International Sovereign Spread Differences and the Poverty of Nations*

Farzana Alamgir
McMaster University.

Alok Johri
McMaster University.

October 2022

Abstract

National poverty head-count ratios and country default risk are positively correlated and robust to controlling for variables including debt, Gini coefficient and per capita GDP. A sovereign default model with poor and non-poor households as well as a social safety net which taxes the non-poor and transfers to the poor delivers this correlation. A political constraint ensuring participation constrains the fiscal choices of the government. The model is calibrated using South African data on household income dynamics, its poverty line and aggregate social transfer rate. An economy with more poor households displays higher default risk than the Benchmark economy. The interaction of international borrowing terms with the social safety net and with the political constraint accounts for this result. Political defaults occur when the non-poor find the safety net too weak to support and this occurs when too large a fraction of taxes will be needed for debt repayment.
JEL classification: F34, F41, G15, H63.
Keywords: Sovereign default, country spreads, poverty rates.

*For comments and suggestions, we thank (without incriminating) César Sosa-Padilla, Viktoria Hnatkovska, Mick Devereaux, Amartya Lahiri, Pau Pujolas, Gajendran Raveendranathan, Marc-Andre Letendre and participants at various seminars and the 2021 Canadian Macro Study Group and 2022 CEA meetings. Johri thanks the Social Sciences and Humanities Research Council of Canada for research funding through grant number 435-2016-0708.
Address: Dept.of Economics, McMaster University, 1280 Main St.West, Hamilton, Ontario L8S4M4, Canada.
E-mails: `alamgirf@mcmaster.ca`, and `johria@mcmaster.ca` (corresponding author).

1 Introduction

Countries vary widely in the number of citizens who live below the poverty line. For example, the proportion of the population that lives on less than \$1.90 per day is around 5 percent in Tunisia while it is about 25 percent in Ghana. Governments facing high rates of poverty may wish to provide a social safety net for these indigent citizens by providing transfers in order to prevent their consumption levels from getting too low. The public revenues of poor nations, however, are severely constrained by the local tax base and they may wish to use international debt markets in order to supplement their resources, especially when facing a recession. Their ability to access international resources depends on the terms at which credit is made available and these terms vary widely from nation to nation, even when averaged over decades. Figure 1 plots the sovereign spread on the government debt of Tunisia and Ghana for the period 2007-2015. As is evident from the graph,the sovereign spread for Tunisia was less than half of that of Ghana at any given year between 2007 and 2015.

In this paper, we begin with an empirical analysis to explore the relationship between the spread paid by a national government and the average poverty rate of that nation, as captured by the poverty head-count ratio. We find a positive correlation between poverty rates and sovereign spreads in international emerging economy data. Next, we build a sovereign default model of a small open economy in the tradition of Eaton and Gersovitz (1981) where a proportion of households are distinguished from the rest by having a much lower level of income than average. The income state of households is not permanent and households can transition into and out of poverty. The government implements a social safety net using a tax and transfer system while also borrowing on international markets using long-duration non-contingent bonds. These bonds are subject to default risk, so international lenders price the debt accordingly. The government maximizes utilitarian social welfare where the utility of the poor and the rest are weighted according to social welfare weights. The level of transfers offered by the social safety net are constrained by a political constraint which ensures that all households are willing to participate in the policy (which is tantamount to an insurance scheme). We calibrate the model parameters using available data from the South African economy on debt, sovereign spreads, real per capita income, the poverty head-count ratio

Figure 1: Sovereign spreads (2007-2015)

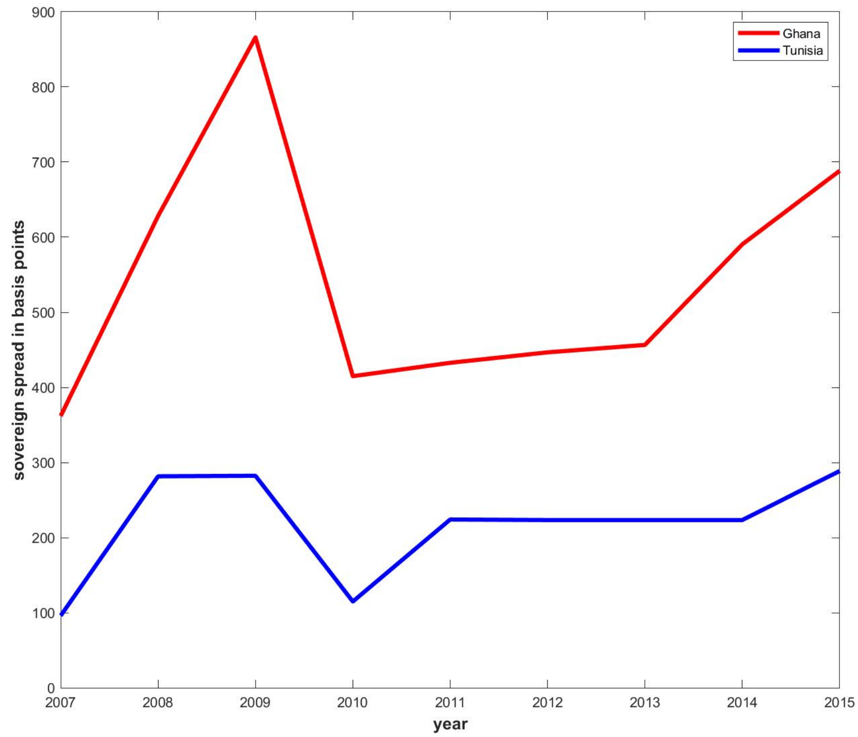

Note: Here spread refers to the difference between the yield of sovereign bonds of the respective country and the yield of US 10-year treasuries in basis points between the period 2007-2015.

and fiscal data. In addition, we use several waves of the National Income Dynamics Study to measure the proportion of households below the poverty line, and the transition rates of non poor into poverty.

In order to use our model to shed light on the positive correlation between poverty rates and spreads documented in the paper, we build a variant of the model – the High-Poverty economy that has a greater proportion of poor households than the Benchmark economy. We show that the High-Poverty economy faces worse borrowing terms and a higher default risk than the Benchmark economy. As a result, it is able to borrow less than the Benchmark economy and this lowers welfare considerably. Poor households have lower consumption levels in the High-Poverty economy than in the Benchmark economy despite households receiving the same endowment stream in both economies. This occurs because transfers to the poor are lower. We decompose these effects into those that can be attributed to the smaller number of households able to pay taxes and those that emerge from the worse credit terms available to the High-Poverty economy. If an international aid agency could replicate the spread-debt menu offered by private lenders to the Benchmark economy, not only would debt, welfare and consumption rise in the High-Poverty economy, defaults would be much less likely to occur.

Since a higher proportion of poor households imply a higher level of inequality in the High-Poverty economy, we also compare the default risk of the High-Poverty economy to an analogous economy with the same level of inequality as the Benchmark economy and show that the additional default risk is primarily driven by the increase in the number of poor households as opposed to the increase in inequality that this creates. We also show that controlling for the relatively lower GDP in the High-Poverty economy does not eliminate the results discussed above.

Why is the High-Poverty economy more likely to default? There are two factors that influence this. First note that the government of the High-Poverty economy would like to provide the same consumption stream to the poor as the government in the Benchmark economy (given the same endowment streams and social welfare weight of the poor) but has access to a smaller tax base due to the smaller number of non-poor households. Second, in the presence of the political constraint, the limited tax base can interact with high borrowing costs to create defaults that occur for political reasons. To understand this, note that non-poor households are willing to pay taxes as payments into a government run consumption insurance scheme whose effectiveness depends on the level of transfers promised by the government. If too large a fraction of tax receipts flow out of the economy as debt repayment, then the social safety net looks less attractive to the non-poor and they limit the level of taxes that the government can charge. In these situations, the government revenues in the High-Poverty economy take a double-hit – not only are the number of tax-payers smaller than in the Benchmark economy, the tax rate is also constrained politically. At low levels of debt, borrowing provides a means to keep transfers high relative to taxes but as the debt burden rises or as income falls, too large a fraction of tax receipts start to flow out of the social safety net towards debt repayment. As this dynamic worsens, non poor households will eventually be unwilling to participate in the tax-transfer scheme. Defaults at this juncture occur to keep the political constraint binding and we call these political defaults because the politically acceptable tax rate cannot support repayment of debt. We show that in our calibrated Benchmark economy, political defaults almost never occur while a large fraction of the higher default risk of the High-Poverty economy are due to political defaults. Of course both the Benchmark and High-Poverty economy also display standard defaults which occur when income is low and the marginal utility of the poor is high. Governments may then prefer to default rather than send precious resources abroad as debt repayment. Defaults in this situation are similar in spirit to the standard sovereign default models that follow Arellano (2008) and these are not the major source of the higher default risk of the High-Poverty economy. The High-Poverty economy is able to sustain a lower level of transfers relative to taxes and for any debt level, it hits that critical ratio when default becomes desirable at higher levels of income. Similarly, for any income level, the political constraint starts to bind at much lower levels of debt than in the Benchmark economy beyond which default is preferable.

Finally, we vary the proportion of poor households in the observed range in order to generate a number of artificial economies that differ from the Benchmark economy in only this aspect. Using simulations from these economies, an artificial panel dataset containing observations on poverty rates, debt to GDP ratios, log GDP deviations and spread levels can be created. We show, using regressions similar to those used in our empirical work that the conditional correlation between spreads and poverty rates is of a similar magnitude and sign as in the international data.

1.1 Related Literature

Our paper is related to recent work that links inequality to default risk and especially to Andreasen et al. (2019) who show that income inequality and regressive taxes make default more likely for a given level of debt, and tighter political constraints reinforce this effect. While our work shares an interest in the role that feasibility of fiscal policies play in determining default risk, not only is our focus on international differences in poverty rates, the specifics of the political constraint and fiscal policies are quite different. They have exogenously given taxes that depend on the income of the agent whereas our taxes are endogenous and may be positive or negative depending on income. Moreover we focus on a tax-transfer scheme that acts like a social safety net whereas they primarily focus on tax receipts to finance debt. For this reason, either everyone pays taxes or receives transfers from the government whereas in our model a fraction may pay taxes while the rest receive transfers. In Andreasen et al. (2019), the fiscal plan of the government must be approved by the population otherwise a default occurs. While this sounds similar to our political constraint there are important modelling differences which are driven by the context of the studies. When deciding whether to participate or not, households in our model compare the expected value of participating in the social safety net with the value of not participating, not with the value available to them in default (as happens in Andreasen et al. (2019)). Indeed, the tax-transfer scheme may continue to be implemented in default in our model whereas this is ruled out by assumption in theirs. We found it unreasonable to assume that there is no fiscal plan in operation while the government is in default in our context whereas it makes sense in theirs since taxes are only collected to repay debt. In our model, the government only proposes tax-transfer schemes that will be approved by households and sometimes these schemes involve defaults on existing debt. The strong assumptions regarding the effect of the political constraint is needed in their paper because the tax structure is set in stone and does not respond to the state of the economy. In contrast, ours responds each period and households evaluate their willingness to participate each period.

