IMF Working Paper

Assessing IMF Lending: A Model of Sample Selection

by Nicolas Mäder, Jean-Guillaume Poulain, and Julien Reynaud

INTERNATIONAL MONETARY FUND

IMF Working Paper

European and Fiscal Affairs Departments

Assessing IMF Lending: A Model of Sample Selection

Prepared by Nicolas Mäder, Jean-Guillaume Poulain, and Julien Reynaud

Authorized for distribution by Donal McGettigan and Catherine Pattillo

July 2019

Abstract

Extending previous work on the determinants of IMF lending in an interconnected world, we introduce a model of sample selection in which both selection and size dimensions of individual IMF arrangements are presented within a unified econometric framework. We allow for unobserved heterogeneity to create an additional channel for sample selection at the country level. The results suggest that higher external financing needs, larger exchange rate depreciation, lower GDP growth, as well as deteriorated global financial conditions, are associated with larger individual IMF arrangement sizes. Using the estimated parameters, Monte Carlo simulation of a wide spectrum of global shock scenarios suggest that the distribution of potential aggregate IMF lending exhibits a substantial right tail. Our approach may provide an insightful input to broader policy discussions on the adequacy of the IMF resources.

1 Introduction

The global financial crisis (GFC) has underscored the need for an adequate global financial safety net (GFSN). Higher trade and financial integration, which has resulted in increased interconnectedness, carry benefits but have also increased the risk of systemic liquidity crises. Without adequate and prompt liquidity provision, even member countries with stronger fundamentals may become vulnerable as crises can propagate quickly. Accordingly, the GFSN has grown significantly since the GFC and has become more multi-layered, reflecting the continued accumulation of reserves and the expansion of official bilateral and multilateral arrangements. Yet, while coverage from Regional Financing Arrangements and Bilateral Swap Lines remains uneven and limited, the IMF continues to play a central role in the GFSN. Furthermore, besides its near-universal membership, what makes the Fund unique is its predictable and reliable financing through an array of lending instruments that continue to improve to meet evolving members' needs [2].

Therefore, a fundamental question when discussing global financial architecture is how much resources the IMF needs to adequately fulfill its role. Modeling potential needs for IMF financial assistance can represent a useful input for assessing the adequacy of Fund resources. This paper contributes to this strand of the analysis by expanding previous work by Poulain and Reynaud [11] on the determinants of IMF lending to model not only which member countries may require Fund financial assistance, but also the potential size of such assistance in a unified framework.

One important methodological challenge faced by researchers investigating IMF lending is the issue of sample selection bias. Sample selection bias arises when the size of arrangements are only observed for a restricted and non-random sample of observations and/or member countries. Specifically, countries approaching the IMF do so because they face economic difficulties such that a higher ex ante likelihood of a program may be mirrored by a higher program size per GDP in expectations. For instance, increased exposure to shocks and a related decreased ability to implement appropriate macroeconomic policies likely increases

the potential use of Fund resources. Failing to account for sample selection at the country level, in addition to the traditional idiosyncratic channel, could thus lead to flawed conclusions regarding the determinants and size of financial assistance. In addition, the selection bias can be deepened by the fact that some countries are entering arrangements with the IMF for signaling purposes (cf. precautionary Stand-by Arrangements and Flexible Credit Lines).

Accordingly, our main methodological contribution is, using panel data, to exploit correlated country effects along the selection and size dimensions, thereby directly allowing for an additional, permanent channel of sample selection. As explained in Section 3, we do so by extending Heckman's selection correction model in the spirit of Mundlak [9] who had pioneeringly modeled unobserved heterogeneity by linearly projecting it onto the cross-sectional unit's observables, a practice now referred to by Wooldridge [13] as correlated random effects[1]. Using the estimated parameters, the model is then used to simulate a wide spectrum of global crises of different intensity and breadth. The results can therefore provide a matrix of potential calls on Fund resources, which may provide some useful insights for assessing the adequacy of Fund resources.

The paper is organized as follows: the next section presents some stylized facts on the evolution and distribution of IMF arrangements size over time. Section 3 details the empirical methodology used to model IMF arrangements. Section 4 provides results of the estimation and discusses the fit of the model. In section 5, we use simulate potential aggregate IMF lending under a range of global shock scenarios. Section 6 concludes.

---

[1]Our analysis focuses on non-concessional lending from General Resource Account (GRA) resources only. Programs financed by the Poverty Reduction and Growth Trust (PRGT) are outside the scope of this analysis. The variables discussed in the paper are not the actual criteria that the Fund uses when deciding on approving a member's access to its resources. Rather, the paper explores the possible indicators of *probability* of country's requiring to the Fund's financing. Fund policies governing the access to Fund financing include the strength of the member's program, member's balance of payments need and capacity to repay the Fund.

2   Stylized facts

The average size of IMF arrangements has increased over time, both in nominal terms and in percent of a requesting member's GDP. But the distribution of arrangement size has also become wider over time. Figures 1a and 1b illustrate that the increase in size is most pronounced in the upper tail of the distribution. The tail of the arrangement size distribution has also become fatter toward larger arrangements. The figures also show how the distribution of new GRA arrangements has evolved over time. From being relatively compressed close to the average size in earlier decades of the IMF history, the distribution has gradually become more dispersed with an increasing probability of observing arrangements in the tail. While only 21 percent of all arrangements were larger than the average during 1957-1979, the corresponding number was 34 percent during 2008-16. Even if this key feature of the distribution has been aggravated since the GFC, the trend started in the 1990s with the first large capital account crisis arrangements. The exceptional access framework that was approved in 2002 was a clear recognition by the IMF of the increased importance of capital account crises and related need for higher access. More broadly, this took place against the backdrop of increased interconnectedness. Figure 2 depicts the evolution of average arrangement size and external financing needs since 1990.

