The Well-Being of Nations: Estimating Welfare from International Migration

Sanghoon Lee* Seung Hoon Lee† Jeffrey Lin‡

September 28, 2018

Preliminary and incomplete

(Most recent version)

Abstract

The limitations of GDP as a measure of welfare are well known. We propose a new method of estimating the welfare of countries using bilateral international migration flows. We estimate a discrete choice model in which everyone in the world chooses a country in which to live. This recovers an estimate of each country's overall quality of life. Our estimates, by relying on revealed preference, improve upon previous estimates of economic well-being that consider only income or a small number of factors, or rely on structural assumptions about how these factors contribute to well-being. We can also allow for heterogeneous preferences by separately estimating country quality of life by education, gender and age.

Keywords: International migration, quality of life

JEL classification: E21, D63, I31, F22, J61

*Sauder School of Business, University of British Columbia, sanghoon.lee@sauder.ubc.ca.

†School of Economics, Georgia Institute of Technology, seunghoon@econ.gatech.edu.

‡Research Department, Federal Reserve Bank of Philadelphia, jeff.lin@phil.frb.org.

Disclaimer: This working paper represents preliminary research that is being circulated for discussion purposes. The views expressed here are solely those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

1 Introduction

A central question in economics is the measurement of well-being across countries. The limitations of GDP as a measure of welfare are well known. There have been many attempts to provide a measure that can better capture other factors omitted or inadequately measured by GDP, including leisure, environmental quality, inequality, and political freedom. Many studies incorporate measured factors into statistical or structural models to estimate country well-being (Jones and Klenow, 2016). Others use survey data that asks respondents about their subjective well-being (Easterlin, 1974).

In this paper, we propose a new method of estimating welfare across countries based on the revealed preference of actual and potential international migrants. We combine gross bilateral migration flows across countries with a discrete-choice model to recover estimates of each country's quality of life. The key idea is that people tend to move from a low-utility country to a high-utility one.

Our estimates, by relying on revealed preference, improve upon previous estimates of economic well-being that consider only income or a small number of factors, or rely on structural assumptions about how these factors contribute to well-being. We also (plan to) allow for heterogeneous preferences separately estimating country quality of life by gender, education, and age. Unlike Jones and Klenow (2016), our approach relies on a discrete-choice framework. It also allows for the incorporation of an arbitrary number of factors, such as climate, other natural amenities, public goods, political climate, personal and economic freedom, etc., and the estimation of the contribution of these factors to overall quality of life. We can also perform counterfactuals on how various factors may affect the migration choices for different populations.

We deal with a number of important challenges, including measurement error in gross population flows and unobserved migration restrictions and policy and welfare factors. Under certain conditions, our estimated welfare rankings are invariant to the presence of unobserved factors. Further, our approach is robust to the many country pairs that have zero observed gross population flows.

1.1 Related work

We contribute to several literatures. First, a central question in economics is the comparative welfare of countries. Much of this literature has emphasized that GDP is a flawed measure of welfare. A number of methods have been proposed to estimate country well-being including principal component analysis of a large vector of factors (Ram, 1982; Slottje, 1991) and surveys of subjective well-being (Easterlin, 1974). Recently, Jones and Klenow (2016) proposed a method of estimating well-being using household microdata on consumption and leisure and a calibrated utility model. Compared with previous work, we propose a new method based on the revealed preference of actual and potential international migrants.

Second, there is a large literature in urban and regional economics that estimates variation in quality of life within a country (Roback, 1982; Albouy, 2008; Diamond, 2016). We estimate quality of life for countries as a whole. Our approach also differs in that we do not assume spatial equilibrium as is customary for within-country analyses. As a result, our measures of welfare or the quality of life incorporate the roles of housing costs and income.

Third, a large literature tries to understand the determinants of migration flows (Grogger and Hanson, 2011; Lewer et al., 2013). Much of this literature emphasizes migration as a human capital investment and that migrants respond to labor market opportunities (Bodvarssen et al., 2015). Instead, our estimates emphasize that migrants may also be responding to other factors, including amenities, consumption, and political freedom.

2 Model and estimation

There are J countries of varying size, with populations $\{N_j\}$. Each individual i living in country $o \in J$ maximizes utility U by choosing a country d, including their present one, to which to migrate.

$$U_{od}^i = u_d - c_{od} + \epsilon_{od}^i \tag{1}$$

Destination country d offers utility $u_d \in \mathbb{R}$ to its residents. Assume $u_d \equiv Z_d'\alpha$ is a linear combination of various factors Z_d (e.g., income, leisure, climate, school quality). Moving costs between origin country o and destination country d are described by $c_{od} \equiv X_{od}'\beta$, where $X_{o,d}$ is a vector consisting of characteristics of the origin-destination country pair d and o (e.g., the distance between country d and o, country o's income, country d's immigration policy towards residents of o, etc.). We normalize c_{od} so that $c_{od} = 0$ if $d = o$.