Our paper is also related to Jeon and Kabukcuoglu (2018) which explores how income inequality matters for government borrowing and default decisions[1]. The authors add time varying inequality shocks to a model with two types of households and impose a taxation regime that can be more or less progressive. They show that rising income inequality within a country increases the probability of default significantly. The idea is that during times of adverse inequality, the marginal utility of the poor is significantly higher than that of the rich and the government is tempted to default in order to try to equate the marginal utility of the two types. In contrast, the tax-transfer scheme in our model is endogenously

[1]Azzimonti et al. (2014) links higher income inequality to higher individual income risk which in turn leads to more public debt. Similarly, Arawatari and Ono (2017) show that higher income inequality causes politicians to reallocate fiscal pain to future generations by using debt to finance part of their government expenditure. Related work by Bi (2012) shows that the government's ability and willingness for debt-repayment depends on the fiscal limit which is determined by factors such as political uncertainty and size of the government.

determined and the amount of taxes that can be charged to the agents is limited by the political constraint discussed above. Moreover, we focus on variation in poverty rates across economies, not inequality shocks that change inequality over time. The implied differences in inequality between economies remain constant as a result. Deng (2021) also explores the role of income inequality on sovereign default risk with an emphasis on worker migration. The main channel through which inequality affects the spread is that the government of an economy with more inequality will use more progressive taxes for redistribution and to reduce consumption inequality. The progressive tax not only discourages labor supply but also induces emigration from the nation and shrinks the tax base, thus reducing government's ability to repay debt. At the same time higher inequality increases the temptation of the government to use defaults as a means to increase the progressivity of the tax base. She finds that income inequality and its interaction with migration explain about one-third of the average government spread. Our work differs in several ways from these studies in that we show that the proportion of poor agents in an economy has an independent role to play beyond inequality in determining the level of sovereign risk. While an increase in the proportion of poor households in an economy increases inequality, we show that the impact of this change on the terms of credit offered by international lenders and the debt choices of the government are robust to controlling for inequality. Moreover, we deliberately switch off any distortions caused by changing taxes either through labour supply distortions or through worker migration in order to highlight the interaction of the political constraint with the international cost of borrowing which lies at the heart of our mechanism.

Distributional incentives also play a role in domestic default risk and debt dynamics in D'Erasmo and Mendoza (2016) and D'Erasmo and Mendoza (2021) but the key heterogeneity among citizens is in their holding of government bonds so that defaults can cause a rearrangement of wealth among them and improve consumption dispersion. In contrast, we focus on income dispersion and especially the role of absolute poverty in determining default risk. In addition, unlike our model, government spending and taxes are exogenous (but transfers vary to satisfy the budget constraint) and there is no political constraint limiting the fiscal policy choices of the government. While our model does not have elections or political turnover, the presence of the political constraint which can induce political defaults links our work to the literature that studies the role that politics plays in defaults, fiscal policy and in sovereign debt dynamics. Some examples from this literature include Amador (2004), Chatterjee and Eyigungor (2019), Cotoc et al. (2021), Cuadra and Sapriza (2008), and Hatchondo and Martinez (2010). Our work is also related to sovereign default models with government spending and taxes that follow Cuadra et al. (2010). Bianchi et al. (2019) studies optimal fiscal policy in a New Keynesian model with unemployment where the government not only runs a tax and transfer scheme, it also offers unemployment insurance. Our work is also related to Alamgir et al. (2022) who find that cross-country differences in sovereign spreads are positively correlated with the average bribe rate in the country. The authors build a sovereign default model to generate a positive relationship between the average amount of diversion of public resources by bureaucrats and the average spread. Unlike this paper, the key mechanism for additional default risk is that defaults actually lower the diversion of government resources and allow the sovereign to increase the provision of public services. That study also differs in that all households in the economy are identical.

The remainder of the paper is structured as follows: Section 2 presents the empirical findings that motivate the theoretical analysis. Section 3 presents the model and section 4 describes the calibration of the model while Section 5 presents the main quantitative implications. Section 6 concludes the paper. An appendix with data sources and computational algorithm is contained in Section 7.

2 Empirical motivation

In this section we show that the positive correlation between the sovereign spread of countries and the proportion of their population that lies below the poverty line discussed in the introduction is a robust feature of international data. We begin with a brief discussion of the main variables.

In our data the international poverty line is set at $1.90 per day in 2011 purchasing power parity terms. This represents the mean of the national poverty lines found in the poorest 15 countries ranked by per capita consumption (World Bank, World Development Indicators). The poverty headcount ratio (our measure of national poverty) identifies the share of a country's population with income less than the poverty line. We use JP Morgan's

Table 1: Summary Statistics

Variable	Mean	Median	Std. dev	Minimum	Maximum	N
Spread	498.30	341.52	518.92	29.76	3926.79	552
Poverty headcount	10.99	8.43	10.89	1.16	59.33	777

Emerging Market Bond Index (EMBI) spreads which capture the difference between the yield of sovereign bonds of a country in basis points and the yield of US 10-year treasuries between the period 1995-2015. We start with all countries in the EMBI data and keep those for which at least five years of data is available. We also limit nations to those that have an average poverty headcount ratio that is greater than or equal to 1%. This leaves us with 37 middle and low income countries which are listed in the appendix in section 7.2.

Table 1 displays some descriptive statistics for sovereign bond spread and poverty headcount ratio. The mean spread is about 498 basis points and ranges from a minimum of 29.76 to a maximum of 3926.79 basis points. For poverty, the mean is around 11 % while the lowest and highest values are 1.16 and 59.33 percentage points respectively.

In order to show that the positive correlation between poverty headcount ratios and sovereign spreads is robust to conditioning on other factors, we use the following empirical specification:

$$spread_{it} = \beta_0 + \beta_1 poverty_i + \Gamma' X_{it} + \Lambda' Z_i + \Psi' Y_t + u_{it}$$

where Y_t refers to a vector of global variables that are common to all countries but vary over time, Z_i refers to vector of country averages that do not change over time, while X_{it} refers to other control variables that vary over time and across countries. The vectors (Γ, Λ, Ψ) consist of coefficient estimates of the corresponding variables. We would like to emphasize that the goal of this specification is not to establish a causal link but rather to establish the robustness of the correlation in question. Data on macroeconomic control variables are collected from the World Development Indicators data set of the World Bank. The global variables, which include the ten year US treasury bill rate and the ten-year to one-year treasury yield curve, are taken from the Federal Bank of St Louis FRED database. We include the debt to GDP ratio as a domestic control variable because higher levels of debt relative to repayment capacity are expected to be correlated with higher spreads both in theory and in previous empirical work. We also include the log deviation of GDP per capita (in constant 2010 dollars) from trend as a measure of the state of the business cycle of the country as a conditioning factor since default risk has been shown to increase in recessions. Given the history of defaults in Latin American countries, we include a dummy variable that takes the value of 1 if a nation is part of that region and 0 otherwise. We will discuss additional variables as we add them below.

Our results are reported in Table 2 where column (i) presents the simplest specification with the poverty headcount ratio and a constant term as the only variables. Column (ii) adds the Latin American dummy and column (iii) further adds the current level of debt as a proportion of GDP and the log deviation from trend of real per capita GDP. Several studies have suggested that global factors play a role in determining spread levels and default risk (see Johri et al. (2022) for a sovereign default model of the impact of the level and volatility of the world interest rate on default risk as well as references to related empirical studies.) As a result we include year fixed effects in this and the following specifications. All three specifications indicate that there exists a positive and significant relationship between sovereign spreads and poverty rates. The results suggest that on average nations with poverty headcount ratios of 40 percent will display average sovereign spreads that are 120 to 160 basis points higher than nations with poverty headcount ratios of 10 percent. The control variables in column (iii) also exhibit their expected signs – higher debt burdens are associated with higher spreads, whereas higher than trend GDP is associated with lower spreads. The Latin American dummy is positive indicating that spreads are higher in that region relative to other emerging economies. The addition of year effects leaves the relationship of interest intact. We conclude that the positive correlation between poverty rates and spreads is robust to the inclusion of the usual conditioning factors.

Countries with more poor people have been shown to have a higher level of income inequality as well as weak institutions resulting in a higher incidence of corruption. These nations also tend to have lower GDP per capita. The existing literature on sovereign debt has discussed some of these as sources of default risk (see for example the role of the bribe rate of a nation in Alamgir et al. (2022) and the impact of inequality inDeng (2021) and Jeon and Kabukcuoglu (2018)). Alamgir et al. (2022) argue that when bureaucrats are known to divert

Table 2: Poverty and Spreads

	Dependent variable:						
	spread						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Poverty prop	3.40*	4.07**	3.16**	4.33***	3.09**	3.14**	3.14**
	(1.93)	(1.88)	(1.39)	(1.16)	(1.40)	(1.39)	(1.36)
LA dummy		85.23*	155.03***	189.44***	146.19***	143.53***	138.62***
		(45.00)	(44.53)	(51.77)	(44.97)	(44.95)	(46.10)
Debt/GDP			15.87***	16.67***	16.25***	16.01***	16.14***
			(1.81)	(1.99)	(1.87)	(1.93)	(1.76)
Log dev of output from trend			−27.26***	−27.16***	−27.00***	−26.28***	−23.50***
			(6.46)	(6.44)	(6.33)	(6.29)	(5.52)
Gini				−5.05			
				(3.27)			
Bribe rate %					19.98**	19.83**	19.94*
					(10.15)	(10.04)	(10.19)
GDP per capita rank					2.41	2.18	2.54
					(2.00)	(2.01)	(2.10)
Default						332.18	
						(292.14)	
10 year US Tbill							13.85
							(16.48)
US yield curve 10/1							38.99
							(25.90)
Constant	459.14***	419.23***	248.39	417.31**	158.23	169.62	−191.84*
	(28.48)	(29.75)	(151.96)	(202.57)	(163.38)	(164.55)	(109.89)
Year FE	No	No	Yes	Yes	Yes	Yes	No
Observations	552	552	472	465	472	472	472
R^2	0.01	0.01	0.43	0.43	0.43	0.43	0.39
Adjusted R^2	0.004	0.01	0.39	0.39	0.40	0.40	0.38
Residual Std. Error	517.82	516.69	418.45	417.91	417.60	416.70	424.09
F Statistic	3.34*	3.38**	13.80***	13.04***	12.94***	12.62***	36.84***

