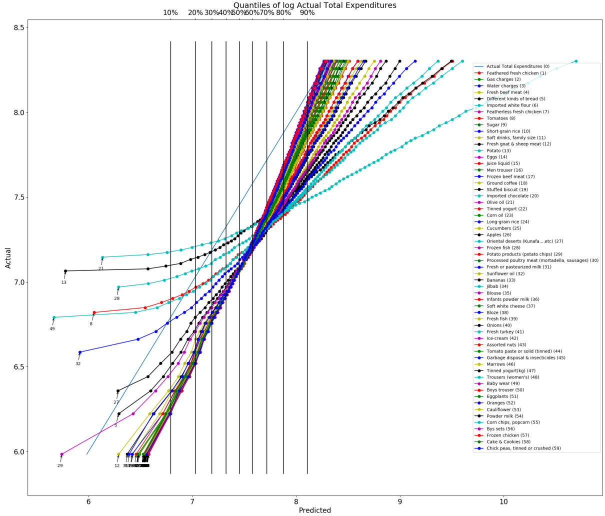

Takeaways (goods with larger shares)

1. This approach is terrible. No subset does well at predicting total expenditures.
2. Mostly subsets dramatically underestimate poverty.
3. Some funny outliers: These are subsets that add one of the highly income elastic goods.
4. Christiaensen, Ligon, and Sohneson (2017) present theoretical results predicting this poor performance.

Using different subsets for different households

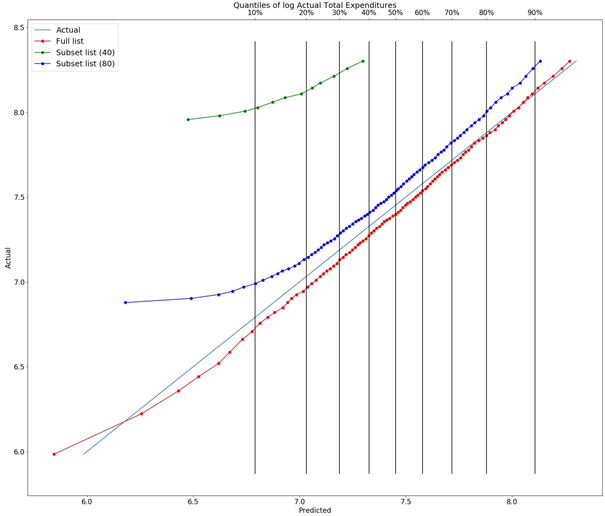

Takeaways (different subsets for different households)

1. Systematically *under*-estimate expenditures at all levels of wealth, but particularly for the very poorest households.
2. Can't even estimate for many households.
3. Issue related to missing data—if we only ask about a small set of goods poorer households may not have purchased many of those. Wealthier households buy more, but also more different kinds of goods.
4. Predictions only for best off households, but prediction is wealthiest are among poorest!

Welfare & Poverty in Palestine

Next we apply methods above to understand welfare & poverty in Palestine. Our focus is on household-level head-count poverty.

We do everything twice:

1. Using full basic basket of 155 goods;
2. Using selected basket of 25 more income elastic goods.

Note on sampling weights

1. We don't use sampling weights in estimation; estimates here are for the *sample*.
2. We do generally weight estimates when we make claims about population statistics.

Neediness & the Marginal utility of expenditures

- Interpret the statistic λ as the value of a shekel to the household, relative to other households.
- Within the sample the average of $\log \lambda$ is zero by construction in every period and every "market" (i.e., WB & GS).
- Larger values of λ indicate greater neediness; giving a shekel to a household with value of $\log \lambda = 0.05$ indicates that giving a shekel to this relatively poor/needy household would have 5% more value than to the average household.
- This is a utilitarian interpretation, and basically requires that households all have the same weight in the social welfare function.

Returns to a shekel (basic basket/reduced basket)

Table: Returns to an additional shekel to households in particular quantiles of the sample log λ distribution; returns to household at mean are zero.