Figure 1a - Distribution of IMF arrangement sizes - Histogram
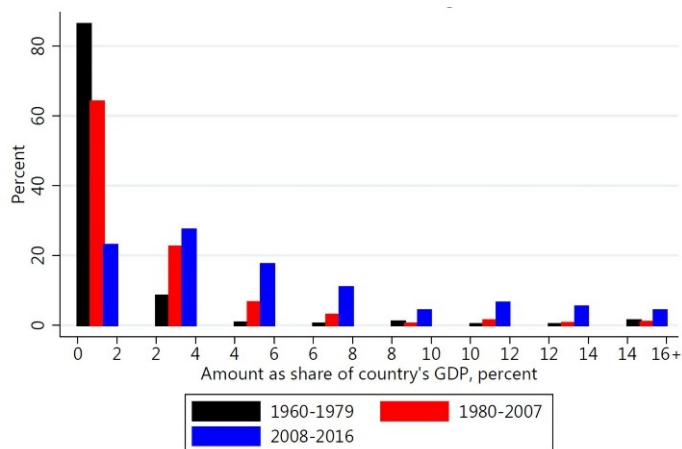
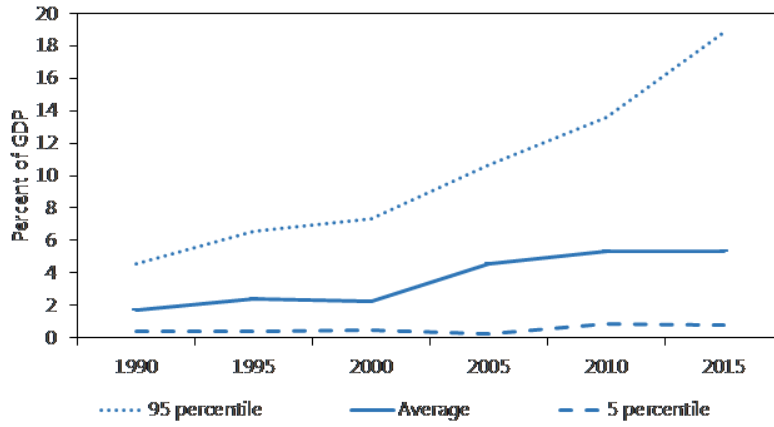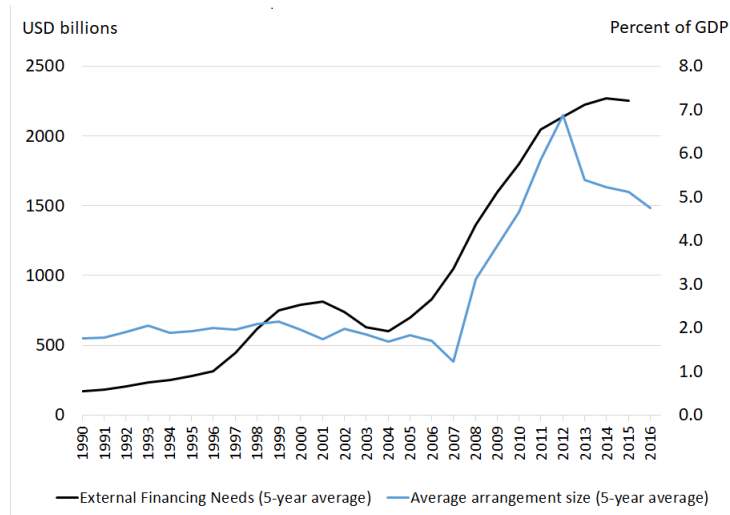Figure 1b - Distribution of IMF arrangement sizes

Figure 2 - Evolution of external financing needs and average arrangement size since 1990

On an aggregate level, correlation between IMF credit outstanding and global risk aversion (proxied by the VIX) has been high. Poulain and Reynaud [11] discuss and quantify this relationship in greater detail, together with a discussion of factors influencing IMF lending.

Figure 3 - Evolution of IMF credit outstanding and risk aversion since 1990
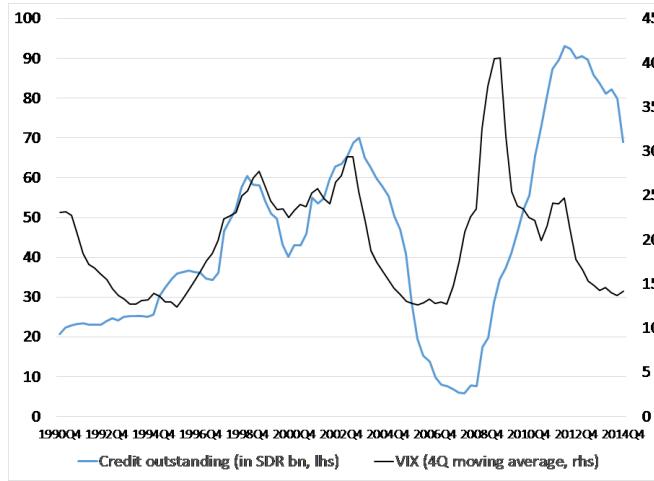
3   Empirical methodology

The aim of this section is to construct and estimate an empirical *forecasting model* for aggregate use of Fund resources. As underscored in Poulain and Reynaud [11], past studies trying to disentangle demand and supply factors have been criticized. IMF lending differs from typical market mechanisms where supply and demand are balanced by variations in prices, not least because the Fund's rate of charge is exogenously set. Accordingly, our analysis focuses on joint determinants of access to resources.