Standard errors are clustered at the country level, and reported in parentheses. ***,**,* represent significance level at 1%, 5% and 10% respectively.

funds from public projects, the government has an additional incentive to default relative to an economy without diversion since default may curtail diversion and increase the share of government services in times when revenue is below average. Jeon and Kabukcuoglu (2018) shows that periods of rising income inequality within a country increases the probability of default significantly. In light of this work, we explore if the poverty head-count ratio could be serving as a proxy for these other variables? Finally, since low per capita GDP and high poverty headcounts are likely to occur together we wish to control for this variable also. In the remaining specifications, we wish to see if the inclusion of these factors influences our results. In specifications (iv), we add the Gini coefficient as a measure of inequality while in column (v) we introduce the bribe rate and add the rank of initial year per capita GDP as an alternate measure of low income. As is evident from Table 2, the poverty headcount ratio continues to be positively correlated with spreads with little change in the magnitude of the relationship. While the bribe rate coefficient is positive and significantly different from zero in all specifications, neither inequality nor income level rank of nations appear to display a significant correlation with the spread[2]. In column (vi) we add a dummy variable which takes the value of 1 in the year that a nation defaults on its debt and 0 otherwise. The inclusion of this variable, which displays a large coefficient has little impact on the correlation of interest[3]. Finally, in column (vii), we replace the year fixed effects with two widely studied global factors – the yield on the ten year US treasury bill as well as the ten year to one year US treasury yield curve as a proxy for global risk. Neither factor is statistically significant though they display the anticipated signs.

Overall, in all specifications, the proportion of population lying below the poverty line is found to be positively correlated with sovereign spreads. Moreover, the results in Table 1 suggest that poverty rates have an additional marginal impact on country spreads beyond the fact that high poverty headcount ratios occur in nations with low per capita income and

[2]We have replaced the initial income rank of nations with their actual per-capital real GDP level in the initial period as well as in the final year of our sample with very similar results in specifications that are not shown here for brevity. Similarly we experimented with one-period lagged debt to GDP ratio with little impact

[3]We have also experimented with a crisis dummy in the spirit of Catão and Mano (2017) where we assign a value of 1 for countries that have ever experienced default, and zero otherwise. This captures the idea that countries which have defaulted in the past appear to experience higher spreads well into the future. This variable emerges as very strongly correlated with higher spreads but leaves the relationship of interest intact.

display high rates of income inequality. As such, we build a model in the next section where an economy with a higher proportion of poor people will display higher default risk.

3 Model

Consider a small open economy inhabited by a continuum of households of measure one and a government. At every point in time, there are two types of households: poor households comprise λ proportion of the population and non-poor households make up $1-\lambda$ proportion of the population. A non-poor household earns an income given by y while a poor household earns income given by $y_p = \alpha y$ with $\alpha \in (0,1)$. Here α represents the income gap between the average non-poor household and the poor household.

To be consistent with the quantitative sovereign default literature following Arellano (2008), preferences of both types of households follow the CRRA form:

$$u(\mathcal{C}) = \frac{\mathcal{C}^{1-\sigma} - 1}{1-\sigma}. \tag{1}$$

where $\mathcal{C} \in \{c, c^p\}$ represents the consumption of the poor c^p or the non-poor c.

A non-poor household pays non-distortionary taxes (transfers) at the rate τ which are set each period by government policy and consumes whatever remains from their endowment y :

$$c = y - \tau, \tag{2}$$

where the endowment process evolves stochastically according to the transition density $f(y', y)$. Similarly, the budget constraint of a poor household is:

$$c^p = y^p + \tau^p, \tag{3}$$

where the notation anticipates that a poor household will receive transfers though this is determined by the government in equilibrium each period and no restrictions are placed on the sign of τ^p. We note that the income gap of the two types of households is always held constant. This modeling choice deliberately distinguishes us from Jeon and Kabukcuoglu (2018) which studies the role of time varying inequality on default risk.

Households of either type do not save and make no dynamic decisions. We note that in this version of the model, we deliberately eschew distortionary taxation in order to focus on the relationship between the proportion of the poor in an economy and its default risk. Naturally, if higher distortionary taxes are required by the social safety net in an economy with more poor households then the ensuing reduction in labor supply could reduce revenue and be an additional source of default risk. This impact of income losses due to distortionary taxes on default risk has already been explored in the literature since Cuadra et al. (2010) and in several other contexts so we eliminate them to isolate our distinct mechanism (see Deng (2021) and Alessandria et al. (2020) for examples where the economy suffers from distortions due to both a decrease in labour supply as well as migration out of the economy).

Before turning to the government, we provide an expression for aggregate output or GDP as the weighted sum of household endowments:

$$Y = \lambda y^p + (1-\lambda)y = (\lambda\alpha + 1 - \lambda)y \tag{4}$$

and we note that aggregate output, Y_t is proportionate to the endowment of the non-poor, y_t. Since λ does not vary over time, aggregate output follows the endowment process, y_t.

3.1 The Government

The government maximizes a social welfare function where the utility of the poor and the non-poor households are weighted according to the parameters λ_{pol} and $1-\lambda_{pol}$ respectively. These social welfare weights used by the government are not necessarily equal to the proportion of poor and non-poor households but govern the *desired* size of the social safety net that the government implements through its tax-transfer scheme. They can be viewed as the outcome of a social process outside the scope of the model or simply as capturing the average ideology of the nation which underpins the government's choice of social policy.

The government is the only agent in the economy who can borrow and lend in international credit markets and borrows from abroad by selling long-duration non-contingent bonds. As in Hatchondo and Martinez (2009), we assume that a bond issued in period t promises an infinite stream of coupons which decrease at a constant rate δ.[4] As such, a bond issued in period t promises to pay $\kappa(1-\delta)^{(j-1)}$ units of tradable goods in period $t+j$ with $j \geq 2$. Let

[4] See also Chatterjee and Eyigungor (2012) and Arellano and Ramanarayanan (2012).

d denote the number of coupons due at the beginning of the current period, while d' denote the number of coupon claims due at the beginning of next period, κ is a parameter driving the coupon payment of bond. Then the number of new bonds issued by the government is given by the following:

$$\iota = d' - (1 - \delta)d$$

Following the literature, we assume that the government cannot commit to repaying its debts. If the government does not default, it repays existing claims and can issue new debt, ι. If it defaults, it defaults on all existing claims and is excluded from credit markets for a stochastic number of periods. The bond price $q(d', y)$ takes default risk into account and will be discussed further below. The transfers made to the poor as part of the social safety net is the only domestic expenditure incurred by the government so its current period budget constraint is given by:

$$\lambda \tau^p + \kappa d = (1 - \lambda)\tau + q(d', y)[d' - (1 - \delta)d] \tag{5}$$

where the left hand side of the equation provides the uses of public funds while the right side provides the sources of public funds.

The government is limited in its ability to redistribute income using taxes by a political constraint which states that households must be willing to participate in the tax-transfer scheme. The idea here is that government's fiscal choices either face a referendum or the government takes into account that it must win elections which are lost if households are unhappy with the utility delivered by fiscal policy.

In order to motivate the non-poor to participate in a tax-transfer policy we offer two possible rationalizations. The first is the need for a social safety net to cover the possibility of an idiosyncratic household income shock. We assume that in every period there is a constant probability, p, that a household will draw a poverty income level receiving αy instead of y that period. In this situation, the tax-transfer scheme acts as an insurance device against the small probability of drawing a low income. We assume that the aggregate proportion of poor remains unchanged over time and government policy operates with the aggregate in mind. The other rationalization assumes that non-poor households care about the utility of poor households to a limited extent and p is the relative utility weight attached to the consumption of the poor while 1-p is attached to own utility. We can attribute this benevolence to feelings of justice, charity, equity etc. While a non-poor household is willing to give up some consumption to raise the consumption of poor households, a fiscal plan that gives too much away to the poor at the expense of the non-poor will have that plan rejected. Similarly a plan that gives the poor too little will also be viewed as unsatisfactory. Having said this, in our calibration of the model we will proceed with the first approach and obtain estimates of the transition probability from non-poor to poor households using South African data in order to tie down p.

Formally, government policy must satisfy the political support constraint:

$$W_\tau \geq W_0 \tag{6}$$

where W_τ is the utility in period t given the proposed fiscal plan (τ, τ^p), while W_0 is the utility without any fiscal program where all households consume their endowment (y, y^p).

$$W_\tau = pu(c^p) + (1-p)u(c) \tag{7}$$

$$W_0 = pu(y^p) + (1-p)u(y) \tag{8}$$

As explained earlier, p is the ex-ante transition probability from being a non-poor household to a poor household in which case W_τ may be thought of as the expected utility of the household under the tax-transfer scheme while W_0 is in its absence. This implementability constraint thus limits the amount of redistribution to the poor. It can, in certain states, also limit the ability of the government to collect sufficient taxes for repaying existing debt. As such, it can influence the government's decision to default or repay, which we discuss below.

We can think of the government's problem in two steps. Given the constraints discussed above, and the amount of resources available to the government from its debt policies, the government chooses the optimal fiscal policies (τ, τ^p) in response to the realization of the endowment shock y. The goal of the government in this step is to smooth consumption across households by minimizing the gap in socially-weighted marginal utilities of consumption between household types. The first-order condition from the government's optimization problem gives us:

$$U_c(c) = U_{c^p}(c^p)(\frac{\lambda_{pol}(1-\lambda)}{\lambda(1-\lambda_{pol})}) \tag{9}$$

Note that this equation ties down relative consumption of the two types, not the absolute consumption levels which also depend on debt policy discussed below.