Finally, ϵ_{od}^i is a person-level idiosyncratic shock. We assume that ϵ_{od}^i follows a Gumbel distribution with location parameter $\ln N_d$ and shape parameter 1. This differs from the standard conditional logit model that uses the standard Gumbel distribution with a location parameter 0. This departure is to account for population size differences across destination countries. A larger destination country tends to draw higher ϵ_{od}^i and thus a worker is more likely to choose a larger country. This assumption helps rationalize larger gross flows in and out of larger countries. One way to think of this is as modeling the number of potential destinations within country d: countries with more current residents offer proportionally more opportunities for potential migrants.

More precisely, this setting is equivalent to the one where for a potential destination country d, the worker receives multiple utility draws following the standard Gumbel distribution, with the number of draws equal to its population size N_i. (See Head and Ries, 2008, or Bury, 1999.) For example, imagine a world with two identical countries except their population sizes, where the deterministic component of utility, $u_d - c_{od}$, is 0 for each country, country 1 has 1 million residents, and country 2 has 2 million residents. This means that each person in the world compares 1 million random draws following the standard Gumbel distribution for country 1 and 2 million draws for country 2. This makes the person twice as likely to choose country 1 than country 2. The resulting population size distribution following this structure would be very close to the original population sizes of 1 and 2 million. On the other hand, if we assume that each person draws a draw following the standard Gumbel distribution for each country, as in the standard model, the resulting country sizes would be nearly equal, at 1.5 million each. Thus, our model offers the attractive property that, if moving costs exactly offset utility differences across countries, then migration choices do not affect country size.

At a first glance, the model resembles a conditional logit model to be estimated separately for each origin country. However, our model differs in that the utilities u_d of alternative destination countries are common across *all* origin countries. Instead, the best way to interpret our model is to consider our model as a conditional logit model on a global scale, where workers in different origin countries face different moving costs to a same destination countries.

A property commonly discussed for a discrete choice model is whether the model prediction is restricted by the Independence of Irrelevant Alternatives (IIA). But the IIA is relatively not so important here because we are not interested in estimating the choices when a new country emerges or an existing country disappears.[1]

[1] In our model, the IIA is binding at each origin country but not at the global level. Suppose that the world consists of three countries A, B, and C of the same population sizes. Countries B and C are identical to each other. Assume that the model has only countries A and B where country A workers choose A and B with 66.6% and 33.3% probabilities (i.e., 2 to 1 ratio) and country B workers choose A and B with 33.3% and 66.6% probabilities. Now suppose that we add country C to the choice set. Country A workers choose A, B, and C with 50%, 25%, and 25%. Country B workers choose A, B, and C with 20%, 40%, and 40%. Note that the IIA holds at an origin country level; the probability ratios of choosing countries A and B remain 2 to 1 for country A or 1 to 2 for country B. However, at a global level the percentage of people choosing countries A and B change from 1 to 1 (33.3%+66.6% vs. 66.6% + 33.3%) to 1 to 1 to 0.92 (33.3 + 50%).

2.1 Estimation

Following these assumptions on utility, the probability that a country o resident chooses to migrate to country d is:

$$\pi_{o,d} = \frac{N_d \exp(u_d - c_{o,d} + \nu_{o,d})}{\sum_d N_d \exp(u_d - c_{d,j} + \nu_{d,j})} \tag{2}$$
$$= \frac{\exp(\ln N_d + u_d - X'_{od}\beta + \nu_{o,d})}{\sum_d \exp(\ln N_d + u_d - X'_{d,j}\beta + \nu_{d,j})} \tag{3}$$
$$= \frac{\exp\{(\ln N_d + Z'_d\alpha + \nu_{o,d}) - X'_{od}\beta\}}{\sum_d \exp\{(\ln N_d + Z'_d\alpha + \nu_{d,j}) - X'_{d,j}\beta\}} \tag{4}$$

where equation 4 substitutes our parameterizations of utility and migration costs. Note that in the first line we use our assumption on the individual error term that random draws for each destination country is equal (or equivalently, proportional) to its population size.

Equation 3 describes McFadden's (1974) alternative-specific conditional logit model. We estimate equations 3 (and later, 4) using individual-alternative data based on all global flows (versus separate estimations using data for each origin country). In other words, each person in the world considers the same set of destination countries, but face different factors X_{od} depending on their origin country, the destination alternative, and the origin-destination pair.

An issue that often arises in estimating discrete choice models is that some choices may have zero shares. Gross international migration flow data also feature many origin-destination pairs with zero observed migrants. However, by using global flows, we avoid identification issues stemming from zero shares. For example, suppose there are zero observed gross flows from the U.S. to Vanuatu. The estimate for Vanuatu's utility is identified as long as there are non-zero flows to Vanuatu from other origin countries, *including* Vanuatu residents who choose to stay in Vanuatu.

Note that the structure of the model implies an offset N_d that enters the estimation with coefficient 1.

