Assessing Idiosyncratic and Aggregate IMF Lending: An Exponential Model of Sample Selection*

Nicolas Mäder†, Jean-Guillaume Poulain‡, Julien Reynaud§

February 2, 2018

PRELIMINARY DRAFT - DO NOT CITE

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, higher 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 is expected to exhibit a substantial right tail. Our results provide an insightful input to policy discussions on the adequacy of size and composition of the IMF general resources.

Keywords: IMF lending, Sample selection

*The views expressed herein are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

†Department of Economics, Vanderbilt University; *nicolas.a.maeder@vanderbilt.edu*

‡International Monetary Fund; *jpoulain@imf.org*

§International Monetary Fund; *jreynaud@imf.org*

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 the official bilateral and multilateral arrangements. Yet, because coverage from Regional Financing Arrangements and Bilateral Swap Lines remain uneven, the IMF continues to play a major 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 [3].

Therefore, a fundamental question when discussing global financial architecture is the extent to which IMF resources can be deemed adequate for the instruction to fulfill its role as a lender of last resort. The aim of this paper is to contribute to this ongoing policy debate by expanding previous work by Poulain and Reynaud [14] on the determinants of IMF lending to model not only which member countries may require Fund financial assistance, but also the 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 IMF arrangements are only observed for a restricted and non-random sample of observations and/or member countries. Specifically, countries that approach the IMF often do so because they are already facing economic difficulties or expect to experience difficulties in the future. Similarly, structural vulnerabilities such as commodity dependence or poor governance may also lead to longer-term use of Fund resources, and thus result in increased exposure to shocks and decreased ability to implement appropriate macroeconomic policies in the face of these shocks, which would increase use of Fund IMF 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 IMF lending.

Accordingly, our main methodological contribution is to exploit the available panel data by allowing for unobserved country effects along both the selection and size dimensions, thereby directly allowing for an additional, permanent channel of sample selection. As explained in detail in Section 3, we do so by extending Heckman's selection correction model in the spirit of Mundlak [12] who had pioneeringly modeled unobserved heterogeneity by linearly projecting it onto the cross-sectional unit's observables, a practice now referred to by Wooldridge as *correlated random effects (CRE)*[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 useful insights not only on the adequate size of Fund resources, but also on the optimal composition between permanent and temporary (borrowed) resources, given the risk tolerance of the international community.

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 and justifies the empirical methodology used to model idiosyncratic IMF arrangements. Section 4 provides results of the estimation and discusses the fit of the model. In section 5, we use our model to simulate potential aggregate IMF lending under a range of global shock scenarios. Section 6 concludes.

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 has been 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 arrange-

[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 *strength of the member's program, member's balance of payments need and capacity to repay the Fund.*

ments 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.

Figure 1a - Distribution of IMF arrangement sizes - Histogram
(in percent of a member's GDP)

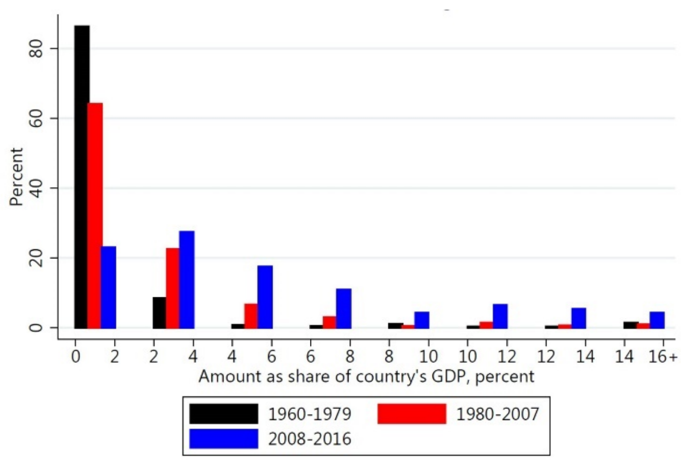

Figure 1b - Distribution of IMF arrangement sizes
(in percent of a member's GDP)

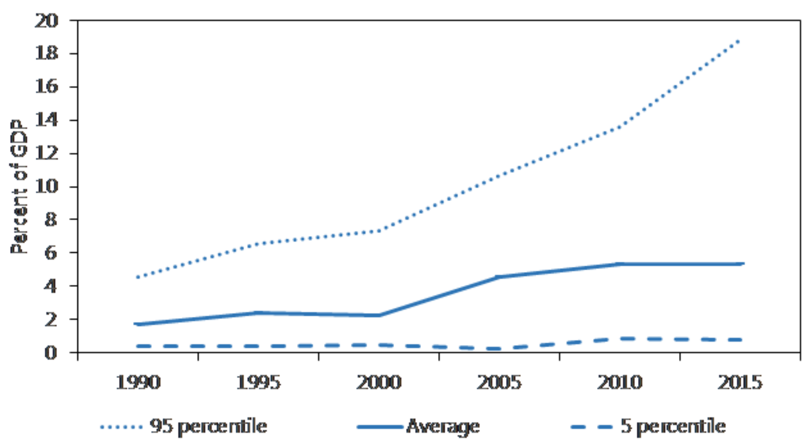

These trends likely reflect increased in countries' interconnectedness together with a shift towards more capital account-based crises. Figure 2 depicts the evolution of average arrangement size and external financing needs since 1990.