Timing The timing in the model is as follows:

1. The output shock y is realized.
2. The government decides whether to repay its debt obligations or default and how much to borrow. Given the debt choices, the government proposes the optimal tax-transfer scheme. These decisions must ensure that equation (10), the political constraint (6) is satisfied.
3. Finally households consume with respect to their types.

3.2 Government value functions and recursive equilibrium

Each period, the government chooses between honoring its outstanding foreign claims or defaulting in order to maximize the weighted welfare of households. When the government defaults, it gets excluded from the credit market for a stochastic number of periods. Following default, the government may regain access to the international credit market with an exogenous probability θ, and it does so with no outstanding claims, $d = 0$. Default also entails direct costs such that output is lower during the periods the government is in autarky.[5] That is, each non-poor person's endowment is

$$y_{def} = h(y) \leq y \tag{10}$$

while each poor person's endowment is

$$y^p_{def} = \alpha h(y) \leq y^p \tag{11}$$

[5] Mendoza and Yue (2012) offer an endogenous interpretation of default costs that are related to domestic access to foreign goods.

The value function when the government has access to the international credit market and owes d, given endowments (y, y^p) is given by $V^0(d, y)$. The government decides whether to default or repay its debts by comparing the value associated with paying back and remaining in the credit market $V^c(d, y)$ with the value associated with defaulting and going to temporary autarky $V^d(y)$. The problem can thus be expressed as:

$$V^0(d, y) = max[V^c(d, y), V^d(y)] \tag{12}$$

The value function when it does not default is given by:

$$V^c(d, y) = max[(1 - \lambda_{pol})u(c^*) + \lambda_{pol}u(c^{*p}) + \beta \int_{y'} V^0(d', y')f(y', y)dy'] \tag{13}$$

subject to

$\lambda\tau^p + \kappa d = (1 - \lambda)\tau + q(d', y)[d' - (1 - \delta)d]$

$c^* = y - \tau$

$c^{*p} = y^p + \tau^p$

$W_\tau \geq W_0$

$U_c(c) = U_{c^p}(c^p)(\frac{\lambda_{pol}(1-\lambda)}{\lambda(1-\lambda_{pol})})$

Similarly, when the government defaults on its debt, the value of default is given by:

$$V^d(y) = max[(1-\lambda_{pol})u(c^*_{def})+\lambda_{pol}u(c^{*p}_{def})+\beta \int_{y'} [\theta V^0(0, y')+(1-\theta)V^d(y')]f(y', y)dy'] \tag{14}$$

subject to

$(1 - \lambda)\tau - \lambda\tau^p = 0$

$c^*_{def} = y_{def} - \tau$

$c^{*p}_{def} = y^p_{def} + \tau^p$

$W_\tau \geq W_0$

$U_c(c_{def}) = U_{c^p}(c^p_{def})(\frac{\lambda_{pol}(1-\lambda)}{\lambda(1-\lambda_{pol})})$.

where we assume that when the economy regains access to financial markets, it starts with zero debt.[6]

[6]In contrast, Yue (2010) offers a model with partial debt recovery and renegotiation between lenders and borrowers.

Later we will show that if repayment implies that equation (6) cannot be satisfied, the government will default in order to prevent too large a share of tax revenue from leaving the country, however the government may also default when the political constraint is satisfied. We also note that while it is not guaranteed in autarky that both 6 and 9 will be satisfied this is true for our parameterization.

The optimal default policy of the sovereign is thus characterized by:

$$D(d,y) = \begin{cases} 1 & V^d(y) \geq V^c(d,y) \\ 0 & otherwise \end{cases}$$

As such, given a level of claims d, the default set can be characterized as :$\mathcal{D}(d) = \{y \in Y : D(d,y) = 1\}$. Furthermore, given any set of claims and current realization of the income process y, the probability of default $\mathcal{P}$ can be inferred from the default set and the transition process for y.

3.2.1 International lenders

There are a large number of identical foreign lenders who can borrow or lend at the world interest rate r^* and participate in a perfectly competitive market to lend to the small open economy. The lenders have perfect information about the economy's endowment process, the distribution of poor households, as well as other parameters governing the economy and can observe y every period. They are risk neutral and they maximize their expected profits. With perfect competition in the credit market, a zero profit condition for the foreign creditor is satisfied. Hence the bond price is given by:

$$q(d',y) = E_{y'|y} \frac{(1-D(d',y'))(\kappa+(1-\delta)q')}{1+r^*}$$

3.2.2 Equilibrium

The recursive equilibrium for this economy is characterized by a set of policy functions for (i) consumption (c and c^p), (ii) international claims on the government d' and default rule D, (iii) government's fiscal choice τ and τ^p, and (iv) a bond price $q(d',y)$

such that:

1. Given the sovereign bond price $q(d', y)$ and the political constraint $W_\tau \geq W_0$, the government's policy set $\{d'(d, y),\ D(d, y),\ \tau(d, y),\ \tau^p(d, y)\}$ satisfies the government's optimization problem.

2. Given the government's policy set $\{d'(d, y),\ D(d, y),\ \tau(d, y),\ \tau^p(d, y)\}$, the household's consumption choices $c(d, y)$ and $c^p(d, y)$ satisfy the political constraint and the resource constraint.

3. The sovereign bond prices $q(d', y)$ reflect the government's default probabilities and are consistent with lender's expected zero profit condition and government's optimization problem.

4 Calibration

The model is solved numerically using value function iteration which requires parameter values to be assigned. We use annual data to calibrate to the South African economy over the period 1995-2015 and refer to this as the Benchmark model.

Table 3: Parameter values

Description	Parameter	Value	Target/Source
persistence	ρ	0.822	Estimation
sd of ε	σ_ϵ	0.016	Estimation
risk aversion	σ	2	Prior literature
world interest rate	r	0.04	Prior literature
re-entry probability	θ	0.154	Prior literature
poverty transition prob.	p	0.06	NIDS data
proportion of poor households	λ	0.245	NIDS Data
income share of poor households	α	0.16	NIDS Data
Coupon decay rate	δ	0.084	Avg duration of South African bonds
Bond coupon	κ	0.119	$(r^* + \delta)/(1 + r^*)$
political weight	λ_{pol}	0.067	Social expenditure/GDP
discount factor	β	0.927	Average debt/GDP
default cost	ψ_0	-0.698	Standard dev of spread
default cost	ψ_1	0.758	Average spread

We assume that the aggregate output follows an AR(1) process:

$$lnY_t = \rho lnY_{t-1} + \epsilon_t, \tag{15}$$

with $|\rho| < 1$ and $\epsilon_t \sim N(0, \sigma_\epsilon^2)$. In order to estimate ρ and σ_ϵ, we use logged and linearly de-trended GDP data for South Africa covering the period 1995-2015 which gives a value of 0.822 for ρ and 0.016 for σ_ϵ. Since the endowment process y follows aggregate output Y due to the assumption of fixed proportion of poor and non-poor households, we can use the estimated parameters above to calibrate the endowment process.

In order to estimate p, we employ South Africa's National Income Dynamics Study (NIDS). NIDS is a panel survey based on the livelihoods of households and individuals tracked over time. The study began in 2008 and continues to be carried out every two years with a nationally representative sample of over 28000 individuals in 7300 households across the country (Brophy et al. (2018)). The fact that the survey contains repeated households over time, allows us to study poverty transitions. Our estimate of p is based on total household income derived from the NIDS data set for Wave 1 (2008), Wave 2 (2010), Wave 3 (2012), Wave 4 (2014-15), and Wave 5 (2017). We begin by calculating the per capita income of all households in all the waves using the survey weights to correct for non-response and systematic bias. In order to adjust for inflation, we use Statistics South Africa's headline CPI index to deflate the nominal income data to their real values. Next, using the upper bound national poverty lines relevant to each wave, we identify the poverty status of each household in the waves. Finally, the transition probabilities are calculated by counting the number of households that changed poverty status from wave to wave. We find that, on average, 12% of those who were non-poor in $Wave_i$, transitioned into poverty in $Wave_{i+1}$. This leads us to attach a value of 0.06 per year to p assuming a steady flow into poverty over the two year period.

We also use the NIDS data to tie down model parameters λ and α to the proportion of South African households that are poor and to their relative income respectively. The value of λ in the data is found to be 0.245 which is the mean proportion of households who lie below the poverty line in wave 1 through 5 of the NIDS dataset. Similarly, we calculate the income gap between the non-poor and poor households, α, using the data on per capita income of households and the national poverty line. Specifically, we express the national poverty lines of each wave as percentage of the average income per capita of the non-poor households in the corresponding wave which generate a value of 0.16 for α. This simplifying assumption ignores that some households may be below the poverty line, not on it, so that the average income of the poor may be less than that implied by the poverty line. On the other hand, there may be non-monetary grants for these households which are incompletely accounted for in NIDS so we opted for a flawed but clearly stated measure of α.

Following the standard values used in the real business cycle literature for small open economies as well as in the quantitative sovereign default literature, we assume a coefficient of relative risk aversion of 2 and a world interest rate of 4 percent annually. The probability of re-entry after default into international credit market is set to 0.154, such that the government remains in financial autarky for a period of approximately 6.5 years on average following a default event, consistent with Chatterjee and Eyigungor (2012). We set $\delta = 0.084$ and use it with the average yield to produce an average bond duration of 7.25 years in the simulation which is consistent with the average average bond duration for South Africa as reported in Bai et al. (2017).[7]

We calibrate the welfare weight of the poor in the government's social welfare function λ_{pol}, so that our model can match the transfers to GDP ratio in South African data. According to the Government Finance Statistics, IMF, public social expenditure comprises approximately 8.1% of GDP in South Africa over the period 1995-2015. This target is achieved by setting the socio-political welfare weight λ_{pol} equal to 0.067. Finally, we are left with three parameters to assign values to: the discount factor β and default cost parameters ψ_0 and ψ_1. We calibrate these parameters to jointly target the following moments from the data:(i) a mean external debt-to-GDP ratio of 25.1%; (ii) a mean sovereign spread of 2.3%, and (iii) a standard deviation of 1.0% for the spread. All the data counterparts correspond to the average values in South Africa over the period 1995-2015 extracted from the World Development Indicators dataset of the World Bank and the JP Morgan's EMBI Global Index.[8]

Table 4 displays these targeted moments along with some moments that were not targeted in the calibration. In order to calculate these moments, we simulate the model by generating

[7] We use the Macaulay duration of bond formula $D = \frac{1+r}{r+\delta}$ where r refers to the average yield generated by the bond.