There are two ways to estimate our model. We can directly estimate all parameters in one step using equation 4, or we can use a two step estimation strategy where we first estimate fixed effects $\hat{u}_d$ for each destination country using equation 3, and then regress it against Z_d to estimate coefficients $\hat{\alpha}$, weighting $\hat{u}_d$ by the inverse of variance estimated in step 1.

We use the two step estimation strategy for two reasons. First, as pointed out in Wooldridge (2003), the two step method avoids omitted variables bias in estimating c_{od}.

Not including the fixed effects u_d may bias estimates of c_{od} if the idiosyncratic error ν_{od} includes an destination-specific factor ξ_d. Without the fixed effects, omitted variables can bias estimates of the c_{od}, or, in the parameterized version of the model, the coefficients on country-pair characteristics β (Moulton, 1990). Second, the two step method provides two sets of welfare measures: $\hat{u}_d$ and $\hat{\hat{u}}_d \equiv Z_d'\hat{\alpha}$. Each measure has its pros and cons. $\hat{u}_d$ encompasses more welfare components including unobservable ones, but is more sensitive to a temporary migration shocks. It also may be more sensitive to unobserved bilateral factors, as described in detail later. $Z_d'\hat{\alpha}$ is less susceptible to these issues, but it only encompasses observable welfare components.

3 Data

We use estimates of bilateral flows between 179 countries from 2005 to 2010 from Abel and Sander (2014). They impute global bilateral migration flows using sequential tabular data on the stock of immigrants by origin and destination country from 1990, 1995, 2000, 2005, and 2010. These stock data are primarily based on place-of-birth responses to censuses. Thus, successive stock tables report the number of people for every country of residence–country of birth pair, in two years t and $t+1$.

Abel and Sander (2014) then impute bilateral flows that are consistent with observed stock tables. (They also account for changes in immigrant stocks from data on births and deaths and refugee movements.) They set the number of stayers in each country to the maximum possible value—thus, if 1 million people are observed in t as having been born in, and residing in, country A, and 0.9 million such people are observed in $t+1$, then (abstracting from natural increase or decrease) Abel and Sander assume that 0.9 million stayed in country A between t and $t+1$. Thus, the remaining flows represent the minimum number of gross flows required to rationalize the evolution of migrant stocks. We investigate the robustness of our results to the Abel and Sander (2014) methodology.

We use data on bilateral country characteristics from the GeoDist database from CEPII (Mayer and Zignago 2006). These data describe for each country pair the presence of a shared border, any shared languages, any past or present colonial relationship, or a number of distance measures. They are standard measures for transportation costs of physical products in international trade, adopted by numerous papers (e.g. Bernard, Redding, and Schott 2011), and also used to measure the moving costs of migration flows across countries (e.g. Beine, Docquier, and Özden 2009.) One measure of moving costs that we include is membership in the Schengen agreement, which allows free mobility between member countries. In 1985, this agreement was initially signed by five countries: Belgium, the Netherlands, Luxembourg, France, and Germany. By 2008, the Schengen Area had expanded to inclue 25 European countries.

Following Jones and Klenow (2016), we also use data on country income and consumption from the Penn World Tables version 9.0 (Feenstra, Inklaar, and Timmer 2015). Other data on country factors that may contribute to welfare—such as inequality, government expenditures, leisure time, and air quality, are drawn from a variety of sources. These are described in the appendix.

Table 1 provides summary statistics for bilateral factors for (176^2 =) 32,041 origin-destination pairs. We report means and standard deviations for bilateral factors conditioned on the origin and destination country being different ($1_{Diff} = 1$).

Table 2 provides summary statistics for country-specific factors for 176 countries. The four factors reported are log GDP per capita; the Gini coefficient of household income; air quality (PM25, or the concentration of atmospheric particulate matter less than 2.5 micrometers); and public health care expenditure as a percent of GDP.

Table 1: Summary statistics for country pairs

	mean/sd
1_{Diff}	0.994
	(0.075)
$1_{Diff}\times$ ln(dist)	8.703
	(1.010)
$1_{Diff}\times$ Sharing Border	0.017
	(0.130)
$1_{Diff}\times$ Common Language	0.148
	(0.355)
$1_{Diff}\times$ Schengen Agreement	0.019
	(0.136)
$1_{Diff}\times$ Colonial Link	0.011
	(0.106)
N of Country Pairs	32041

Table 2: Summary statistics for (destination) countries

	mean/sd
Log(GDPc)	8.039
	(1.624)
Gini	39.918
	(8.218)
PM25	28.862
	(19.346)
Pub. Exp. Share	9.585
	(7.691)
N of Countries	176

4 Results

4.1 First-stage estimation

We first present results of estimating equation 3 by maximum likelihood. The first step recovers estimates of β, the coefficients on bilateral factors measuring moving costs, and u_d. Table 3 presents estimates for a model including 6 country-pair factors: (1) whether the destination country is the same as the origin country, i.e., a choice to stay (1_{Diff}); (2) the log of the distance between the pair; (3) whether the pair are signatories of the Schengen agreement; and (4) whether the pair share a border; (5) whether the pair share a common language; (6) whether the pair share a (past or present) colonial relationship. Factors (2)–(6) enter as interactions with the different-country indicator.