Quantile	Pooled	West Bank	Gaza Strip
0.5	$-0.01/-0.04$	$-0.00/-0.03$	$-0.02/-0.06$
0.6	0.24/0.22	0.25/0.23	0.21/0.19
0.7	0.52/0.49	0.54/0.54	0.49/0.43
0.8	0.83/0.85	0.86/0.89	0.79/0.72
0.9	1.27/1.30	1.29/1.33	1.23/1.22
0.95	1.64/1.67	1.64/1.71	1.61/1.58
0.99	2.36/2.49	2.45/2.57	2.27/2.30

Distribution of $\log \lambda$ (basic basket)

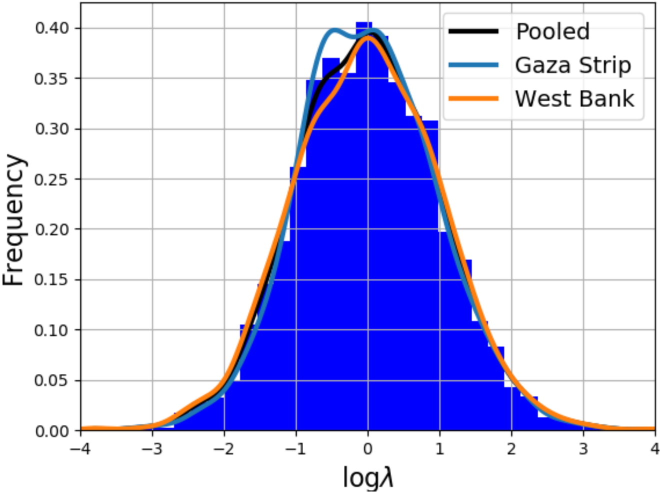

Distribution of $\log \lambda$ (reduced basket)

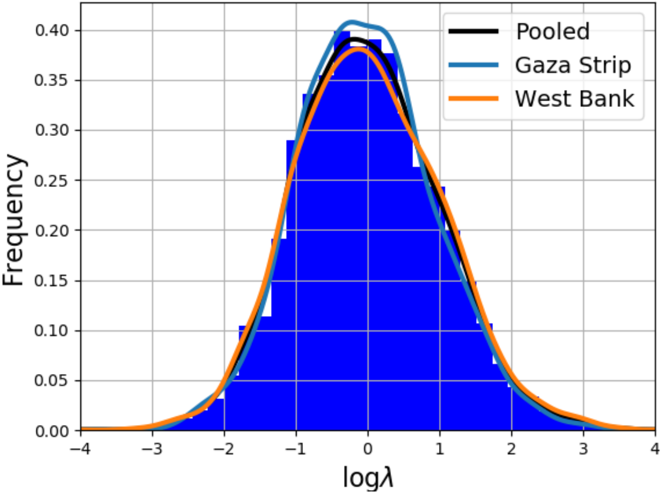

Comparing different areas

The distributions of $\log \lambda$ are separately centered at zero in the West Bank & Gaza Strip; this allows for different prices in the two areas.

- We obtain shadow prices from our demand system.
- Prices highly correlated (92%) across GS & WB;
- Most prices higher in the West Bank

Shadow Prices

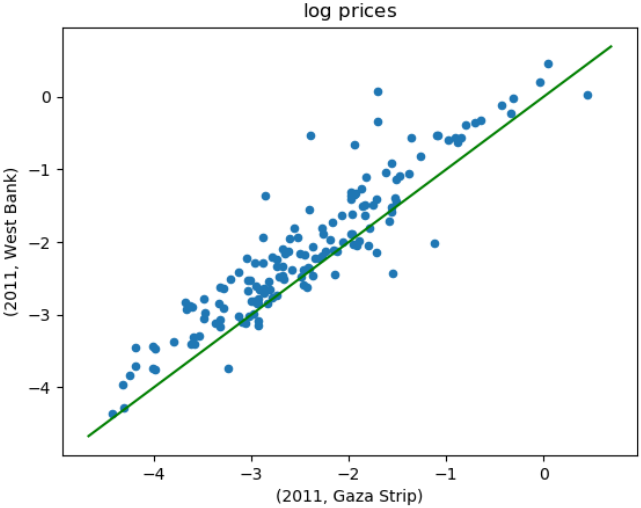

Goods with extreme price differences

Expenditure Item	Gaza Strip	West Bank	Difference
Sewage fees	−2.393	−0.533	1.860
Featherless fresh chicken	−1.699	0.075	1.773
Short-grain rice	−2.862	−1.358	1.505
Olive oil	−1.698	−0.339	1.359
Different kinds of bread	−1.945	−0.657	1.288
Tinned yogurt(kg)	−2.879	−1.930	0.949
Juice liquid	−2.404	−1.556	0.848
Assorted spices	−3.674	−2.830	0.844
Treated thyme	−3.047	−2.229	0.817
Fresh fish	−1.360	−0.553	0.807
Coffee substitutes (nescafe)	−2.562	−1.811	0.750
Other expenditure on dwelling	−1.553	−2.426	−0.873
Imported white flour	−1.125	−2.007	−0.882

The Optimal Price Index

If a particular household in Gaza suddenly faced the prices prevailing in the West Bank, how much would its expenditures have to increase in order to maintain the same $\log \lambda$? *Since we've estimated the demand system* we can answer this question.