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

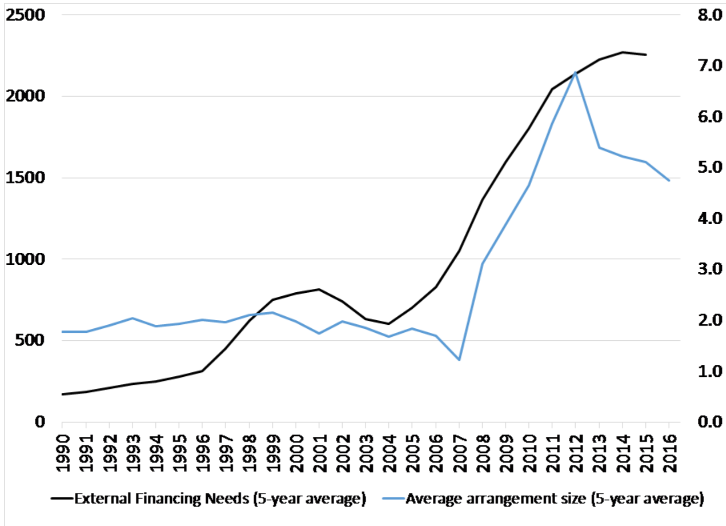

On an aggregate level, correlation between IMF credit outstanding and global risk aversion (proxied by the VIX) has been high. Poulain and Reynaud [14] discuss this relationship in greater detail, together with a discussion of factors influencing IMF lending.

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

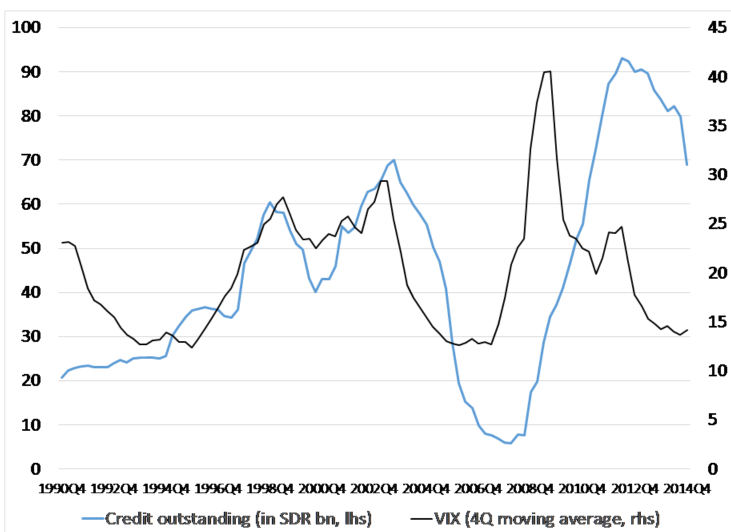

3 Empirical methodology

Sample selection

The literature has used various approaches to address sample selection bias, with the aim of constructing credible counterfactuals using instrumental variable strategies and more recently difference in differences techniques. Another approach is Heckman's selection correction model, which assumes correlated unobserved factors along the selection and size dimensions. In the context of IMF lending, Heckman's approach suggests the use of a probit model to predict the probability of a new Fund arrangement engagement in a first stage, followed by the inclusion of the corresponding inverse Mills ratio (IMR) as a regressor in the second stage to account for sample selection (see Przeworski and Vreeland [15]; Barro and Lee [1]; and Stubbs and Kentikelenis [18], for example). Because the IMR is approximately linear over a majority of its domain, any causal studies using the Heckman model [6] faces, similar to IV, the difficulty of having to provide a causal exclusion restriction' - a variable that influences selection into an IMF arrangement but not the outcome variable of interest. Without such a variable, multicollinearity renders identification of causal parameters imprecise at best or impossible at worst (Lang [9]; Wooldridge [19]). However, particularly because sample selection models are most popular within the inherently causal field of applied microeconometrics, it is important to stress that our goal is actually to simulate potential use of IMF resources, rather than inferring causality[2]. In this context, imposing exclusion restrictions only requires a zero conditional correlation, but not actual exogeneity. Before delving into the empirical estimation, we thus follow the methodology advocated by Pearl [13] and first precisely define our target quantity:

1. **Define**: Our target quantity is given by the conditional density of the dependent variable y given some (potentially endogenous) realization of the covariates Z. Crucially, since we are neither able, nor interested in exogenously influencing the covariates, notice that our parameters of interest explicitely do not represent causal effects, but, since we project linearly, partial correlations[3].

[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?

[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.