[8] We thank César Sosa-Padilla for sharing sovereign spread data with us.

Table 4: Targeted and non-targeted moments

	Data	Model
Targeted moments		
Mean spread (%)	2.3	2.3
Mean debt/GDP (%)	25.1	25.2
Standard deviation of spread	1.0	1.0
Mean transfer/GDP (%)	8.1	8.1
Non-targeted moments		
corr(spread, GDP)	-0.4332	-0.8289
corr(tb/y, GDP)	-0.4104	-0.8624
corr(spread, tb/y	0.3566	0.8523

500 sample paths of 1500 periods each, and then discard the first 500 periods to eliminate the influence of initial guesses. Moments are then calculated from the detrended series of the remaining data after excluding the default episodes. [9] The calibrated model comes close to the targeted values. The model also does quite well in generating the expected cyclical properties found in the literature. As observed in the data, both spreads and trade balance/GDP are countercyclical, and spreads are positively correlated with trade balance to GDP.

5 Results

In this section we will use variants of the Benchmark model economy to understand the economic forces that may lie behind the observed positive correlation between poverty head-count ratios and spreads. To begin, we modify the Benchmark economy by setting $\lambda = 0.42$ while keeping all other parameters unchanged. We will refer to this economy as the High-Poverty economy. Figure 2 depicts the default regions associated with these two economies where the endowment level is shown on the vertical axis and the level of debt normalized by mean y is on the horizontal axis. The green line refers to the Benchmark economy, while the red line refers to the High-Poverty economy. Each line divides the space into two regions such that the government of that economy will default in the represented states under the respective curve. The main message of the paper is immediately apparent from Figure 2 – the

[9] We detrend the log of income, consumption and trade balance using the Hodrick-Prescott filter with a smoothing parameter of 100.

Figure 2: Default regions

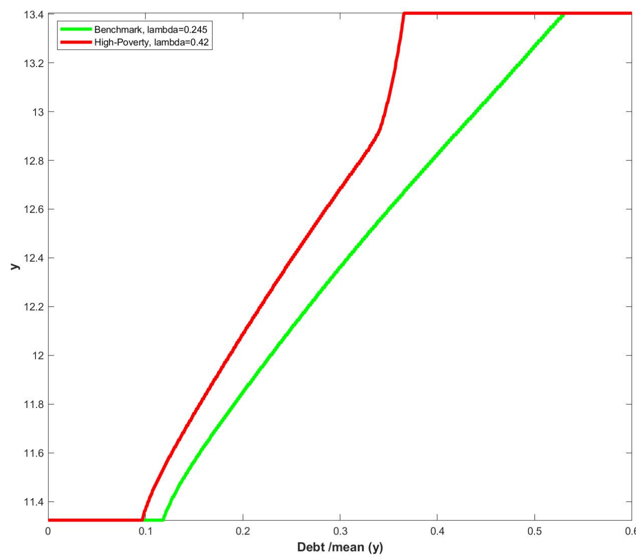

High-Poverty economy's default set is larger and will be associated with higher default risk. Moreover, focusing for the moment on the green line only, for any level of debt/mean(y) above 0.11, a sufficiently low realization of income leads to default. As the level of debt rises from that point, the rising green line implies that the income level that is needed to avoid default also rises. Moreover, while not shown, the government will default at even the highest income level when debt/mean(y) is sufficiently high. Finally, note that at a sufficiently low level of debt, the government will not default at even the lowest income levels. These properties of the default region of an economy are quite standard in the quantitative sovereign default literature. Turning to the red line, we see that the High-Poverty economy can support a much lower amount of debt without defaulting at even the highest income level. On the other hand, the level of debt at which default can be avoided at the lowest income level is smaller than the Benchmark economy. Above that level (where the red line intersects the horizontal axis), the red line is always above the green line indicating that identical debt levels would require higher y levels for the High-Poverty economy to avoid default compared to the Benchmark economy.

The implications of the larger default region of the High-Poverty economy are reinforced using Figure 3 which shows the menu of debt (normalized by mean(y)) and interest rate

Figure 3: Spread-debt menu at mean(y)

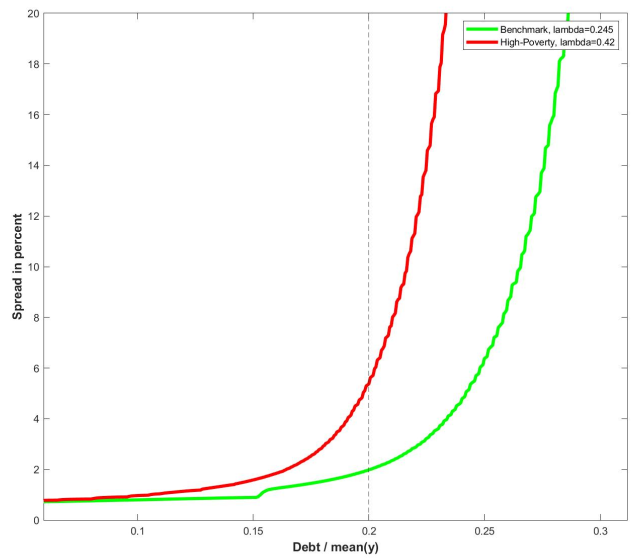

spreads on offer to the two economies at mean endowment levels. Note that the interest rate spread is calculated as the interest rate currently paid by an economy minus the world interest rate so that when there is very little debt in either economy, the spread is small. Since debt is long-lasting, lenders take into account the possibility of future indebtedness levels and future default possibilities so the spread may not be exactly zero even with little currently owed by the government. From the bottom left of the figure, as we move rightward, the spread in either economy is increasing in the chosen amount of debt, which is consistent with the default regions discussed above. At any given level of normalized debt in Figure 3, the red line lies above the green line, implying that the High-Poverty economy will pay higher spreads than the Benchmark economy. We can also see that the debt limits are endogenously imposed by the equilibrium menus (since the spread rises to extreme levels at debt levels approaching this limit), and that this limit (the nearly vertical regions of the lines) is much lower for the High-Poverty economy. Figure 3 shows the much harsher terms of credit available to the High-Poverty economy compared to the Benchmark economy at the average debt/mean(y) level of 0.2 (see vertical line) – the spread is 5.4 percent vs. 1.9 percent. As a result, the government of the High-Poverty economy will choose lower equilibrium debt levels than the Benchmark economy and this can be seen in Figure 4 which plots the current obligations due relative to mean y on the horizontal axis and the corresponding choice of next period obligations on the vertical axis while the endowment is held at the mean. At high debt levels, first the economies start to lower debt below current levels and eventually the almost vertical lines indicate the point at which no debt is sustainable due to extremely high default risk and hence high spreads. This happens at lower debt levels for the High-Poverty economy.

Figure 4: Debt choice

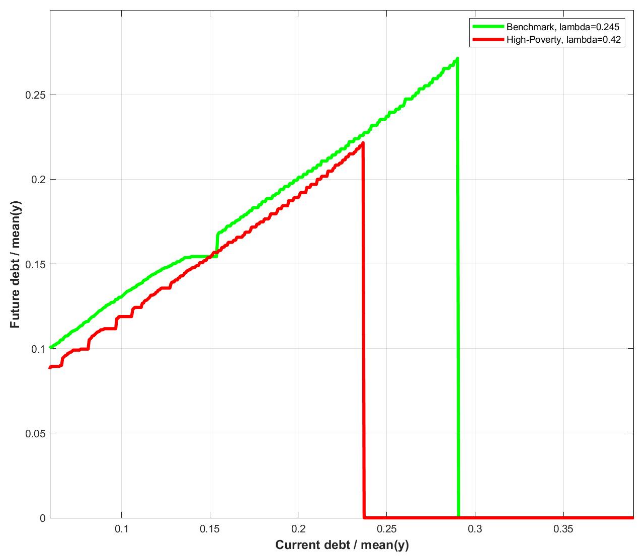

Since the High-Poverty economy faces worse borrowing terms, it will choose lower debt levels and therefore potentially lower equilibrium spreads. In order to see if the equilibrium behavior of the model is consistent with the empirical correlations discussed earlier which were conditioned on debt levels, we run a simple regression of spread on the proportion of the population in poverty, debt to y and log deviation of y from its mean using model-simulated data. The simulated data is generated from the Benchmark economy with a λ of 0.245 as well as twenty-four other economies with different values of λ, ranging from 0.11 to 0.59. For each economy we simulate 500 sample paths with 1500 periods each, and then drop the periods in default. As is evident from Table 5, the estimated coefficient on the proportion of poor households is positive and statistically significant at 1% level of significance after controlling for the level of debt to y and log deviations in y. This is consistent with our empirical results presented in Table 2. We have reproduced the most directly relevant empirical result in column 2 for comparison ease.

Table 5: Poverty and Spread: simulations vs. international data

	Dependent variable: spread	
	Model	Data
poverty prop	5.69***	3.17**
	(0.02)	(1.32)
LA dummy	NA	149.80***
		(42.94)
debt/GDP	21.76***	15.82***
	(0.09)	(1.69)
log dev of output from mean	−33.00***	−25.01***
	(0.07)	(5.59)
constant	−276.18***	11.71
	(1.99)	(44.74)

Note: Standard errors are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

In order to understand why economies with a higher proportion of poor households are more likely to default, we examine the properties of the equilibrium tax-transfer scheme. Recall that fiscal policy serves as an insurance device in the event that a non-poor household draws a poverty level income shock. The effectiveness of the insurance scheme depends on the value of the transfers received by the poor relative to the taxes paid by the non-poor. To augment transfers, the government can also borrow from international lenders, however, when indebtedness is high relative to income, the terms of credit will worsen and this can lead to situations where a large fraction of taxes flow out of the country to repay international lenders. A weakening ability to fund the consumption of poor makes defaults look attractive to the government when this occurs. Moreover, it makes the tax-transfer scheme less attractive as an insurance scheme to the non-poor and this brings the political constraint into play. We will discuss these in turn.