The Optimal Price Index

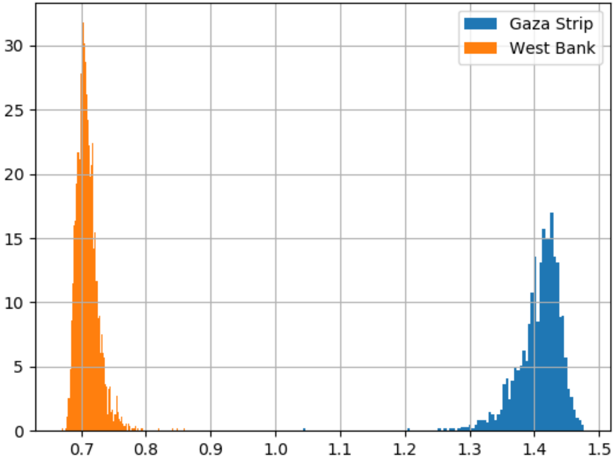

Total exp. changes by 0.34 log points between GS & WB

The Optimal Index (restricted basket)

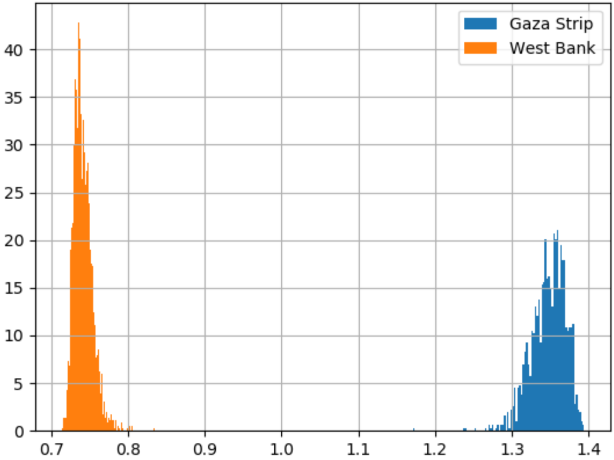

Total exp. changes by 0.30 log points between GS & WB

Poverty line

- As in Palestinian Central Bureau of Statistics and World Bank (2003) we use a single poverty line for both the West Bank and Gaza. Difference in log-price levels of 0.34 (0.30); equivalent to a 0.34 difference in $\log \lambda$.
- (Palestinian Central Bureau of Statistics 2018) reports a 2011 poverty rate for Palestine of 25.3%.
- Translate the distribution of $\log \lambda$ so that (population-weighted) 25.3% of all households have a value of $\log \lambda$ greater than zero.
- With this translation, "poverty line" in $\log \lambda$ always zero!

Distributions of translated log λ (basic basket)

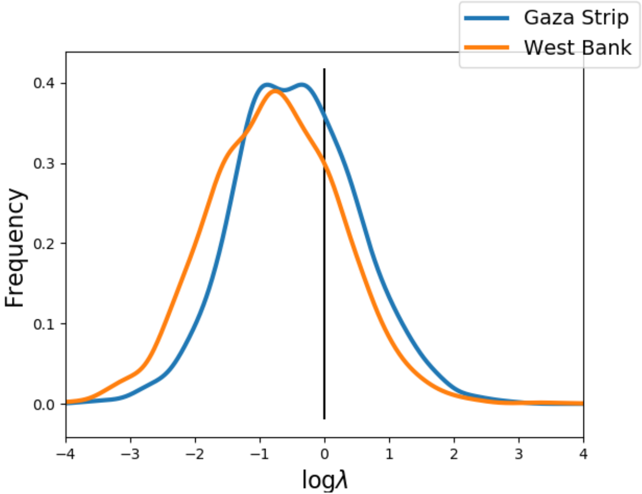

Distributions of translated log λ (25 goods)