Estimated standard errors are reported in parentheses. They are clustered by origin country, to allow for within-country correlation in destination choice.

The signs of the coefficients are as expected, if imprecisely estimated. Country pairs that share a border, a language, or a colonial link, or that are both members of Schengen, have higher migration flows. Countries that are more distant have lower migration flows. Same-country gross flows are significantly larger compared with different-country gross flows.

Note the large number of observations reported in the first-step regression. The unit of observation is each potential destination (176 countries) for each person in the world

(6.3 billion), yielding a sample size of (176 $\times$ 6.3 billion =) 1.14 trillion. Observations vary in factors X_{od} depending on their origin country, the destination alternative, and the origin-destination pair.

Table 3: Country-pair factors predict migration flows

	Set 1	Set 2	Set 3	Set 4	Set 5
1_{Diff}	-3.144	0.607	0.804	-2.786	-3.231
	(3.098)	(2.026)	(2.216)	(3.361)	(3.264)
$1_{Diff}\times$ ln(dist)	-0.983^b	-1.370^c	-1.391^c	-0.997^b	-0.971^b
	(0.418)	(0.302)	(0.316)	(0.451)	(0.444)
$1_{Diff}\times$ Schengen Agreement	-0.342		-0.390	-0.366	-0.277
	(1.063)		(0.894)	(1.048)	(1.166)
$1_{Diff}\times$ Sharing Border	1.513^a			1.699^a	1.568
	(0.913)			(1.031)	(0.963)
$1_{Diff}\times$ Common Language	0.693				0.884
	(0.643)				(0.638)
$1_{Diff}\times$ Colonial Link	1.424				
	(0.969)				
N	1.14e+12	1.14e+12	1.14e+12	1.14e+12	1.14e+12

Notes: Standard errors clustered by origin country. ^{a}—$p < 0.10$; ^{b}—$p < 0.05$; ^{c}—$p < 0.01$.

Table 4 shows the largest 36 countries by population ranked by $\hat{u}_d$, as recovered by the first-step regression. (Smaller countries are omitted for presentation purposes.) Canada comes out on top, followed by Italy, Spain, the U.S., and South Africa. Among the largest countries, Pakistan, Indonesia, China, India, and Bangladesh lag.

Compared with GDP per capita, U.S. welfare is lower compared with other countries according to our measure. Countries in south and east Asia (China, India, Thailand, the Philippines, Japan, South Korea) have lower relative welfare according to our measure compared with per capita GDP. Russia, Iran, Mexico, and Brazil are also lower by our ranking. On the other hand, east African countries like Egypt, Kenya, Tanzania improve markedly in relative welfare according to our measure.

Figure 1 displays a scatter plot of first-stage estimates of country welfare u_d against GDP per capita in 2005, for the 36 largest countries only. The fitted line illustrates that the two measures are positively correlated, though there is dispersion around the fitted line.

Table 4: Welfare ranks for largest countries

	Stage 1	Stage 2	Jones and Klenow	GDP per capita
1	Canada	France	United States	United States
2	Italy	Germany	France	United Kingdom
3	Spain	Japan	United Kingdom	Japan
4	United States	Italy	Japan	Canada
5	South Africa	United Kingdom	Canada	France
6	United Kingdom	Spain	Italy	Germany
7	Germany	Canada	Spain	Italy
8	France	United States	Germany	Spain
9	Poland	South Korea	South Korea	South Korea
10	Russia	Poland	Poland	Poland
11	Ukraine	Turkey	Mexico	Mexico
12	Thailand	Ukraine	Turkey	Turkey
13	Algeria	Russia	Argentina	Russia
14	Japan	Iran	Russia	South Africa
15	Turkey	Mexico	Iran	Argentina
16	Kenya	Thailand	Ukraine	Brazil
17	South Korea	Argentina	Brazil	Colombia
18	Tanzania	Morocco	Thailand	Iran
19	Iran	Egypt	Colombia	Algeria
20	Colombia	China	Egypt	Thailand
21	Egypt	Pakistan	China	Morocco
22	Ethiopia	Brazil	Indonesia	Ukraine
23	Sudan	Colombia	Morocco	China
24	Nigeria	Vietnam	Philippines	Indonesia
25	Argentina	Tanzania	South Africa	Philippines
26	Morocco	Philippines	Pakistan	Egypt
27	Myanmar	Ethiopia	Vietnam	Nigeria
28	Vietnam	South Africa	India	Pakistan
29	Brazil	India	Sudan	India
30	Philippines	Kenya	Bangladesh	Vietnam
31	Mexico	Bangladesh	Nigeria	Sudan
32	Pakistan	Myanmar	Kenya	Kenya
33	Indonesia	Sudan	Tanzania	Bangladesh
34	China	Nigeria	Ethiopia	Tanzania
35	India	Indonesia	Algeria	Myanmar
36	Bangladesh	Algeria	Myanmar	Ethiopia