The core innovation of our approach is that combining unobserved heterogeneity with full information joint maximum likelihood (of selection and size) allows for sample selection at both the idiosyncratic and the country level. In addition to allowing for selection and size to correlate contemporaneously, countries with a permanently higher likelihood of entering into an arrangement with the IMF may also have permanently larger programs. Since sample selection models are most popular within the inherently causal field of applied microeconomics, it is important to stress that our goal is not to elicit causality[2]. Therefore, imposing exclusion restrictions only requires a zero conditional correlation, but not exogeneity. Before delving into the empirical estimation, we first precisely define our target

---

[2]The difference can be illustrated by way of the following two questions: 1. Forecasting: "What can we infer about Y if we were to observe X?" 2. Applied Micro: "How does an exogenous change in X affect Y?"

quantity (Pearl [10]):

1. **Define**: Our target quantity is given by the conditional density of aggregate Fund resources given a set of covariates. Crucially, since the IMF is not capable to instantaneously influence the covariates, notice that our parameters of interest explicitly do not represent causal effects, but partial correlations[3].

2. **Assume**: Since we are not estimating causal effects, no further *causal* assumptions are required. When choosing a particular specification including the exclusion restrictions[4], we thus rely on the model selection criteria AIC and BIC as frequently used in time series econometrics.

3. **Identify**: While our parameters of interest are technically identified through the model's nonlinearity, credibility of identification is increasing in the number of exclusion restrictions which, as indicated below, are chosen according to the AIC and BIC metrics.

4. **Estimate.** Derivation of the likelihood and its maximization using our panel data are discussed in appendix A.

Consider a Heckman type model in which the size of individual IMF arrangement (as a percent of a member country's own GDP) y_{it} is observed if and only if some indicator variable s_{it} is equal to unity.[5] Allowing for *sample selection*, e.g. that some unobserved variable affects both the indicator s_{it} as well as the outcome y_{it}, we assume,

---

[3]In essence, given some model $y = \beta x + u$ with $u \sim \mathcal{N}(0, \sigma^2)$, an analysis of the conditional density $y|x$ may yield meaningful insights, even if x does not cause y.

[4]In our context, interpreting the exclusion restrictions is complicated by the fact that our parameters do not signify causal effects, but partial correlations. In particular, recall that while correlation famously does not imply causation, the absence of causation similarly does not imply a zero correlation if the latter is conditional (Berkson's paradox [1]).

[5]Crucially, non-negativity of program size is addressed via the exponential form of the size equation. Notice that data is neither censored nor truncated. It is not censored because all program size data that is observed is observed accurately. It is not truncated because all dimensions of a country other than size are observed. However, we do not know what the counterfactual program size would have been, had a specific country been selected, a classic sample selection problem akin to the canonical wage offer example.

$$
y_{it}^o = \begin{cases} y_{it} & \text{if } s_{it} = 1 \\[2ex] \cdot & \text{if } s_{it} = 0 \end{cases}
$$

$$
y_{it} = \exp(x_{i1t-1}\beta_1 + x_{i2t}\beta_2 + c_i + u_{it})
$$

$$
s_{it} = 1(z_{i1t-1}\gamma_1 + z_{i2t}\gamma_2 + d_i + v_{it} \geq 0)
$$

where the exponential form is motivated by a positivity requirement on y_{it}, (c_i, d_i) denote potentially correlated unobserved effects, and $x_{it} \equiv (x_{i1t-1}, x_{i2t})$, $z_{it} \equiv (z_{i1t-1}, z_{i2t})$ are covariates either entering with a lag (for country-specific variables) or contemporaneously (for global variables)[6]. Letting $Z_{it} = [Z_{i1t}, Z_{i2t}] \equiv [(x_{i1t}, z_{i1t}), (x_{i2t}, z_{i2t})]$, we assume $u_{it}, v_{it} | Z_{i1t-1}, Z_{i2t}$ to be joint normal[7]. It is well understood that, given the setting described so far, idiosyncratic sample selection arises if $\text{corr}(u_{it}, v_{it}) \equiv \rho_{uv}$ is different from zero, in which case the econometrician must address, in the respective moment conditions or the likelihood function, that the data represent a non-random subsample of the population [5]. Letting S denote a scenario, or a particular realization of the random variable Z_{2t}, we aim to recover the conditional distribution over aggregate resources,

$$
f_{y|S} \equiv P\left(\sum_{i=1}^{n} y_{it} s_{it} = y | Z_{1t-1}, Z_{2t} = Z^S\right)
$$

which requires an estimate of the conditional joint density $P(y_{it}, s_{it} | Z_{it-1}, Z_{2t} = Z^S)$. Importantly, Z_{2t} may include aggregate variables, which are constant across countries and thus generate a channel for stochastic co-dependencies.

The inverse Mills ratio is approximately linear over a majority of its domain such that any causal studies using the original Heckman model [5] must provide a *causal* exclusion

---

[6]We follow the literature by incorporating all country specific variables using lags, thereby mitigating concerns arising from potential contemporaneous correlations between the covariates and the error.