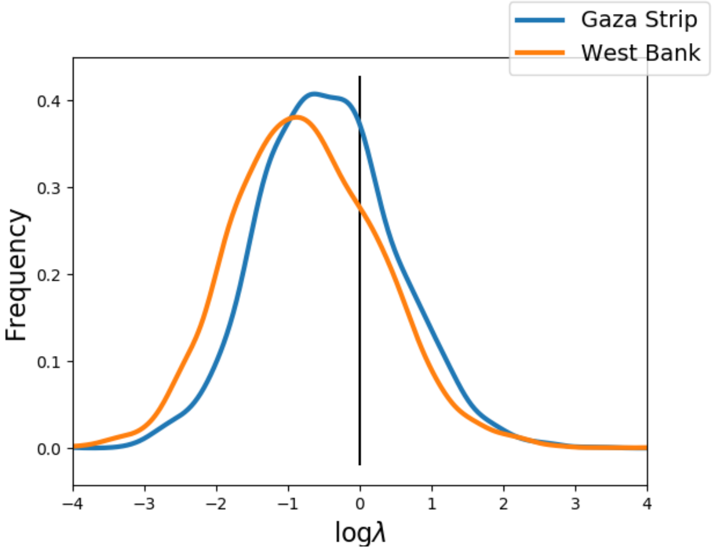

Poverty results

Table: Population-weighted estimates of 2011 Head-count poverty, by region

Region	155 Goods	25 goods	PCBS
Gaza Strip	32.7	30.7	38.9
West Bank	21.4	21.9	17.6
Pooled	25.3	25.3	25.7

Conclusions

Using demand system

1. If we observe even a *small* set of expenditures for different goods across households we can get good estimates of relative welfare.
2. We also get an optimal price index, allowing comparisons across time or places.
3. Natural interpretation of $\log \lambda$ as social return to transfer.

Conclusions

Experiments

1. Using a subset of goods chosen because they have large expenditure shares dramatically underestimates poverty rates (by overestimating expenditures).
2. Asking about all expenditure items, but only asking a given household only a random subset also performs poorly. It becomes very difficult to compare across households, and exacerbates missing data problems.

Conclusions

Poverty in Palestine, 2011

1. For the 2011 PECS data we can go from over 700 goods to 25, with almost no reduction in the accuracy of our estimates.
2. Can estimate poverty separately in West Bank & Gaza.
3. Both reduced (25 goods) and basic basket (155 goods) yield very similar estimated poverty rates.

Recommendations

Two main innovations:

1. Focus on estimating the latent wealth index $\log \lambda$ instead of total expenditures. Even if ultimately interested in total expenditures:
   - Gives a superior way to deal with differences in household size;
   - Convenient & efficient way to impute missing item expenditures.
   - Allows for calculation of theory consistent price index.
2. Use reduced expenditure basket to *frequently* measure expenditures. In case of 2011 PECS, entire survey could be reduced from 58 pages to 19 pages, with corresponding reductions in enumerator time.

Further Practical Suggestions & Experiments

- In the first implementation experiment: randomly split sample into two groups, one of which would receive the full survey; the second of which would receive a 25-good version of the survey. Use this to understand how different approaches affect response rates & time.
- Demand behavior may change over time; a best practice would be to field a short survey frequently, and after a few years to re-validate the approach by fielding a survey with a full expenditure module.

Further Practical Suggestions & Experiments

- With a small expenditure module alternative approaches to data collection become feasible. For example, imagine recruiting individuals with mobile smart phones to a panel, then eliciting expenditures on 25 goods every month or every quarter. Near ‘real-time’ measurement of household welfare!
- Harmonizing expenditure items with the UN “Classification of Individual Consumption According to Purpose” (COICOP) would make international comparisons much easier!
- An alternative elicitation might exploit the COICOP to start by asking about broad categories (“Have you purchased any fish or other seafood during the last week?”) before drilling down to less aggregate categories. This is a method that might be well-suited to CAPI or smart-phone data collection approaches.

Christiaensen, Luc, Ethan Ligon, and Thomas Pave Sohneson. 2017. Using reduced consumption aggregates to track and analyze poverty. Manuscript.

Filmer, Deon, and Lant H. Pritchett. 2001. Wealth effects without expenditure data—or tears: an application to educational enrollments in states of India. *Demography* 38 (1): 115–132.

Khawaja, Marwan. 1998. Poverty assessment in Palestine. Unpublished Ms.

Palestinian Central Bureau of Statistics. 2018. *Main findings of living standards in Palestine (expenditure, consumption and poverty), 2017.* Technical report. Palestinian Central Bureau of Statistics. `http://www.pcbs.gov.ps/Downloads/book2368.pdf?date=7_5_2018`.

Palestinian Central Bureau of Statistics and World Bank. 2003. *Poverty in the West Bank and Gaza after three years of economic crisis.* Technical report. Palestinian Central Bureau of Statistics and World Bank.