2. **Assume**: Since we are not assessing causal effects, no further *causal* assumptions, in the form of a path diagram for example, are required.

3. **Identify**: While our parameters of interest are technically identified through the model's nonlinearity, credibility of identification is well understood to be increasing in the number of exclusion restrictions[4]. In our context, imposing 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 [2]).

4. **Estimate**

Rather than relying on Heckman's original two-step, contemporary estimation of heckit models is traditionally implemented via full information maximum likelihood (FIML)[5]. Particularly in face of unobserved heterogeneity, the likelihood approach is preferred because it allows for a straightforward way of integrating out the corresponding effects [7]. 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[6]. Specifically, we proceed in the spirit of Chamberlain [4] who followed Mundlak [12] in linearly projecting unobserved heterogeneity by linearly projecting it 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.

Model

Consider a model in which the primary variable of interest - here 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. 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,

[4] Notice that exclusion restrictions in our context are only peripherally related to the quality of instruments when eliciting causal effects.

[5] See Wooldridge [19] for a FIML treatment of Heckman's original model.

[6] 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.

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

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

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

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 log or contemporaneously[7]. 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[8]. It is well understood that, given the setting described so far, 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 [6]. Letting S denote a scenario[9], we ultimately aim to recover the conditional distribution,

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

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 across i.

Estimation

We use Poulain and Reynaud's panel dataset of 92 advanced, emerging, and frontier market economies over the period 1992-2014, covering 119 arrangements financed by the IMF General Resource Account (GRA). See Poulain and Reynaud [14] as well as Annex C for description of the

[7] 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.

[8] 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. [20] for a discussion of truncated logistic distributions.

[9] If our model's stochastic environment is captured by the probability space $(\Omega, \mathcal{W}, \mu)$, then we may condition on Z_S by assuming the occurrence of a scenario $S \subset \Omega$ whose image under Z_{2t} is given by the singleton Z^S.

data, sources, and countries in the sample.

In spirit of the methodology presented by Heckman [7], the likelihood is constructed by first conditioning on and subsequently integating out the unobserved effects[10],

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

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[11],

$$\begin{aligned}
l_{it} &\equiv l(y_{it}|x_{it}, z_{it}, d_i, \varepsilon_i^c; \theta) \\
&= 1(s_{it} = 0) \log \left[1 - \Phi(z_{it}\gamma + d_i)\right] + \\
&\quad 1(s_{it} = 1)\{\log \left[\Phi \left(\frac{z_{it}\gamma + d_i + (\rho_{uv}/\sigma_u)(\log(y_{it}) - x_{it}\beta - \rho_{cd}(\sigma_c/\sigma_d)d_i - \varepsilon_i^c)}{(1-\rho_{uv}^2)^{0.5}} \right) \right] + \\
&\quad \log \left[\phi \left(\frac{\log(y_{it}) - x_{it}\beta - \rho_{cd}(\sigma_c/\sigma_d)d_i - \varepsilon_i^c}{\sigma_u} \right) \right] - \log(\sigma_u)\}
\end{aligned}$$

and ϕ and Φ denote the probability and cumulative density functions of the standard normal density respectively. Before integrating out the unobserved effects[12], 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[13] over the parameter space Θ,

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

[10]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.

[11]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).

[12]Independence and joint normality of (d_i, ε_i^c) allow for an implementation of Gauss-Hermite quadrature.

[13]We use `fminunc` in Matlab.

Exclusion restrictions

Quoting Little and Rubin [11], Puhani notes that "for the [Heckman] method to work in practice, variables are needed in [z_{it}] that are good predictors of [s_{it}] and do not appear in [x_{it}], that is, are not associated with [y_{it}] when other covariates are controlled" [16]. Similarly, Wooldridge highlights that while our parameter of interest is identified even when the covariate vectors are equivalent, this is only the case because the inverse Mills ratio is nonlinear in z_{ijt}. However, since the latter is approximately linear over a majority of its domain, employing the same set of regressors can introduce "severe collinearity" and thus large standard errors [19]. In spite of not explicitly relying on the inverse Mills ratio, notice that the FIML estimator is similarly as sensitive to multicollinearity as Heckman's original Two-Step [16]. Therefore, since in face of a high degree of multicollinearity β is identified through a distributional assumption at best and unidentified, when the vector of inverse Mills ratios is numerically indistinguishable from a linear combination of the covariates, at worst, it is highly recommended to impose at least one exclusion restriction. In short, the higher the number of and variation in excluded selection variables, the more credible is identification along the size dimension. However, notice that in our context, imposing exclusion restrictions is complicated by two factors. First, the introduced framework does not allow for separate identification of supply and demand, an issue already pointed out by Ghosh et al. [5]. Second, our parameters of interest do not represent causal effects, but partial correlations. As for the latter, 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 [2], $y_{it} \perp\!\!\!\perp Z_{ijt}|Z_{ikt}$ may fail as it is not implied by $y_{it} \perp\!\!\!\perp Z_{ijt}$. For these two reasons, we follow the forecasting literature by assessing relative performance of varying specifications by way of the Akaike and Bayesian information criteria[14], thereby only relying on economic intuition as a complementary measure.