Figure 5 plots the transfer to tax ratio (measured as total transfers made relative to total

Figure 5: Transfer to tax ratio at mean(y)

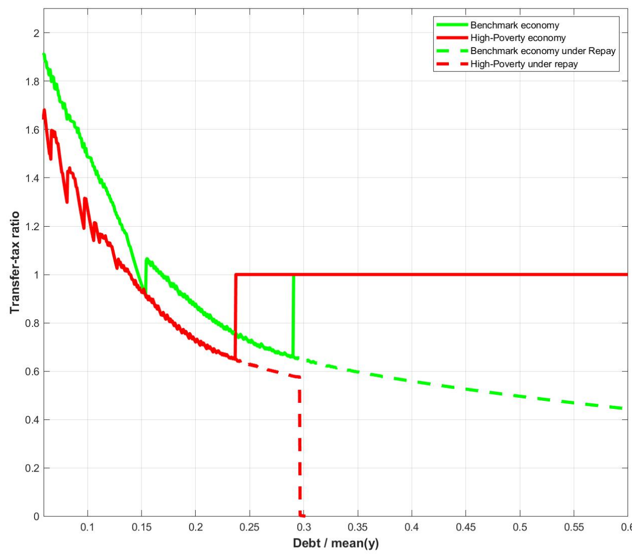

tax collected) for both economies, holding y at the mean level. The green line in the figure refers to the Benchmark economy while the red line represents the High-Poverty economy. The dotted lines show what the transfer to tax ratio would have been if the two economies always chose to repay regardless of the optimal decision. Note that when it is optimal to repay, the solid and dashed lines overlap but when the government prefers to default, the solid line diverges from the dashed line. We first note that the Benchmark economy can usually support a higher level of transfers relative to taxes collected than the High-Poverty economy. This occurs because of the more favourable credit terms available to it. In addition, it can be seen that at lower levels of debt, the transfer to tax ratio is well above 1 for both the economies. This implies that the excess transfer payments above the tax revenue collected are financed by borrowing from international lenders. As the debt to endowment ratio increases or as the aggregate income in the economy falls (not shown), the transfer to tax ratio falls below unity. In each economy, there comes a level of the transfer to tax ratio at which the government prefers to default rather than repay its obligations because repayment takes too large a share of tax revenue away from transfers. This event can be seen in the vertical jump in the transfer to tax ratio and the opening of a gap between the solid and dotted lines. The horizontal section of the transfer-tax ratio line at unity is associated with autarky since the

Figure 6: Political constraint at debt/mean(y)=32%

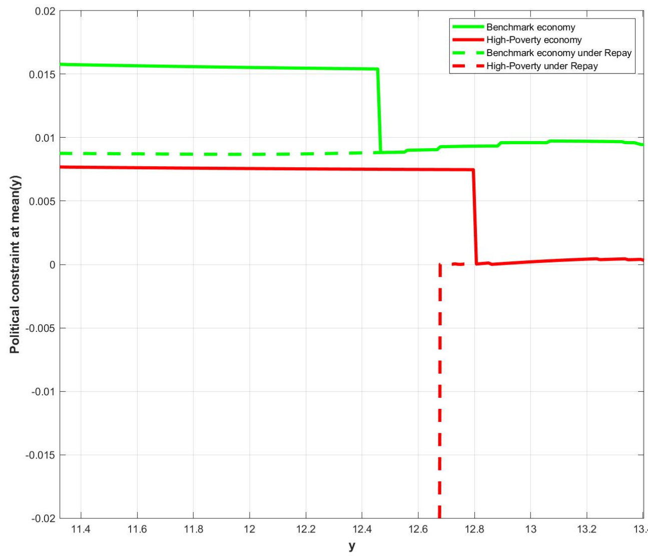

government cannot borrow having just defaulted on its debts. We note that the Benchmark economy defaults at a higher transfer to tax ratio than the High-Poverty economy. We also note that the jump in the red line is larger than in the green line which signifies a greater gain in transfers relative to taxes by defaulting for the High-Poverty economy.[10]

The model can generate two types of defaults. The typical default seen in the Benchmark economy occurs because governments wish to maintain the consumption levels of poor households relative to non poor households through transfers. When credit terms take a turn for the worse, taxes start to fund too little of transfers and the government chooses what may be called non-political or regular defaults in order to free up fiscal space. In addition, the economy can display a political default. This occurs when the political constraint is violated under repayment but not violated under default. In order to understand this type of default we display the behaviour of the political constraint in the two economies in Figure 6 at a debt to mean(y) of 32% for different values of y. The political constraint is formulated by taking the difference between the ex-ante expected utility of a household with and without the tax-transfer policy in operation. Households are in favor of the tax-transfer policy as long as this difference is positive. Going from right to left, the endowment level falls and

[10]We have cut the vertical axis at zero in order to focus on the economically relevant region. This is why the negative part of the dashed red line is not visible.

the political constraint falls as well in both economies. The green political constraint graph for the Benchmark economy in Figure 6 is always positive, implying that the fiscal policies implemented by the government are politically feasible. In contrast, the red graph hits zero when the endowment level is around 12.8 roughly. This triggers a political default in the High-Poverty economy. The default allows the government to raise transfers relative to taxes so that the non poor are once again willing to participate in the scheme. We can also see a non-political default occur in the Benchmark economy at slightly lower endowment levels. Once again, the dashed lines in the figure display the constraint if the government was to repay at all endowment levels so that a divergence between the solid and dashed lines indicates a situation of default. Moreover, it is worth noting that for the High-Poverty economy, the red dotted line is negative indicating that if the government chose to repay rather than defaulting, that would violate the political constraint. In contrast, for the Benchmark economy, the dotted line is still positive below the endowment level at which default occurs, implying that the political constraint would still be satisfied if the government chose to repay. So, default in the Benchmark economy occurs because the government wishes to maintain the consumption levels of its citizens.

5.1 Controlling for GDP and inequality

The increase in the proportion of poor households in the High-Poverty economy has two additional implications relative to the Benchmark economy. First, inequality as measured by the mean value of the Gini coefficient rises from 0.1956 in the Benchmark economy to 0.3162 in the High-Poverty economy. Second, overall GDP is reduced from an average value of 7.97 in the Benchmark economy to 6.48 in the High-Poverty economy. In order to show that the main underlying factor increasing the default risk of the High-Poverty economy is the rise in the proportion of poor households as opposed to changes in inequality or GDP, we will control for these effects next.

We control for inequality by raising α until the Gini coefficient is restored to the value in the Benchmark economy. Note that in this economy there are more poor people but they are not as poor as in the Benchmark economy. Similarly, we control for the fall in GDP in the High-Poverty economy by simultaneously increasing the income of all households, (by raising

Figure 7: Default regions

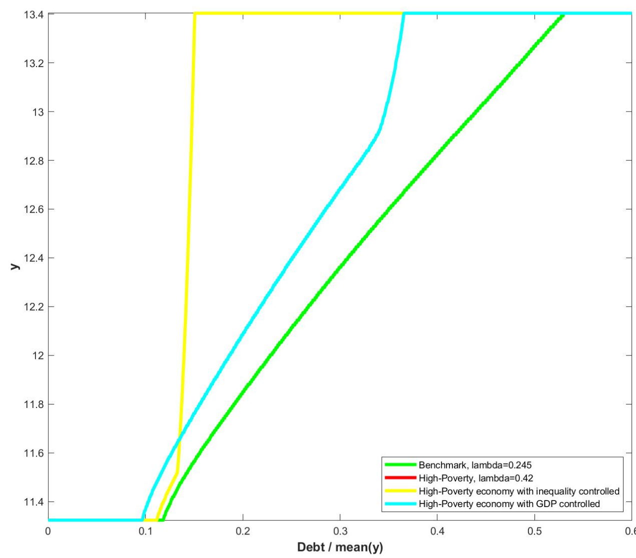

mean(y), so that mean GDP remains unchanged from its value in the Benchmark economy.

Figure 7 plots the default regions for the Benchmark and High-Poverty economy along with the economies where either GDP or Gini are held fixed when λ is increased. The cyan line for the GDP fixed economy sits exactly on top of the red line for the High-Poverty economy showing that the default risk of the two economies is identical. The default region of the Inequality fixed economy, shown by the yellow line is larger than that of the High-Poverty economy which suggests that the pure impact of the change in λ is even larger than when inequality is allowed to change.[11] We also note that the income of poor households in the Inequality fixed economy is higher than in either the Benchmark economy or the High-Poverty economy. This weakens the desire of the non-poor to participate in the tax-transfer scheme making the political constraint tighter. As a result, political defaults become more likely which result in a larger default region.

[11] We conjecture that one reason that inequality seems to reduce default risk in our quantitative results is because the tax-transfer system does not impose output losses on the economy. The effect of distortionary taxes on default risk in the presence of inequality is already explored in Jeon and Kabukcuoglu (2018) and has been turned off in our work.

5.2 The role of poverty parameters

In this subsection we will briefly discuss the impact of varying two key parameters related to poverty: p and α. The former, governs how likely it is for a non-poor household to become poor in any given period, while the latter governs the income of poor relative to non-poor households. In Figure 8, we display the default region of the Benchmark economy as well as the Low-p economy. The only difference between the two is that the probability of becoming poor is half of the Benchmark economy (0.03 vs. 0.06) in the Low-p economy. A glance at the figure immediately reveals that the Low-p economy has a much larger default region. This occurs because the political constraint is much more binding in the Low-p economy, leading to political defaults in states where none would occur in the Benchmark economy. An example of this can be seen in Figure 9 where the solid purple line hits zero while the green line (Benchmark economy) does not. This increased prevalence of political defaults is easy to understand. A lower p implies that an income shock is less likely to push a household into poverty and this makes the insurance scheme provided by tax payments less attractive than in the Benchmark economy. Of course, the higher default probability raises borrowing costs which makes it even more likely that transfers will be low relative to taxes and this worsens the desirability of the tax-transfer scheme.