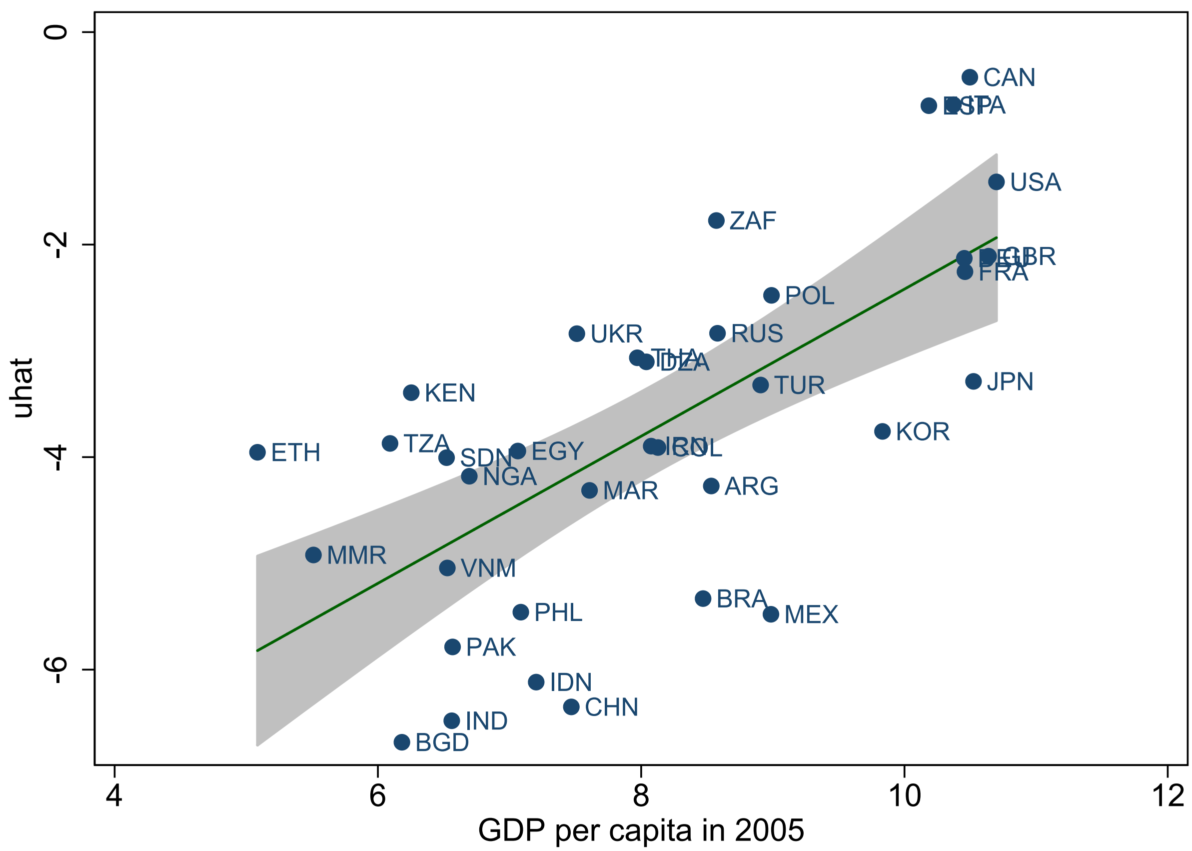

Figure 1: $\hat{u}_d$ versus GDP per capita for largest countries

4.2 Identification

Note that conditioned on measuring moving costs c_{od} precisely, we can recover consistent estimates of country welfare u_d. But, if there are omitted bilateral factors, then our estimates of u_d may be biased. For example, stricter immigration policies will reduce inflows, reducing estimated welfare. Alternatively, strict *emigration* controls will increase stayers, increasing estimated welfare. Data is a challenge here—there is no comprehensive database on all migration policies in the world. Even if we could measure migration policy exactly, there are many non-policy migration costs that are difficult to measure, e.g., the social costs of leaving family and friend networks may vary by sending or destination country.

To address these concerns, we pursue two complementary strategies. One, we project $\hat{u}_d$ onto a vector of country factors Z_i (e.g., per capita GDP, air quality, inequality) and use its projected value $\hat{\hat{u}}_d$ as a second welfare measure. For example, North Korea has strict emigration controls. The lack of outmigration inflates $\hat{u}_i$, as people are unable to leave — the model interprets this is as a choice due to preferences rather than a consequence of constraints. But the projected estimate $\hat{\hat{u}}_d$ reduces this bias as North Korea scores poorly on a number of dimensions, including per capita GDP. Two, we use better measures of migration policy measured at the destination country level. By directly controlling for migration, we hope to reduce bias due to omitted migration policy.