Scenarios

Given the setup presented herein, as rightfully pointed out by Ghosh et al. [5], future individual and aggregate access are not predetermined because they are functions of the random variables $(Z_{2t+1}, v_{t+1}, u_{t+1})$ and the unobserved quantities (c, d). While the errors (v_{t+1}, u_{t+1}) exhibit a

[14]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.

persistence of zero by construction, we conversely should not expect $Z_{2t+1} \perp Z_t$ to be true[15]. More rigorously, unbeknownst the true stochastic environment of our economy as given by the probability space $(\Omega, \mathcal{W}, \mu)$, the model may be simulated by either employing estimate densities $\hat{\Pr}(c, d, u_{t+1}, v_{t+1})$, $\hat{\Pr}(Z_{2t+1}|Z_t)$ or alternatively by postulating the realization of a specific subset $S \in \Omega$.[16]

4 Results

In addition to borrowing their dataset, we build upon the model of Poulain and Reynaud [14] which serves as a natural starting point as we iteratively determine our model's specification[17]. Favored by the information criteria, the following specification summarized in Table 1 serves as our benchmark model. Notice that EFN per GDP serves as our only Mundlak variable, making an appearance along both dimensions. The estimate's precision is reported using the asymptotic z-statistic, derived from the negative Hessian of $\mathcal{L}$ at $\hat{\theta}_{MLE}$.

Table 1 - Benchmark model

	θ	Meaning	$\hat{\theta}_{MLE}$	$\lvert\hat{z}_H\rvert$	$\hat{p}_H$
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	-	-

[15] However, given a set of regularity conditions, $Z_{it+h} \perp Z_{it}$ may hold asymptotically. Intuitively, as h increases, conditioning on current values loses explanatory power such that the conditional density converges to its unconditional counterpart, if the latter exists. Please refer to *mixing* for a rigorous concept of asymptotic independence.

[16] Please refer to Appendix C for a brief definition of a scenario.

[17] 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 by Poulain and Reynaud. Please refer to their paper for a discussion of the data and an explanation of the panel's variables.

	Countries/observations	129	3741
	Used	92	1533
	Uncensored/censoreed	119	1414
IC	Akaike/Bayesian	918	1,025
R^2	Adjusted	0.30	

In addition to all selection variables matching the signs found by Poulain and Reynaud [14], the coefficient 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 its IMF arrangement would be larger.

Since our Mundlak variable, as defined in footnote 7, aims to capture *permanent* cross-country differences, intuitively likely most closely related to Reinhart, Rogoff, and Savastano's notion of "debt (in)tolerance" [17], 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, note that 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 $\hat{\rho}_{uv}$.

Exclusion restrictions

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. [5]. 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 [14]). 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 in 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) many constraint supply in some cases.
- *Credit-to-GDP gap*: Likely mirrored by higher demand, an increased credit-to-GDP gap may also lead to a decrease in supply.

In order to mitigate concerns arising from the joint identification of supply and demand, recognizing that our benchmark model is overidentified, we follow Zimran [21] by re-adding some of the initially excluded variables, thereby directly testing whether our exclusion restrictions were warranted in the first place. Proceeding as such, we find limited evidence for including interconnectedness and government stability, and no evidence for including past arrangements and the credit-to-GDP gap in the size equation. However, while individually or jointly including interconnectedness and government stability has little impact on the other coefficients, simulated aggregate access is affected in an economically significant manner (see appendix A).

Model fit - Selection

Given the estimate $\hat{\theta}_{MLE}$, we have,

$$\begin{aligned}\Pr(s_{it}=1|Z_{it};\hat{\theta}_{MLE}) &= \mathbb{E}[\Phi(z_{it}\hat{\gamma}+d_i)] \\ &= \int \Phi(z_{it}\hat{\gamma}+d_i)dF(d_i)\end{aligned}$$

where the integral may be evaluated using a Gauss-Hermite quadrature. Accordingly, the accuracy of 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

We calculate,

$$\begin{aligned}\mathbb{E}[\log(y_{it})|s_{it}=1,Z_{it};\hat{\theta}_{MLE}] &= x_{it}\hat{\beta}+\mathbb{E}[\mathbb{E}[c_i|v_{it}>-z_{it}\hat{\gamma}-d_i]]+\mathbb{E}[\mathbb{E}[u_{it}|v_{it}>-z_{it}\hat{\gamma}-d_i]] \\ &= x_{it}\hat{\beta}+\hat{\rho}_{uv}\hat{\sigma}_u\int \lambda(-z_{it}\hat{\gamma}-d_i)dF(d_i)\end{aligned}$$