Figure 8: Default regions with varying p

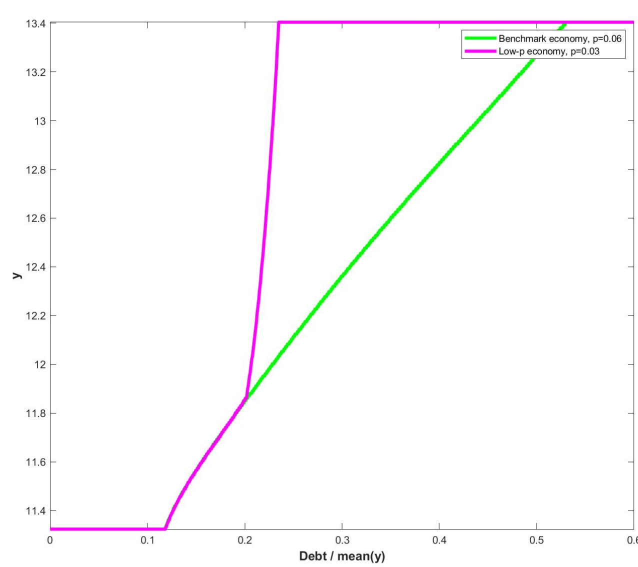

Turning attention to the impact of changing α, we draw the readers attention to Figure 7 again. A comparison of the default region of the High-Poverty economy (cyan and red overlap) to the default region of the inequality fixed economy (yellow) displays the pure impact of raising α while keeping λ at the level of the High-Poverty economy. It is clear from this comparison that an increase in α makes political defaults more likely because poverty is no longer as 'bad' as in the High-Poverty economy since the income gap is lower. As a result, non-poor households will demand a higher transfer to tax ratio in order to participate in the insurance scheme.

Figure 9: Political constraint with varying p

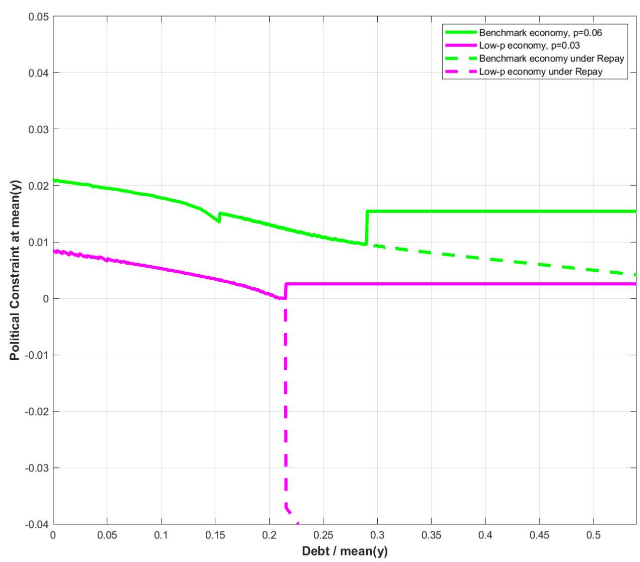

5.3 Welfare

The worse borrowing terms faced by the High-Poverty economy come with welfare losses relative to the Benchmark economy. Figure 10 reveals that the transfers received by poor households in the High-Poverty economy (red line) are much lower than in the Benchmark economy (green line). Figure 11 shows how this translates into large consumption differences for the poor in the two economies while Figure 12 plots the welfare calculated using the social weights used by the government (see equation 12). Recall that these weights are held constant in both economies. All three figures plot the respective variables for different debt levels while holding y fixed at it's mean level.

Figure 10: Transfers to poor households

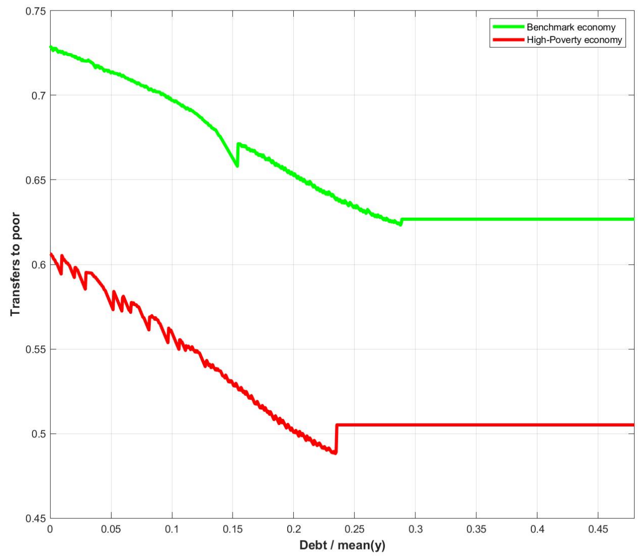

Figure 11: Consumption of poor households

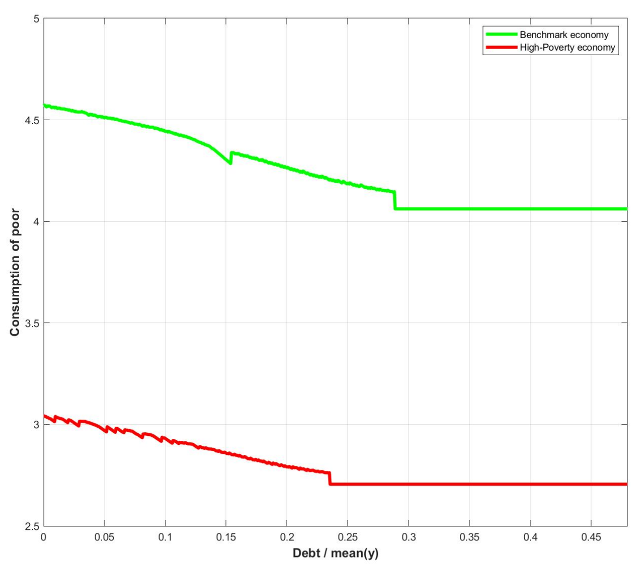

Figure 12: Welfare at average GDP

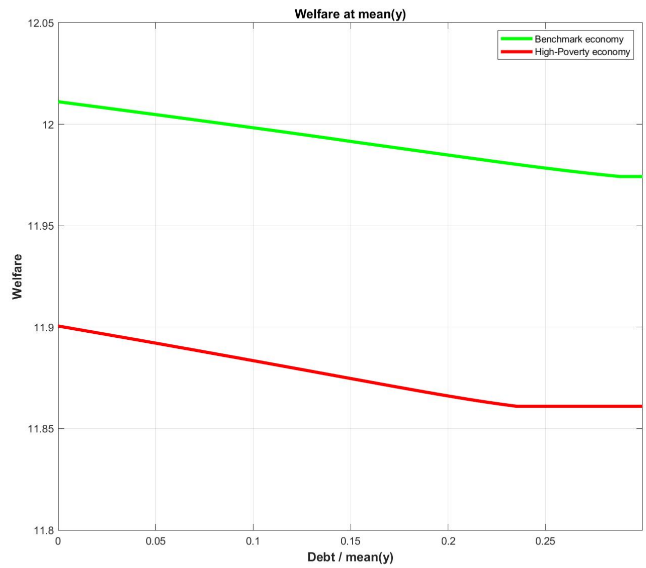

5.4 Decomposing the effect of credit terms

As we saw earlier, the High-Poverty economy suffers from heightened default risk both due to the worse fiscal situation in the country due to higher poverty and the greater risk of political defaults. This additional default risk leads to worse credit terms as exemplified by Figure 3. In order to decompose the additional default risk into the economy specific local factors and the impact of the external borrowing terms which respond to these factors, we solve for the equilibrium choices of the government in the High-Poverty economy using the spread menu of the Benchmark economy. We will refer to this economy as the Counterfactual economy. Any resulting change in the default region in the counterfactual economy relative to the High-Poverty economy will isolate the effect of the worse credit terms on inducing additional default risk. Similarly comparing the counterfactual economy to the Benchmark economy will display the pure impact of λ. We can think of the Counterfactual economy as a situation where a poor nation is offered concessional borrowing terms by an international aid agency.

Figure 13 presents the default regions associated with the three economies: the Benchmark economy, the original High-Poverty economy and the High-Poverty economy projected with Benchmark economy's bond prices, ie. the Counterfactual economy. It is evident from the

Figure 13: Default regions of High-Poverty economy: decomposition

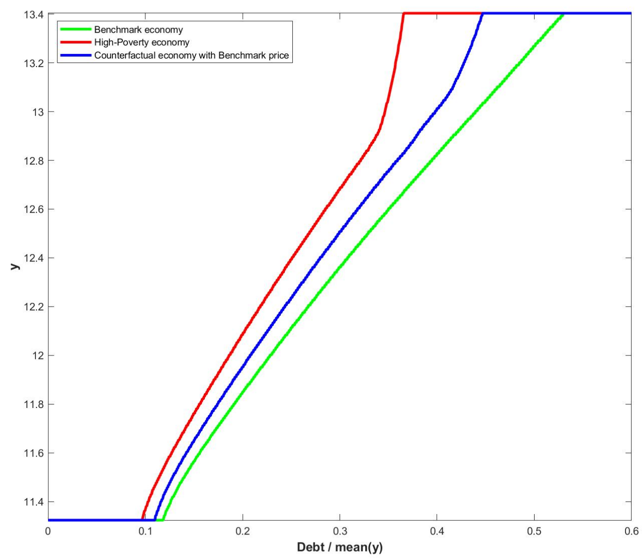

figure that the income level needed to avoid default at any given level of debt is significantly lower for the counterfactual economy than the actual High-Poverty economy. This gives rise to a smaller default region for the counterfactual economy. The resulting decrease in default risk for the Counterfactual economy can be attributed to the improvement in borrowing terms.

Figure 14 shows another immediate implication of the worse borrowing terms faced by the High-Poverty economy. Again it is evident that the harsher credit terms faced by the original High-Poverty economy forced them to choose a lower level of debt for any current debt level than the counterfactual economy. Indeed the endogenous debt limit above which debt is no longer sustainable is also much higher for the counterfactual economy than the High-Poverty economy.