First, we parameterize country welfare $u_d = Z_d'\alpha$ as a function of many observable factors in vector Z_i as in equation 4. Equation 4 still potentially omits many relevant factors. However, the resulting *predicted* values $\hat{u}_d$ may be unbiased estimates of the welfare of countries if omitted factors are related to the observed factors in the same way.

Note that our interest is in predicting country welfare u_d rather than doing counterfactual exercises with respect to country level attributes. Thus, biases from reverse causation do not pose as much concern. To obtain better predicted values, we plan to include additional factors and use LASSO to select relevant variables.

Table 5 reports estimates from a second-step regression of u_d on several destination country factors Z_d. We include controls for (1) the logarithm of per-capita GDP; (2) the Gini index; (3) the concentration of particulate matter (PM) with diameters of 2.5 micrometers and less, a measure of air pollution; and (4) public expenditure share on health.

Second, we also include several measures of destination-level immigration policy. These data are from the UN World Population Policies Database 2015. The survey asks member and non-member states about government policies with respect to population. Of interest are a series of questions characterizing government policy on the integration of non-nations, integration of immigration, naturalization, irregular immigration, return migration, dual citizenship, and diaspora. The potential responses are typically yes/no (e.g., does the government have a policy on this?) or offer a choice of responses (e.g., penalties for employers of irregular migrants versus fines, detention, or deportation of irregular migrants versus regularization of legal status). We encode these responses flexibly as a vector of indicator variables.

Columns (1) and (2) include log per capita GDP and air quality as regressors. Columns (3) and (4) add the Gini index and public spending on health. (The latter two factors are only observed for 128 countries.) Further, columns (2) and (4) add the vector of controls for destination country immigration policy.

The coefficients are precisely estimated and in accord with previous results. Welfare is increasing in income and decreasing in inequality. Air quality and public expenditure on health do not appear to be major factors in welfare. Controlling for destination immigration policy has little effect on the estimates. This suggests that omitted factors may not play an important role in estimating either $\hat{u}_d$ or $\hat{\hat{u}}_d$.

Table 5: Destination-country factors predict welfare

	without Policy	with Policy	without Policy	with Policy
Log(GDPc)	0.584^c	0.654^c	0.406^c	0.432^c
	(0.068)	(0.106)	(0.102)	(0.146)
PM25	0.006	-0.010	-0.003	-0.013
	(0.006)	(0.006)	(0.007)	(0.008)
Gini			-0.041^b	-0.042^b
			(0.017)	(0.019)
Pub. Exp. Share			0.005	-0.009
			(0.025)	(0.030)
Constant	-7.341^c	-6.823^c	-4.212^c	-1.751
	(0.637)	(1.815)	(1.107)	(2.635)
R-squared	0.318	0.595	0.330	0.647
N	173	172	128	128

One may wonder the effect on the results of using only a subset of countries in our estimation. Our sample shrinks as we include more country level factors. To explore this issue, we analyze changes in the relative ranking of country welfare by $\hat{u}_d$ when we select countries based on population size.

Figure 2 shows rankings as we narrow the sample from countries larger than 10 million to countries larger than 14 million, in 1-million-population increments. Lines connect country labels across the five subsamples shown in columns. Note the high degree of correlation between estimates across subsamples. In fact, rankings are highly correlated for all subsamples of countries based on size in 1-million-population increments up to 30 million+. The correlation coefficients between rankings across subsamples are always higher than 0.89.

The strong correlation across subsamples tests whether IIA holds. Under IIA, the ratio of probabilities for any two alternatives is the same, whether or not other alternatives are available. If IIA holds in reality, then the coefficient estimates obtained from a subsample of countries (and hence the ranking of countries) will not be significantly different from those obtained from the full sample of countries. (This is the spirit of the McFadden and Hausman (1984) test of IIA.) Thus, the strong correlation between rankings obtained on various subsamples is consistent with IIA holding in this setting.

4.3 Comparison with Jones and Klenow

Jones and Klenow (2016) estimate the welfare of countries using a calibrated utility model and data from household surveys in all countries on four factors: consumption, leisure, mortality, and inequality. Earlier, Table 4 showed a comparison of welfare rankings by the Jones-Klenow estimates for the largest countries by population.

Figure 3 shows correlation coefficients among four measures of county welfare: (1) GDP per capita; (2) the Jones-Klenow estimates; (3) our first-stage estimates of $\hat{u}_d$; and (4) our second-stage estimates of $\hat{\hat{u}}_d \equiv Z_d'\hat{\alpha}$.

All four measures of strongly positively correlated. Our second-stage estimates are highly correlated with the Jones-Klenow estimates and GDP per capita. This is unsurprising given that consumption is an important input into their measure. The divergence between our first-stage estimates of country welfare $\hat{u}_d$ and GDP per capita suggests that other factors not accounted for by the JK framework may also contribute to welfare. In future work, we plan to refine the list of country factors used to predict $\hat{\hat{u}}_d$ to further understand the divergence between $\hat{u}_d$ and income.