Recognizing that $\mathbb{E}[f(x)] = f(\mathbb{E}[x])$ is generally false for any nonlinear function f, the corresponding level moment is recovered using Monte Carlo integration[18],

$$\mathbb{E}[y_{it}|s_{it}=1,Z_{it};\hat{\theta}_{MLE}] \approx \frac{1}{J}\sum_{j=1}^{J}\exp(x_{it}\hat{\beta}+\hat{\rho}_{cd}(\hat{\sigma}_c/\hat{\sigma}_d)\tilde{d}_{ij}+\tilde{\varepsilon}^c_{ij}+\hat{\rho}_{uv}\hat{\sigma}_u\tilde{v}_{ijt}+\tilde{\varepsilon}^u_{ijt})$$

[18]Since the sum $c_i + u_{it}|d_i, s_{it}$ is neither normal nor truncated normal, see Lipow et al. [10] for an analysis of the resulting density, an analytical derivation of $\mathbb{E}[y_{it}|s_{it}=1,Z_{it};\hat{\theta}_{MLE}]$ is nontrivial at best.

where $(\tilde{\varepsilon}_{ij}^c, \tilde{\varepsilon}_{ijt}^u, \tilde{d}_{ij}, \tilde{v}_{ijt})$ are drawn from independent normal densities, the last of which is truncated below by $z_{it}\hat{\gamma} - \tilde{d}_{ij}$.

Figure 4 - Residual Plot

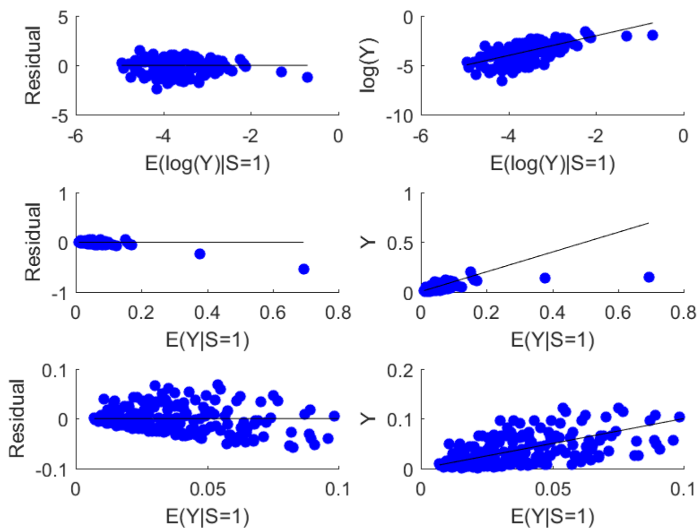

When examining figure 1, note that the plots depict both the log moment $\mathbb{E}[\log(y_{it})|s_{it} = 1]$, the one fitted by our maximum likelihood algorithm, and the level moment $\mathbb{E}[y_{it}|s_{it} = 1]$. Before analyzing the second and third rows, the latter of which is simply an enlarged 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.

As emphasized before, the sum of a normal and a truncated normal random variable is neither normal nor truncated normal. 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. Further examining the two graphs, the reader may wonder what could potentially account for the two seemingly very large negative shocks experienced by the observations at the left tail. The loans corresponding to those residuals are the Greek arrangements of 2010

Figure 5 - Residual Distribution

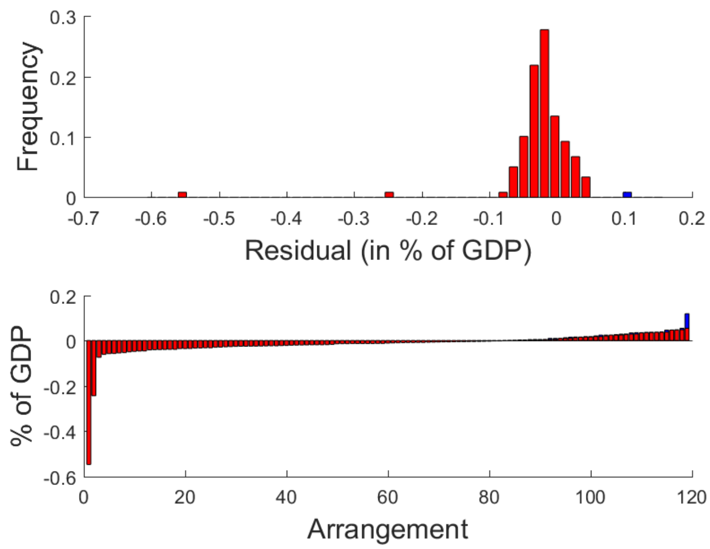

and 2012. Crucially, note that both of these arrangements were to a large extent financed by European facilities. When accounting for the sizable European contributions, the two residuals exhibit a substantial shift to the right (as captured by the blue observation) meaning that the Greek arrangements were in fact rather large conditionally speaking.

5 Scenarios

In this section, we use the model described earlier in the paper to simulate potential aggregate IMF lending in 2015 under a wide range of global shock scenarios following two different approaches[19].