[7]Without loss of generality, we impose $\sigma_v = 1$ and thus have $u_{it} = \rho_{uv}\sigma_u v_{it} + \varepsilon_{it}^u$ with $\varepsilon_{it}^u \sim \mathcal{N}(0, (1 - \rho_{uv}^2)\sigma_u^2)$. While it is technically possible to instead opt for a logit specification with v_{it} being distributed logistically, it would be important to note that $v_{it} | s_{it} = 1$ being truncated logistic implies that $\mathbb{E}[u_{it} | s_{it} = 1]$ is not equal to the inverse Mills ratio $\lambda(-z_{it}\gamma)$. See Xu et al. [14] for a discussion of truncated logistic distributions.

restriction - a variable that *causes* selection into an IMF arrangement but does not affect the outcome variable of interest. Without such a variable, multicollinearity renders identification of causal parameters imprecise at best or impossible at worst (Lang [7]; Wooldridge [13]; Little and Rubin [8]). In spite of not explicitly relying on the canonical inverse Mills ratio, FIML estimators may also be sensitive to multicollinearity in the sense that the likelihood is virtually flat in certain regions of the parameter space. Therefore, the higher the number of excluded selection variables, the more credible is identification along the size dimension. In our context, attributing economic interpretation and imposing exclusion restrictions is complicated by the fact that our parameters of interest do not represent causal effects, but partial correlations[8]. We therefore follow the forecasting literature by atheoretically assessing relative specification performance by way of the Akaike and Bayesian information criteria.[9]

4   Results

We use Poulain and Reynaud's panel dataset of 92 advanced, emerging, and frontier market economies over the period 1992-2014, covering 119 GRA arrangements. In addition to borrowing their dataset, we build upon their model which serves as a natural starting point as we iteratively determine our model's specification. Favored by the information criteria, the following specification summarized in Table 1 serves as our benchmark model.[10] Notice that EFN per GDP serves as our only "Mundlak variable", making an appearance along both dimensions.

---

[8]Consider a case in which we have successfully established $y_{it} \perp\!\!\!\perp Z_{ijt}$, but both variables cause some third variable Z_{ikt}. In this case, as initially pointed out by Berkson [1], $y_{it} \perp\!\!\!\perp Z_{ijt}|Z_{ikt}$ may fail as it is not implied by $y_{it} \perp\!\!\!\perp Z_{ijt}$.

[9]While both criteria were derived in an effort to compare the performance across a class of models, the Bayesian Information Criterion (BIC) penalizes overparameterization more harshly than does the Akaike Information Criterion AIC. The latter's corrected version ($AICc$) further corrects the asymptotic statistic, which is similar but more general than a likelihood ratio test, for the finite size of any sample.

[10]See Poulain and Reynaud [11] as well as Appendix B for description of the data, sources, and countries in the sample. Note that in our context of joint selection and size, it is unsurprising that the employed information criteria favor a (selection) specification other than the one originally proposed in Poulain and Reynaud. Please refer to their paper for a discussion of the data and an explanation of the panel's variables.

Table 1 - Benchmark model

	θ	Meaning	$\hat{\theta}_{MLE}$	$|\hat{z}_H|$	$\hat{p}_H$
Size	β_1	EFN per GDP	3.77	4.60	0.000***
		GDP per capita	0.28	2.45	0.014**
		Exchange rate variation	-1.25	3.05	0.002***
		Growth	-0.04	2.38	0.017**
	β_2	**VIX**	0.04	3.59	0.000***
	β_3	EFN per GDP	(-)		
Selection	γ_1	EFN per GDP	2.53	3.21	0.001***
		GDP per capita	-0.42	3.92	0.000***
		Government Stability	-0.16	3.41	0.000***
		Interconnectedness	-0.33	2.02	0.043**
		Past Arrangement	0.13	2.78	0.005***
		Credit-to-GDP gap	0.01	2.46	0.014**
		Exchange rate variation	-0.92	2.21	0.027**
		Growth	-0.04	2.86	0.004***
	γ_2	**VIX**	0.04	3.18	0.001***
	γ_3	EFN per GDP	(-)		
Errors	ρ_{uv}	Sample selection parameter	0.49	2.38	0.017**
	σ_u	Shock variation	0.83	-	-
	σ_c	Effect variation	0.70	-	-
	σ_d	Effect variation	0.23	-	-
		Countries/observations	129	3741	
		Used	92	1533	
		Uncensored/censoreed	119	1414	
	IC	Akaike/Bayesian	918	1,025	
	R^2	Adjusted	0.30		

The coefficients of all selection variables match the signs found by Poulain and Reynaud [11]. Furthermore, the coefficients of variables entering the size equation also carry signs which are supported by economic intuition. Higher external financing needs, lower GDP growth, and increased levels of global risk aversion (proxied by the VIX) are mirrored by higher idiosyncratic arrangement size. Larger exchange rate depreciations are also associated with a larger arrangement size, consistent with the likely need to bolster foreign exchange reserves and restore confidence. The case of GDP per capita is interesting: its coefficient is negative in the selection equation but positive in the size equation. This may be interpreted as follows: the richer a country, the less likely it is to require Fund financing; but if it does, the extent of the crisis it is facing may be such that it would need a larger access to resources.

Our information criteria favor a specification in which government stability, interconnectedness, past arrangements, and the credit-to-GDP gap only enter along the selection, but not the size dimension. Of course, since both criteria are statistical artifacts unrelated to the economic phenomenon at hand, validity of these exclusion restrictions should be critically assessed using economic intuition. It is an acceptable proposition that some factors may influence the decision to lend (e.g. spillovers, government stability) but not the size of the arrangement, which is primarily driven by the balance of payment need of the member country. Furthermore, in our context, such an intuitive evaluation is convoluted by the fact that the introduced framework does not allow for separate identification of supply and demand, an issue already pointed to by Ghosh et al. [4]. For example, it is certainly conceivable for a potential covariate to constitute both a demand and a supply shifter without affecting observed outcomes in a statistically significant way.

- *Interconnectedness.* A priori, one may expect interconnectedness to enter with the same sign in both equations as the IMF aims to prevent spillovers (see Poulain and Reynaud [11]). However, notice that higher interconnectedness may be mirrored by a greater availability of private finance and thus smaller need for IMF resources if access to capital markets is not completely shut down.