4.4 Future work

1. Sensitivity to Abel and Sander method.

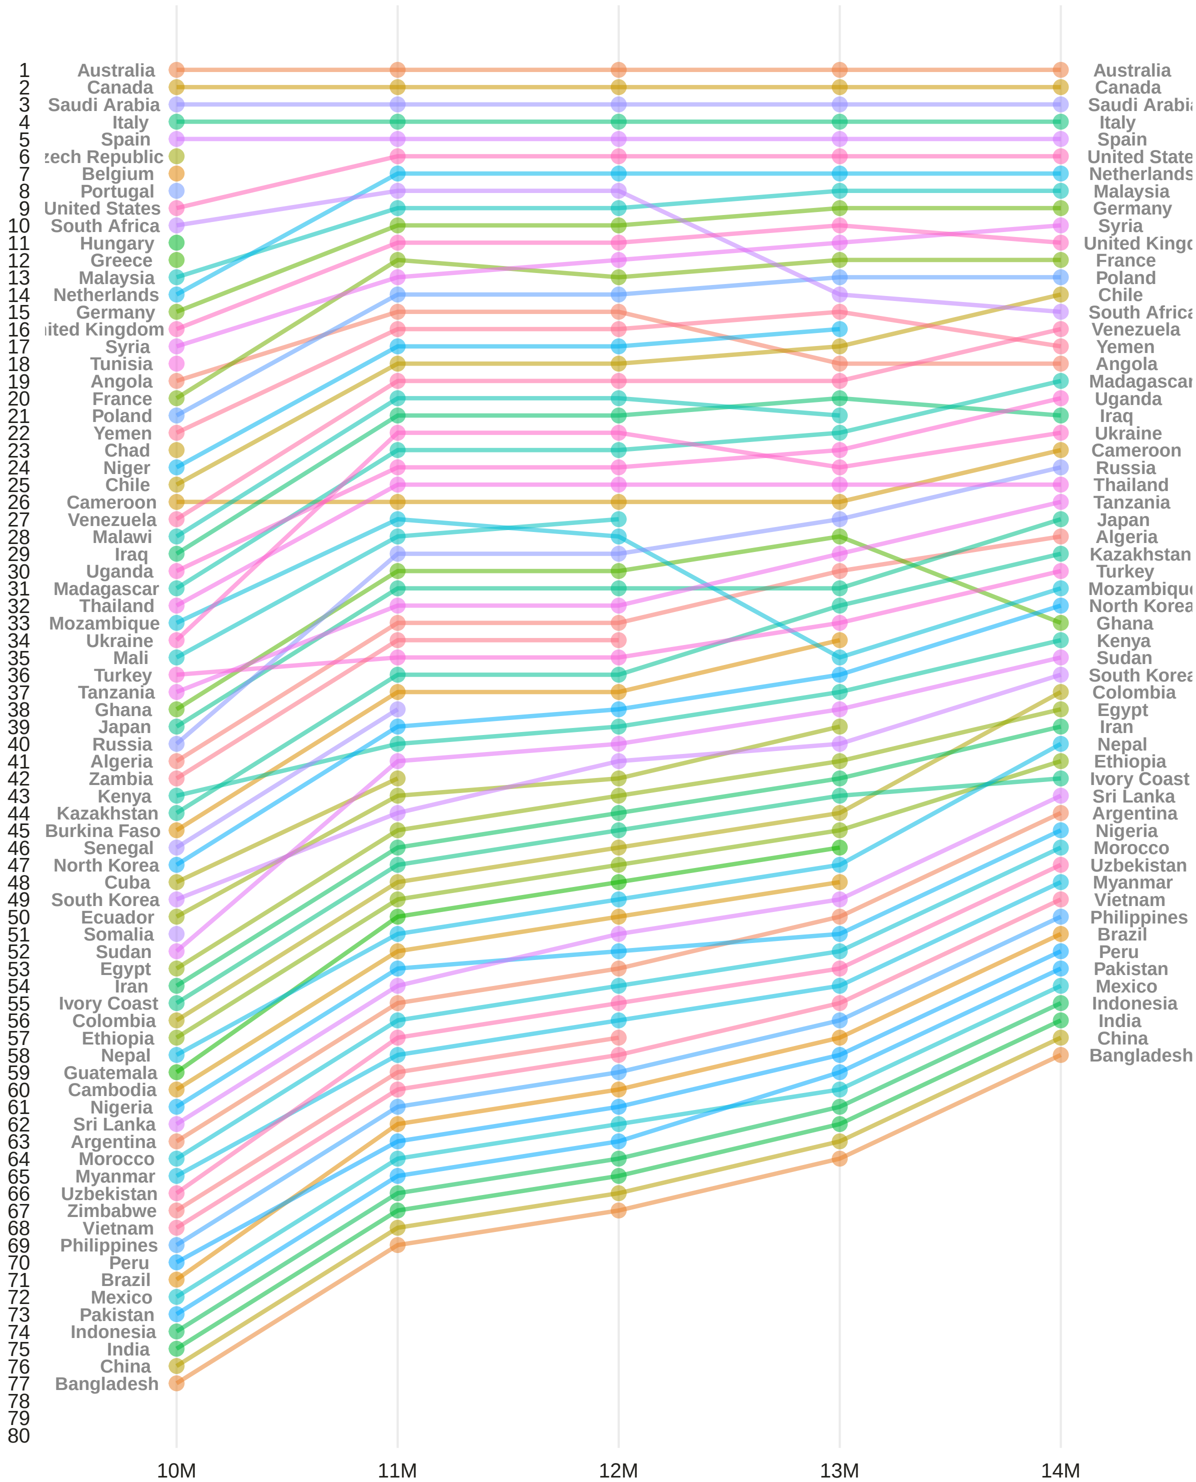

Figure 2: Country rankings by $\hat{u}_d$ are stable for subsamples

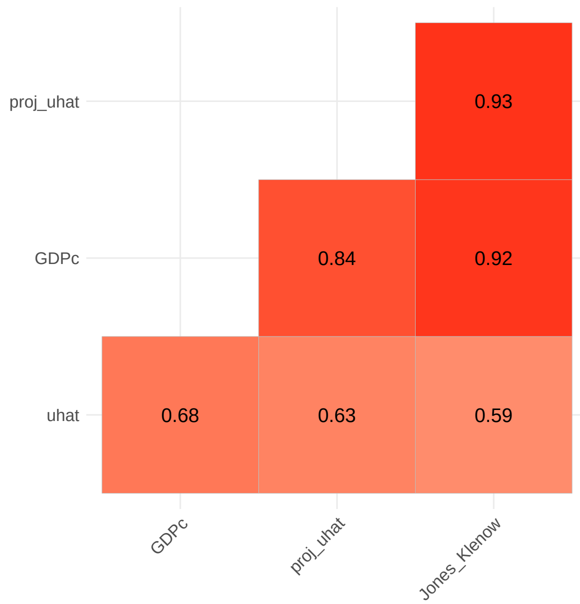

Figure 3: Correlation of welfare measures

2. Estimates by education, gender, and age.

5 Conclusion

This paper proposes a new method of estimating the welfare of countries based on international migration patterns. The key idea is that people tend to move from low-utility places to high-utility ones.

6 Works cited

Albouy, David (2008). "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas." NBER Working Paper 14472.

Beine, Michael, Frédéric Docquier, and Çağlar Özden (2011). "Diaspora." *Journal of Development Economics* 95: 30-41

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott (2011). "Multiproduct Firms and Trade Liberalization." *The Quarterly Journal of Economics*, Volume 126, Issue 3, Pages 12711318

Bodvarsson, Örn B., Nicole B. Simpson, and Chad Sparber (2015). “Migration Theory.” In *Handbook of the Economics of International Migration*, edited by Barry R. Chiswick and Paul W. Miller, 3–51. Amsterdam: North-Holland.

Diamond, R. (2016). “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980–2000,” *American Economic Review* 106(3): 479–524.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. “The Next Generation of the Penn World Table.” *American Economic Review*, 105 (10): 315082.