- 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 constrained to be equal to some realization Z^1_{2t}.
- In the second approach, a country is selected if its estimated selection probability exceeds a certain threshold $s^1_{it} = 1(\int \Phi(z_{it}\hat{\gamma} + d_i)dF(d_i) \geq t^1)$, given some realization Z^2_{i2t}, where the size of each arrangement is pinned down by $u^2_{it} = \hat{\rho}_{uv}\hat{\sigma}_u v^2_{it} = -\hat{\rho}_{uv}\hat{\sigma}_u z_{it}\gamma$.

[19]Of course, it would be possible, in a separate third scenario, to sample from the unrestricted sample space Ω by integrating out the random variable Z_{2t}. While such an approach would allow for a fairly general statement regarding the marginal distribution over aggregate lending, probabilistic accuracy would inherently hinge on the additional estimate density $\hat{\Pr}(Z_{2t}|Z_{t-1})$. Since the objective of our paper explicitly does not lie in the derivation of a (unconditional) marginal density, we purposely omit a scenario of this sort.

The first approach is conducted in style of a typical Monte Carlo simulation, whereas the second alternative simply assumes the realization of what is perceived as an interesting benchmark case, namely that all countries with a selection likelihood exceeding some threshold are in fact selected. While the former is methodologically closer to the spirit of the original empirical model, the latter's results may be perceived as more concrete and thus more intuitive.

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

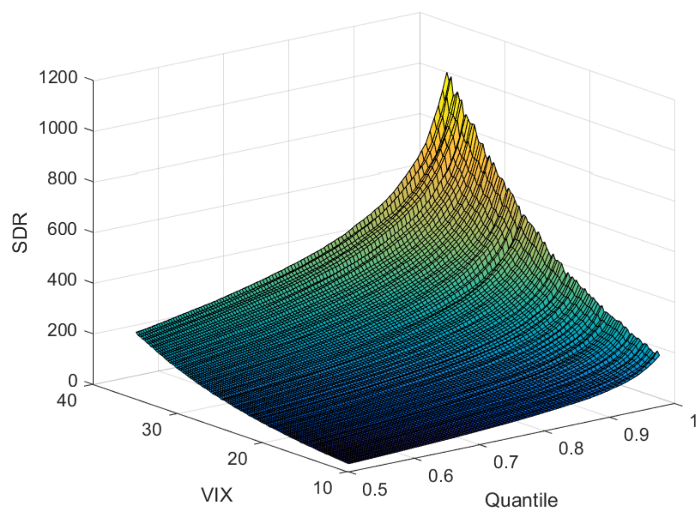

Figure 7 - Potential aggregate IMF lending under the second approach
(fix VIX, fix errors)

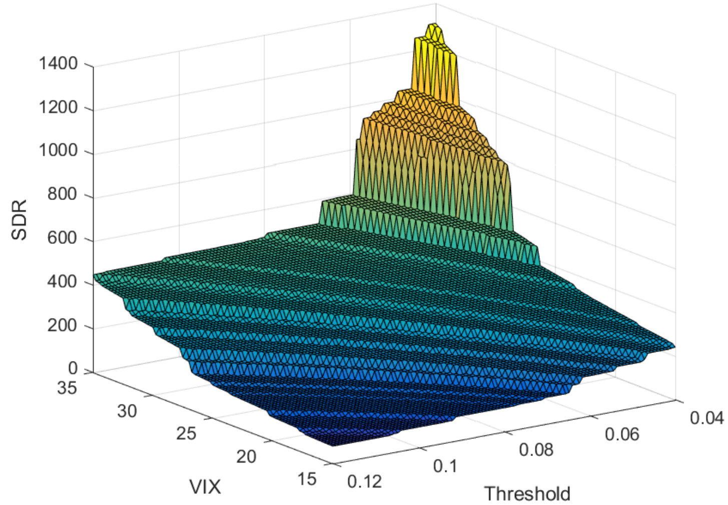

While both scenarios generate a similar range of conceivable aggregate access, the second approach is characterized by a fixed ordinal ranking of country selection whereas a scenario of the first type yields some additional statistical insights. In particular, as suggested by the first graph, the (conditional) density over aggregate access features a substantial right tail, a speculation indubitably confirmed by median third and fourth moments of 6 and 87. Unsurprisingly, the distribution of aggregate access is thus characterized by a sizeable wedge between the conditional expectation, a moment sometimes emphasized in *Value-at-Risk* analyses, and our primary moments of interest, the upper quantiles.