- *Government stability.* Higher government stability is presumably mirrored by lower demand, but also less constrained supply.

- *Past arrangement.* Many member countries have a history of protracted financial engagement with the Fund, mirroring deeply-rooted economic problems. Conversely, Fund policies governing access to resources (exceptional access, conditionality) may constraint supply in some cases.

- *Credit-to-GDP gap.* A large credit-to-GDP gap has been identified as an early warning signal of banking crisis, which may increase demand for Fund resources, but also potentially reduce supply if adequate safeguards are not in place.

Since our Mundlak variable aims to capture *permanent* cross-country differences, intu-

10

itively likely most closely related to Reinhart, Rogoff, and Savastano's notion of "debt (in)tolerance" [12], its signs are to be interpreted with utmost caution. Accordingly, since it is very tempting to interpret them in an inaccurate "within country" sense, the table does not show the estimated parameter's value. However, our Mundlak variable carries the same sign across both selection and size dimensions, which interestingly indicates the existence of a sample selection channel at the country level: permanently higher likelihoods of selection are mirrored by larger arrangements per GDP. Additionally, there is limited evidence of idiosyncratic sample selection as measured by the estimate $\hat{\rho}_{uv}$.

Model fit - Selection

The accuracy of the selection may be tested by creating bins of ex ante selection probabilities, as fitted by our model, and computing the proportion of arrangements in each bin. Of course, if the predicted likelihoods are accurate and the number of observations in a particular bin is large, the corresponding proportion of actual arrangements must fall into the bin's probability interval.

Table 2 - Selection performance

Fitted probability	$[0, 5\%]$	$[5, 10\%]$	$[10, 15\%]$	$[15, 20\%]$	$[20, 25\%]$
Proportion with actual arrangement	1%	10%	12%	18%	22%
Fitted probability	$[25, 30\%]$	$[30, 35\%]$	$[35, 40\%]$	$[40, 45\%]$	$[45, 50\%]$
Proportion with actual arrangement	33%	47%	33%	46%	50%

Note that the number of observations in the higher probability bins is too small for the law of large numbers to apply such that the performance of our model should be viewed as satisfactory at worst.

Model fit - Size

The plots in Figure 4 depict both the log moment $\mathbb{E}[\log(y_{it})|s_{it} = 1]$, the one fitted by our maximum likelihood algorithm, and the level moment $\mathbb{E}[y_{it}|s_{it} = 1]$. Before analyzing the second and third rows, the latter of which is simply an enlarged version of the former, recall that even an untruncated lognormal random variable's expected value exceeds its median by a factor of proportionality $\exp(\sigma^2/2)$. While $y_{it}|s_{it} = 1$ is neither lognormal nor

truncated lognormal, we should nevertheless expect more mass to lie below the mean such that the number of dots is not equal on both sides of the black line. However, as the sum's density nevertheless closely resembles a normal distribution, $y_{it}|s_{it} = 1$ may approximately be thought of as a lognormal random variable. This intuition is confirmed by the residuals' negative mode and the large mass of residuals below zero depicted in figure 5. The two seemingly very large negative shocks experienced by the observations at the left tail refers to two post-GFC large loans that were to a large extent financed by non-IMF sources. When accounting for these sizable contributions, the two residuals exhibit a substantial shift to the right (as captured by the blue observations).
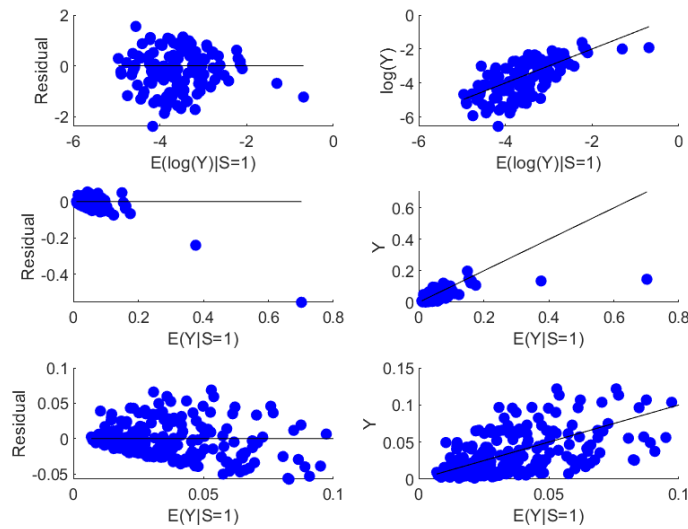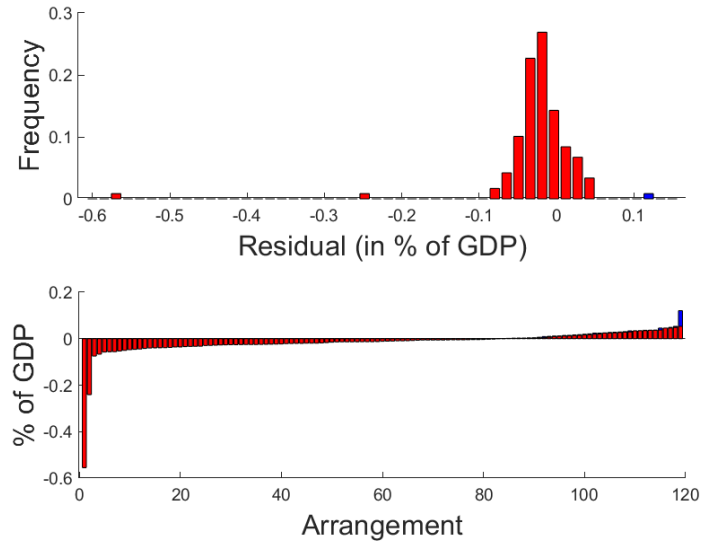
Figure 4 - Residual Plot

Figure 5 - Residual Distribution

5   Scenarios

In this section, we use the model described earlier to simulate potential aggregate IMF lending in the first post sample year under a wide range of shock scenarios, constructed as global volatility shocks, whereby outcomes are simulated for different VIX levels. It is further assumed that all other variables remain unchanged. In practice, a global volatility shock would likely be accompanied – among other things – by a drop in GDP growth and an increase in external financing needs, which would both increase the potential size of aggregate Fund financing. On the other hand, some non-modeled considerations could work in the opposite direction (e.g. stigma, conditionality).