Easterlin, Richard (1974). “Does Economic Growth Improve the Human Lot? Some Empirical Evidence.” In *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, edited by Paul A. David and Melvin W. Reder, 89–125. New York: Academic Press.

Grogger, Jeffrey and Gordon H. Hanson (2011). “Income maximization and the selection and sorting of international migrants.” *Journal of Development Economics* 95: 42–57.

Hausman, J. and D. McFadden (1984). “Specification tests for the multinomial logit model,” *Econometrica* 52, 1219–1240.

Jones, Charles I., and Peter J. Klenow (2016). “Beyond GDP? Welfare across Countries and Time.” *American Economic Review*, 106 (9): 2426–57.

Lewer, Joshua J., Gail Pacheco, and Stephanié Rossouw (2009). “Do Non-Economic Quality of Life Factors Drive Immigration?” IZA Discussion Paper 4385.

Mayer, Thierry and Soledad Zignago (2006), “Notes on CEPII’s Distance Measures,” Working Paper, *CEPII*.

Ram, Rati (1982). “Composite Indices of Physical Quality of Life, Basic Needs Fulfillment, and Income.” *Journal of Development Economics*, 11: 227–47.

Roback, Jennifer (1982). “Wages, Rents, and the Quality of Life,” *Journal of Political Economy*, 90 (6): 1257–78.

Slottje, Daniel J. (1991). “Measuring the Quality of Life across Countries.” *The Review of Economics and Statistics*, 73 (4): 684–93.