6 Conclusion

In this paper, we provide a methodology to correct for an important methodological challenge, namely the issue of country level sample selection bias when estimating of the potential need for IMF arrangements. Exploiting the insights provided by Mundlak [12], Chamberlain [4] and Heckman [7], we use the model of Poulain and Reynaud [14] as a starting point for our estimation. In addition to all coefficients in the selection equation matching the signs found by Poulain and Reynaud [14], 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 potential future use of IMF resources under a wide range of global shock scenarios. From a qualitative perspective, our simulations yield non-normal aggregate access densities. While the proposition of aggregate non-normality generally aligns well with economic intuition, recall that we know the sum of independent and identically distributed random variables, whose means and variances are finite, to be asymptotically normal by the Central Limit Theorem. In fact, even if our sequence of interest is not identically distributed, aggregate normality may still hold as long as independence is not violated[20]. Accordingly, if some econometrician's ex ante aim were to generate aggregate non-normality, a straightforward and intuitive approach would be to break with independence[21]. While introducing cross-sectional

[20] See *Lyapunov condition*.

[21] A common approach to model stochastic dependence in mathematical finance is the use of multivariate probability distributions called copulas.

stochastic dependence could technically be achieved through either the covariates or the errors, the distributional assumptions imposed by our model render the latter channel impractical. Of course, in our specification, the potential for aggregate non-normality arises from the existence of the common random covariate, the VIX. Presumably further proliferated by heterogeneity in cross-country means, non-normality should therefore rather be viewed as following from our assumptions rather than constituting a result in its own right. In essence, since we assume individual arrangements to be neither independent, nor identically distributed, we should indeed expect their sum to be non-Gaussian.

From a quantitative perspective, both sets of scenarios suggest that while most crises could be accommodated with a lending capacity of SDR 600 billion, a more extreme crisis could conceivably lead to much larger needs, north of SDR 1 trillion. However, recall that our analysis has focused on the conditional upper quantiles of a random variable, whose heavy tail drives a sizeable wedge between its expected value and the upper quantiles. For example, conditional on current low levels of global risk aversion, note that expected aggregate access is below SDR 150 billion, consistent with the idea that the global financial cycle is a major driver of IMF lending cycles. Moreover, it should be emphasized again that the primary objective of the paper has been limited to the estimation of conceivable quantities, but not the derivation of a marginal density. While assigning estimated probabilities to outcomes is certainly possible, credibility would likely suffer from potentially error prone distributional assumptions.

Finally, from a policy perspective, bearing in mind the methodological caveats outlined above, these results may provide useful insights not only regarding the adequate size of Fund resources, but also their optimal composition between quotas and borrowed resources. What kind of crisis scenarios should Fund resources be able to accommodate? Among these, which should be covered by permanent quota resources and borrowed resources? This is eventually a matter of informed judgment, based on risk tolerance as well as the political and financial cost of providing different types of resources.

References

[1] Robert J. Barro and Jong-Wha Lee. IMF programs: Who is chosen and what are the effects? *Journal of Monetary Economics*, 2005.

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

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

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

[5] 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.

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

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

[8] Frank H. Knight. *Risk, Uncertainty, and Profit*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Co., 1921.

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

[10] Myron Lipow, Nathan Mantel, and J.W. Wilkinson. The sum of values from a normal and a truncated normal distribution. *Technometrics*, 1964.

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

[12] Yair Mundlak. Pooling of time series and cross section data. *Econometrica*, 1978.

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

[14] Jean-Guillaume Poulain and Julien Reynaud. IMF lending in an interconnected world. *IMF Working Paper*, 2017.

[15] Adam Przeworski and James Raymond Vreeland. The effect of IMF programs on economic growth. *Journal of Development Economics*, 2000.

[16] Patrick Puhani. The heckman correction for sample selection and its critique? A short survey. *Journal of Economic Surveys*, 2000.

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

[18] Thomas Stubbs and Alexander Kentikelenis. Targeted social safeguards in the age of universal social protection: The IMF and health systems of low-income countries. *Critical Public Health*, 2018.

[19] Jeffrey M. Wooldridge. *Econometric analysis of cross section and panel data.* MIT Press, 2010.

[20] 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.

[21] Ariell Zimran. Does sample-selection bias explain the antebellum puzzle? Evidence from military enlistment in the nineteenth-century united states.

A Alternate model

Recall that our benchmark model features four exclusion restrictions. When imposing these restrictions, we heavily rely on standard model selection criteria because a corresponding economic evaluation is complicated by the fact that our approach does not allow for separate identification of supply and demand. Since solely relying on statistical quantities may be viewed as unsatisfactory, we also consider the following alternate specifications.