6 Conclusions

International data suggests that there is a positive correlation between the proportion of extremely poor households in an economy and the average level around which its sovereign spread fluctuates. Even after controlling for a number of country specific and time varying factors, we find that a country with 40% of their population below the poverty line is asso-

Figure 14: Debt choices

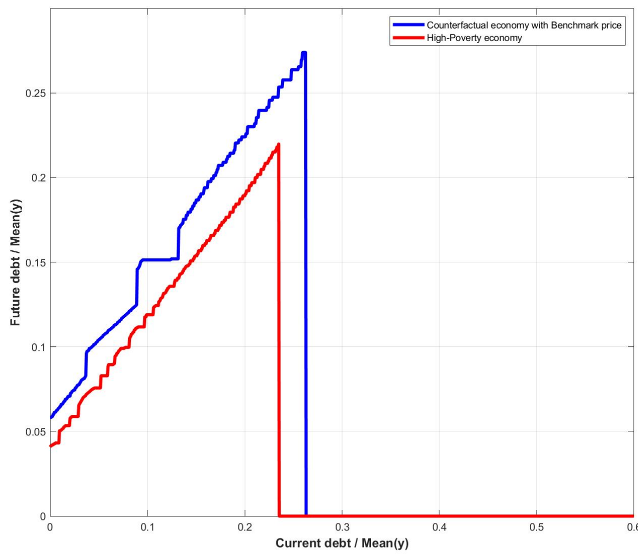

ciated with average spreads that are more than 120 basis points higher than a country with only 10% of population below the poverty line. We build a sovereign default model with poor and non-poor households and a government that wishes to run a social safety net using a tax-transfer scheme. The government must ensure that proposed taxes on the non-poor leave these households willing to participate in the insurance against income shocks provided by the scheme. While international borrowing allows the government to limit taxes relative to transfers when debt is low relative to aggregate income, when either debt is high or income is low, too high a proportion of tax revenue may need to flow out of the country to repay debt leading to two types of default. Regular or non-political defaults occur when the political constraint is satisfied but the government chooses to default in order to free up fiscal space to achieve the desired consumption levels of the two types of households. Political defaults occur when the government realizes that the unconstrained optimal tax-transfer levels it wants to achieve with repayment of debt would be unacceptable to households. As a result it chooses to default. The tightness of the political constraint is an equilibrium function of the proportion of poor households in an economy, their income relative to the non-poor as well as the probability of transitioning into poverty. The higher the proportion of poor, or the lower the risk of poverty, the more default risk in the economy. Moreover, the smaller the gap between poor and non-poor, the more likely political defaults are to occur since the insurance provided seems less valuable in the face of an income shock. Variants of the model generate quantitatively similar correlations between the proportion of poor households and spreads as seen in international data.

The higher default risk associated with a higher proportion of poor households implies much worse credit terms from international lenders. This has important implications for the economies in question – they pay higher interest rates and carry less debt which, in turn, leads to lower transfers and consumption for the poor and lower overall welfare. We decompose the additional default risk associated with the High-Poverty economy into two parts. The first is intrinsic to the higher poverty while the second is a consequence of the interaction of the worse borrowing terms with fiscal policy and the political constraint. If the High-Poverty economy could borrow from an international aid agency at the same terms as the Benchmark economy, external borrowing would be much higher and defaults would be much less likely to occur.

7 Appendix

7.1 Data sources and variable definitions

7.1.1 Income and poverty dynamics of South Africa

We make use of the Household income variable in the public-release dataset of the National Income Dynamics Study (NIDS). NIDS is a panel survey based on the livelihoods of households and individuals over time. The study began in 2008 and continues to be carried out every two years with a nationally representative sample of over 28000 individuals in 7300 households across the country. Our estimation is based on all existing waves which covers survey periods ranging from 2008 to 2017. The household income variable as reported in the NIDS is an aggregation of all income sources received by households on a monthly basis, net of taxes.

7.1.2 Other macroeconomic variables

1. Spread: We obtained spread data from Cesar Sosa-Padilla. It is extracted from JP Morgan's EMBI Global Index which consists of weekly observations ranging from the

first week of 1995 to May 29, 2015 for 53 countries. We converted them to annual frequency and computed annual averages for each country.

2. Poverty: As a measure of poverty, we use the average poverty head count ratio at 1.90 dollars a day which identifies the share of a country's population with income less than the poverty line.

3. Bribe rate: This represents the 'value of gift expected to secure a government contract' taken from the World Bank Enterprise Survey.

4. GDP: Extracted from the World Bank's World Development Indicator database.

5. Debt-to-GDP ratio: We use total external debt stocks (current US$) extracted from the World Bank's Global Development Finance database divided by GDP (current US$) to compute the debt-to-gdp ratio.

6. 10-year US Tbill and US yield curve: Taken from the Federal Bank of St. Louis FRED database.

7. Gini coefficient: The Standardized World Income Inequality database.

8. Default events: Taken from the online appendix of Catão and Mano (2017)

9. Social benefit/GDP: This represents the public social expenditure as a percentage of GDP, extracted from the Government Financial Statistics, IMF.

7.2 List of countries included in the Empirical analysis

Argentina, Belarus, Belize, Brazil, Bulgaria, Chile, China, Columbia, Cote D'Ivoire, Dominican Republic, Ecuador, Egypt, El Savador, Gabon, Georgia, Ghana, Indonesia, Iraq, Jamaica, Kazakhstan, Lithuania, Mexico, Morocco, Namibia, Nigeria, Pakistan, Panama, Peru, Philippines, Russian Federation, Serbia, South Africa, Sri Lanka, Tunisia, Turkey, Venezuela, RB and Vietnam

7.3 Computational algorithm

We solve the benevolent government's problem using value function iteration. The aggregate output shock Y follows an AR(1) process and is discretized using 201 equally spaced grid points following Tauchen's method. Similarly, a discretized state space for debt d consisting of 1000 equally spaced grid points is used. Upon realization of output shock Y, the government updates the repayment value and default value , and decides whether to repay or default and how much to borrow d'. Next, the government chooses an optimal tax-transfer scheme τ, τ^p given that the political constraint is satisfied. Any fiscal tax-transfer and debt policy that violates the political constraint is discarded from the value function iteration process by assigning a sufficiently low value to consumption.

The model is solved using the following algorithm:

1. Create grids and discretize the state space for debt d with 1000 grid points, and for output y with 201 grid points.

2. Start with a guess for the bond price schedule $q^o(d, y)$ for all d' and y.

3. Given the bond price schedule, solve for the debt choice d' using value function iteration. The optimal choice for consumption of the two household types $c(d, y)$ and $c^p(d, y)$ follow from the government's budget constraint and the optimality condition 9. The taxes and transfers $\tau(d, y)$ and $\tau^p(d, y)$ are then determined from the household budget constraints.

4. Check that these optimal choices c, c^p, d satisfy the political constraint; if not, reject this d'

5. Update the repayment value $V^c(d, y)$ and default value $V^d(d, y)$ for each iteration.

6. Based on the updated value function, update the defaulting decision, bond price schedule $q(d', y)$ and the optimal value function of the government $V^0(d, y)$.

7. Iterate on the value function until the distance between the updated value function and the one from last iteration for a given $q^0(d, y)$.

8. Finally compute new bond price schedule $q^1(d, y)$ based on the default sets and repayment sets and keep on iterating until the distance between $q^1(d, y)$ and $q^0(d, y)$ convergences to the tolerance level.

References

Alamgir, F., Cotoc, J., Johri, A., 2022. The bribe rate and long run differences in sovereign borrowing costs. Available at SSRN 3767115 .

Alessandria, G., Bai, Y., Deng, M., 2020. Migration and sovereign default risk. Journal of Monetary Economics 113, 1–22.

Amador, M., 2004. A Political Model of Sovereign Debt Repayment. Technical Report. Society for Economic Dynamics.

Andreasen, E., Sandleris, G., Van der Ghote, A., 2019. The political economy of sovereign defaults. Journal of Monetary Economics 104, 23–36.

Arawatari, R., Ono, T., 2017. Inequality and public debt: A positive analysis. Review of International Economics 25, 1155–1173.

Arellano, C., 2008. Default risk and income fluctuations in emerging economies. The American Economic Review 98, 690–712.

Arellano, C., Ramanarayanan, A., 2012. Default and the Maturity Structure in Sovereign Bonds. Journal of Political Economy 120, 187–232.

Azzimonti, M., De Francisco, E., Quadrini, V., 2014. Financial globalization, inequality, and the rising public debt. American Economic Review 104, 2267–2302.

Bai, Y., Kim, S.T., Mihalache, G., 2017. The payment schedule of sovereign debt. Economics Letters 161, 19–23.

Bi, H., 2012. Sovereign default risk premia, fiscal limits, and fiscal policy. European Economic Review 56, 389–410.

Bianchi, J., Ottonello, P., Presno, I., 2019. Fiscal Stimulus under Sovereign Risk. Working Paper 26307. National Bureau of Economic Research.

Brophy, T., Branson, N., Daniels, R.C., Leibbrandt, M., Mlatsheni, C., Woolard, I., 2018. National income dynamics study panel user manual. Technical Note Release .

Catão, L.A., Mano, R.C., 2017. Default premium. Journal of International Economics 107, 91–110.

Chatterjee, S., Eyigungor, B., 2012. Maturity, indebtedness, and default risk. American Economic Review 102, 2674–99.

Chatterjee, S., Eyigungor, B., 2019. Endogenous political turnover and fluctuations in sovereign default risk. Journal of International Economics 117, 37–50.

Cotoc, I., Johri, A., Sosa-Padilla, C., 2021. Sovereign Spreads and the Political Leaning of Nations. Working Paper 29197. National Bureau of Economic Research.

Cuadra, G., Sanchez, J., Sapriza, H., 2010. Fiscal Policy and Default Risk in Emerging Markets. Review of Economic Dynamics 13, 452–469.

Cuadra, G., Sapriza, H., 2008. Sovereign default, interest rates and political uncertainty in emerging markets. Journal of International Economics 76, 78–88.

Deng, M., 2021. Inequality, Taxation, and Sovereign Default Risk. Discussion Papers. Department of Economics, Simon Fraser University.

D'Erasmo, P., Mendoza, E.G., 2016. Distributional incentives in an equilibrium model of domestic sovereign default. Journal of the European Economic Association 14, 7–44.

D'Erasmo, P., Mendoza, E.G., 2021. History remembered: Optimal sovereign default on domestic and external debt. Journal of Monetary Economics 117, 969–989.

Eaton, J., Gersovitz, M., 1981. Debt with potential repudiation: Theoretical and empirical analysis. The Review of Economic Studies 48, 289–309.

Hatchondo, J.C., Martinez, L., 2009. Long-duration bonds and sovereign defaults. Journal of international Economics 79, 117–125.

Hatchondo, J.C., Martinez, L., 2010. The politics of sovereign defaults. FRB Richmond Economic Quarterly 96, 291–317.

Jeon, K., Kabukcuoglu, Z., 2018. Income inequality and sovereign default. Journal of Economic Dynamics and Control 95, 211–232.

Johri, A., Khan, S., Sosa-Padilla, C., 2022. Interest rate uncertainty and sovereign default risk. Journal of International Economics 139.

Mendoza, E.G., Yue, V.Z., 2012. A general equilibrium model of sovereign default and business cycles. The Quarterly Journal of Economics 127, 889–946.

Yue, V.Z., 2010. Sovereign default and debt renegotiation. Journal of International Economics 80, 176–187.