We use two different approaches to simulate potential calls on Fund resources based on the above shocks. As rightfully pointed out by Ghosh et al. [4], future individual and aggregate access are not predetermined because they are functions of the random variables and the unobserved quantities. More rigorously, unbeknownst the true stochastic environment of the global economy as given by the probability space, the model may be simulated by either employing estimate densities or alternatively by postulating the realization of a specific event.

- In the first approach, the vectors (c, d, u_t, v_t) are simulated given their estimated joint density as implied by our model while Z_{2t} is held fixed to some realization Z_{2t}^1. This is conducted in style of a typical Monte Carlo simulation.

- In the second approach, a country is selected if its estimated selection probability exceeds a certain threshold, given some realization Z_{i2t}^2, where the size of each arrangement is pinned down by $u_{it}^2 = \hat{\rho}_{uv}\hat{\sigma}_u v_{it}^2 = -\hat{\rho}_{uv}\hat{\sigma}_u z_{it}\gamma$. This approach is characterized by a fixed ordinal ranking of country selection and is thus more intuitive to interprete.

Both scenarios generate a similar range of conceivable aggregate access and a substantial right tail. In particular, the distribution of aggregate access under the first approach is characterized by a sizeable wedge between the conditional expectation, a moment sometimes emphasized in *Value-at-Risk* analyses, and the upper quantiles. For illustrative purposes, let us look for instance at a global shock in which the VIX would be at 30 on average for a year (only slightly below the average level of 32 reached during the GFC).

- Under the first approach, potential needs are estimated to reach about SDR 280 billion at the 85th quantile, SDR 560 billion at the 95th quantile. Figure 6 illustrates that needs would be much larger for higher VIX and/or higher quantiles.

- Under the second approach, potential needs are estimated to reach about SDR 400 billion at the 7.5 percent threshold (equal to the frequency of events in our sample) and quickly increase as the threshold is reduced. At the 4 percent threshold, potential needs are estimated at SDR 825 billion.

Figure 6 - Potential aggregate IMF lending under the first approach
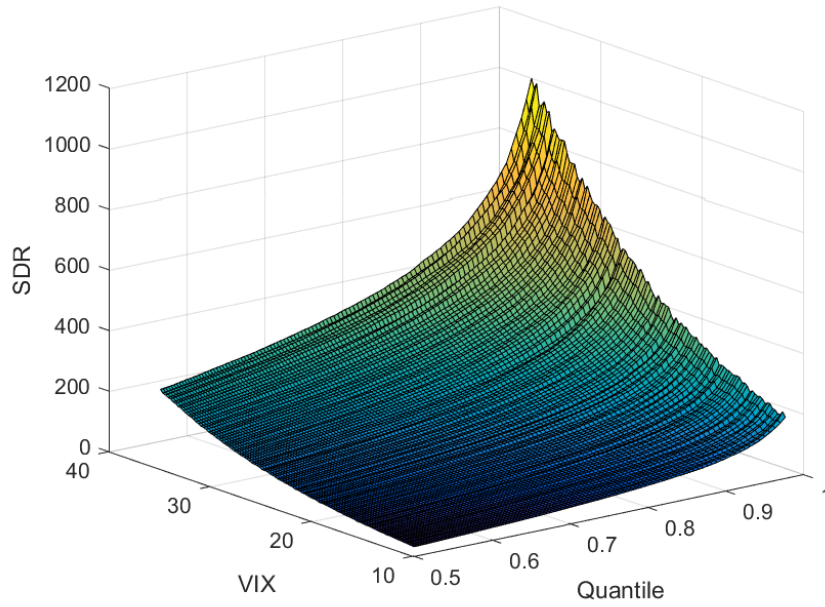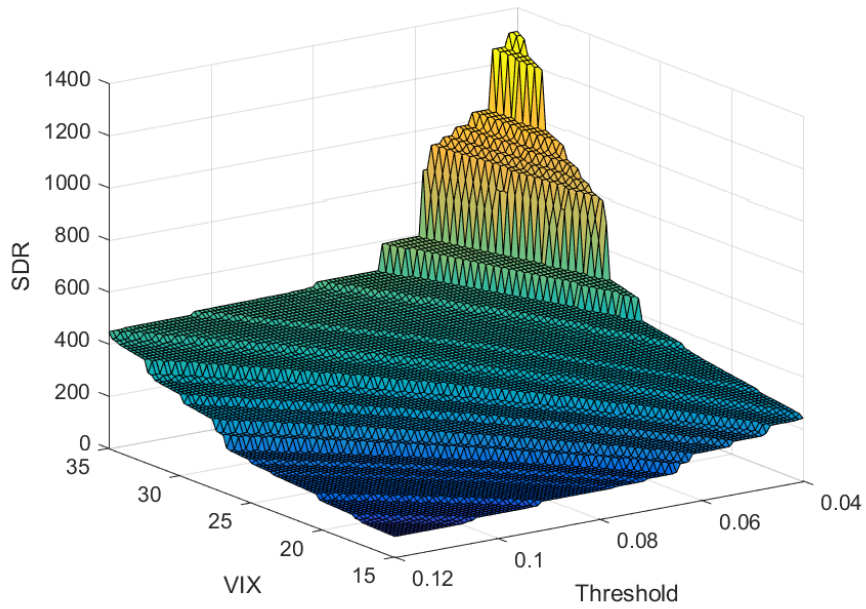(fix VIX, simulate errors)

Figure 7 - Potential aggregate IMF lending under the second approach
(fix VIX, fix errors)

6   Conclusion

In this paper, we provide a methodology to correct for an important empircal challenge in estimating the potential need for Fund financing, namely the issue of sample selection bias. We extend Heckman's selection correction model in the spirit of Mundlak [9] who had pioneeringly modeled unobserved heterogeneity by linearly projecting it onto the cross-sectional unit's observables. In other words, we exploit correlated country effects along the selection and size dimensions, thereby directly allowing for an additional, permanent channel of sample selection.

We use the selection model of Poulain and Reynaud [11] as a starting point for our estimation. In addition to all coefficients in the selection equation matching the signs found by Poulain and Reynaud [11], the coefficients in the size equation also carry signs which are supported by economic intuition. In particular, higher external financing needs, exchange rate depreciation, global risk aversion, and lower GDP growth are mirrored by higher idiosyncratic arrangement size.

We further use the estimated version of our model to simulate future potential needs for IMF resources under a wide range of global volatility shock scenarios. Our scenarios suggest that most crises could be accommodated with a lending capacity of about SDR 600-800 billion. However, the distribution of aggregate access exhibits a substantial right tail, and accordingly a more extreme crisis could conceivably lead to significantly larger needs. Yet, at current low levels of global risk aversion, the expected aggregate access would be much lower, consistent with the idea that the global financial cycle is a major driver of IMF lending cycles. The approach outlined in this paper may be one of many inputs for assessing the adequacy of Fund resources, bearing in mind that this complex issue is a matter of informed judgment, based on the risk tolerance of the international community as well as the political, governance, and financial considerations.

References

[1] Joseph Berkson. Limitations of the application of fourfold table analysis to hospital data. *Biometrics Bulletin*, 1946. 4, 8

[2] International Monetary Fund. Adequacy of the global financial safety net. Technical report, International Monetary Fund, 2016. 1

[3] Chamberlain Gary. Multivariate regression models for panel data. *Journal of Econometrics*, 1982. A

[4] Atish Ghosh, Manuela Goretti, Joshi Bikas, Thomas Alun, and Juan Zalduendo. Modeling aggregate use of IMF resources - Analytical approaches and medium-term projections. *IMF Staff Papers*, 2008. 4, 5

[5] James J. Heckman. Sample selection bias as a specification error. *Econometrica*, 1979. 3

[6] James J. Heckman. *Structural Analysis of Discrete Data and Econometric Applications*, chapter Dynamic Discrete Probability Models. The MIT Press, 1981. A

[7] Valentin Lang. The economics of the democratic deficit: The effect of IMF programs on inequality. *Discussion Paper Series, University of Heidelberg*, 2016. 3

[8] Roderick J.A. Little and Donald Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, 1987. 3

[9] Yair Mundlak. Pooling of time series and cross section data. *Econometrica*, 1978. 1, 6, A

[10] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 2009. 3

[11] Jean-Guillaume Poulain and Julien Reynaud. IMF lending in an interconnected world. *IMF Working Paper*, 2017. 1, 2, 3, 10, 4, 6, B

[12] Carmen M. Reinhart, Kenneth S. Rogoff, and Miguel A. Savastano. Debt intolerance. *Brookings Papers on Economic Activity*, 2003. 4

[13] Jeffrey M. Wooldridge. *Econometric analysis of cross section and panel data*. MIT Press, 2010. 1, 3, 11

[14] Feng Xu, Ron C. Mittelhammer, and Allen Torrell. Modeling nonnegativity via truncated logistic and normal distributions: An application to ranch land price analysis. *Journal of Agricultural and Resource Economics*, 1994. 7

A Appendix A - Likelihood

Rather than relying on Heckman's original two-step, contemporary estimation of heckit models is traditionally implemented via full information maximum likelihood (FIML)[11]. Particularly in face of unobserved heterogeneity, the likelihood approach is preferred because it allows for a straightforward way of integrating out the corresponding effects [6]. As far as the nature of the heterogeneity is concerned, we follow Wooldridge by interpreting a fixed effect's key feature not to be its non-randomness, but rather its non-zero correlation with the covariates[12]. Specifically, we proceed in the spirit of Chamberlain [3] who followed Mundlak [9] in linearly projecting unobserved heterogeneity onto the covariates, a practice now referred to by Wooldridge as correlated random effects (CRE). As indicated by its name, while technically falling under the contemporary fixed effects umbrella, a correlated random effect is best understood as a natural variation of a random effect. In spirit of the methodology presented by Heckman [6], the likelihood can then be constructed by first conditioning on and subsequently integating out the unobserved effects[13],

$$\mathcal{L}(y|Z;\theta) = \sum_{i=1}^{N} \log\left(\iint \exp\left(\sum_{t=1}^{T} l_{it}\right) dF(d_i) dF(\varepsilon_i^c)\right)$$

where l_{it} denotes the conditional log-likelihood of observing (or not observing) y_{it} given the covariates, the parameter vector θ, **and** the unobserved effects d_i and c_i[14],

---

[11]See Wooldridge [13] for a FIML treatment of Heckman's original model.

[12]While the term "fixed effect" may have originated from the view that the effect takes the form of a parameter, it is well understood that a literal interpretation thereof regularly leads to incidental parameter problems. Of course, if the effect were truly a parameter, the corresponding correlation would be zero by definition.

[13]While constituting a random effects setup, recall that our algorithm nests the *pooled* case, when $\sigma_c = \sigma_d = 0$, as well as a subset of *fixed effects* in which the latter are centered around a linear combination of the longitudinal means of a subset, the "Mundlak variables", of the observables x and z,

$$c_i = \alpha_x + \beta_3 \bar{x}_{i3} + \ddot{c}_i$$
$$d_i = \alpha_z + \beta_3 \bar{z}_{i3} + \ddot{d}_i$$

As the Mundlak variables $(\bar{x}_3, \bar{z}_3)$ are added as regressors, $(\alpha_x + \beta_x \bar{x}_{i3}, \alpha_z + \beta_z \bar{z}_{i3})$ may be interpreted as *known* unobserved heterogenity whereas $(\ddot{c}_i, \ddot{d}_i)$ denote the new random effects to be integrated out. While not necessary from a practical perspective, we further assume the unobserved effects $(\ddot{c}_i, \ddot{d}_i)$ to be conditionally uncorrelated.

[14]Again, joint normality implies $c_i = \rho_{cd}(\sigma_c/\sigma_d)d_i + \varepsilon_i^c$ with $\varepsilon_i^c \sim \mathcal{N}(0, (1-\rho_{cd}^2)\sigma_c^2)$ such that c_i is entirely determined given (d_i, ε_i^c).

$$l_{it} \equiv l(y_{it}|x_{it}, z_{it}, d_i, \varepsilon_i^c; \theta)$$

$$= 1(s_{it} = 0) \log\left[1 - \Phi(z_{it}\gamma + d_i)\right] +$$

$$1(s_{it} = 1)\left\{\log\left[\Phi\left(\frac{z_{it}\gamma + d_i + (\rho_{uv}/\sigma_u)(\log(y_{it}) - x_{it}\beta - \rho_{cd}(\sigma_c/\sigma_d)d_i - \varepsilon_i^c)}{(1 - \rho_{uv}^2)^{0.5}}\right)\right] + \right.$$

$$\left. \log\left[\phi\left(\frac{\log(y_{it}) - x_{it}\beta - \rho_{cd}(\sigma_c/\sigma_d)d_i - \varepsilon_i^c}{\sigma_u}\right)\right] - \log(\sigma_u)\right\}$$

and ϕ and Φ denote the probability and cumulative density functions of the standard normal density respectively. Before integrating out the unobserved effects[15], we must first evaluate the conditional level-likelihood of each country's time series, as captured by the interior sum $\sum_{t=1}^{T} l_{it}$. Aside from allowing for unobserved heterogeneity, the only difference between our approach and the full maximum likelihood estimator of Heckman's linear model is the logarithmic entry of y_{it}. Given a working likelihood function $\mathcal{L}(y|Z;\theta)$, we maximize[16] over the parameter space Θ,

$$\hat{\theta}_{MLE} = \text{argmax}_{\theta \in \Theta} \ \mathcal{L}(y|Z;\theta)$$

Inference is conducted using the asymptotic z-statistic, derived from the negative Hessian of $\mathcal{L}$ at $\hat{\theta}_{MLE}$.

---

[15]Independence and joint normality of (d_i, ε_i^c) allow for an implementation of Gauss-Hermite quadrature.
[16]We use `fminunc` in Matlab.

B Appendix B - Data

Table B1 - Data - Variables and Sources

Variable	Source	Explanation
New arrangement	IMF	Indicator equals one when new arrangement starts
New arrangement size	IMF	Size of new arrangement, percent of GDP
Past arrangement	IMF	5-year moving average of active arrangement indicator
EFN	WEO	External financing needs; see [11]
GDP Growth	WEO	In percent
GDP per capita	WEO	log of level
GDP	WEO	log of level
Credit-to-GDP gap	BIS; WDI; WEO	Deviation of credit-to-GDP from 5-year moving average
Exchange rate variaton	WEO	Variation of bilateral nominal exchange rate vs. USD
Government Stability	ICRG	Government unity, legislative strength, popular support
Interconnectedness	WEO; DOTS	See [11]
3-month interest rate	WEO	In percent
VIX	CBOE	Measure of risk aversion
Oil price	WEO	Deviation from 5-year moving average

Table B2 - Countries in the sample

Albania	Ecuador	Kazakhstan	Romania
Algeria	Egypt	Korea	Russia
Angola	El Salvador	Kuwait	Saudi Arabia
Argentina	Estonia	Latvia	Singapore
Armenia	Finland	Lebanon	Slovak Republic
Austria	France	Libya	Slovenia
Australia	Gabon	Luxembourg	South Africa
Azerbaijan	Germany	Malaysia	Spain
The Bahamas	Greece	Malta	Sri Lanka
Bahrain	Guatemala	Mexico	Suriname
Belarus	Guyana	Morocco	Sweden
Belgium	Hungary	Namibia	Switzerland
Botswana	Iceland	Netherlands	Syria
Brazil	India	New Zealand	Thailand
Canada	Indonesia	Norway	Trinidad and Tobago
China	Iran	Oman	Tunisia
Colombia	Iraq	Pakistan	Turkey
Costa Rica	Ireland	Panama	Ukraine
Croatia	Israel	Peru	United Arab Emirates
Cyprus	Italy	Philippines	United Kingdon
Czech Republic	Jamaica	Poland	Uruguay
Denmark	Japan	Portugal	Venezuela
Dominican Republic	Jordan	Qatar	