Table 3 - Alternate model

	θ	Meaning	$\hat{\theta}_{MLE}$	$\lvert\hat{z}_H\rvert$	$\hat{p}_H$
Size	β_1	EFN per GDP	3.36	4.10	0.000***
		GDP per capita	0.25	2.29	0.022**
		Exchange rate variation	-1.12	2.76	0.006***
		Growth	-0.05	2.96	0.003***
		Interconnectedness	0.41	1.51	0.13
		Government Stability	-0.10	1.81	0.07*
	β_2	**VIX**	0.04	3.59	0.000***
	β_3	EFN per GDP	(-)		
Selection	γ_1	EFN per GDP	2.51	3.19	0.001***
		GDP per capita	-0.42	4.00	0.000***
		Government Stability	-0.17	3.84	0.000***
		Interconnectedness	-0.30	1.79	0.073*
		Past Arrangement	0.13	2.78	0.005***
		Credit-to-GDP gap	0.01	2.45	0.014**
		Exchange rate variation	-0.91	2.17	0.03**
		Growth	-0.04	2.87	0.004***
	γ_2	**VIX**	0.04	3.33	0.001***
	γ_3	EFN per GDP	(-)		
Errors	ρ_{uv}	Sample selection parameter	0.53	2.67	0.008***
	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	1,033	
	R^2	Adjusted	0.31		

As aluded to above, while the magnitude and significance levels of the other coefficients are comparable to the benchmark model, simulated aggregate access is affected in an economically significant manner. Having increased at all point of its domain, the conceivable range of aggregate access in fact exceeds SDR 2 trillion (see figures A1 and A2).

Figure A1 - Potential aggregate IMF lending under approach 1
(fix VIX, simulate errors)

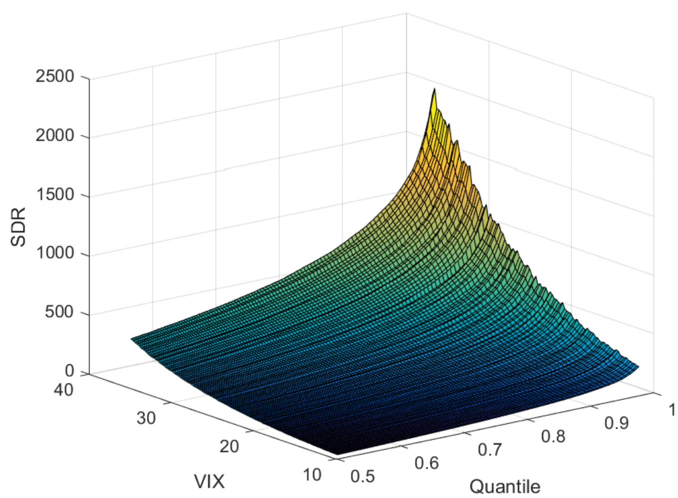

Figure A2 - Potential aggregate IMF lending under approach 2
(fix VIX, fix errors)

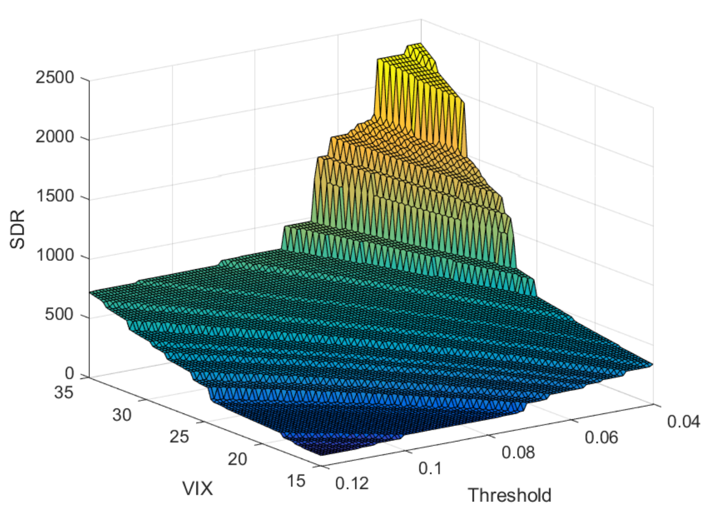

B A Scenario

B.1 Definition

Suppose our economy's stochastic environment is given by the probability space $(\Omega, \mathcal{W}, \mu)$. If misspecification concerns realistically render an estimation of the measure μ impractical, the econometrician is faced with a Knightian type of uncertainty [8]. Therefore, rather than relying entirely on a likely error prone best estimate $\hat{\mu}$, we may thus prefer to condition on the occurrence of a particular subset $S \subset \Omega$ by constructing an alternate space $(\Omega, \mathcal{W}, \hat{\mu}_S)$ with $\hat{\mu}_S(\Omega \setminus S) = 0$.

Definition. Consider a probability measure $\mu_S : \mathcal{W} \mapsto [0,1]$ satisfying $\mu_S(A) = \mu(A)/\mu(S)\ \forall A \in \mathcal{S}$, where $\mathcal{S} \subset \mathcal{W}$ denotes the σ-algebra of the smallest set in $\mathcal{W}$, called S, satisfying $\mu_S(A) = 1$. Then, let a scenario, constructed by way of the alternate space $(\Omega, \mathcal{W}, \mu_S)$, be defined as the set S.

Defining the alternate space $(\Omega, \mathcal{W}, \hat{\mu}_S)$ with estimate measure $\hat{\mu}_S$ allows for the recovery of an estimate density $\hat{f}_{X|S}(x) \equiv \hat{\Pr}(X = x|\omega \in S)$ over any $\mathcal{W}$-measurable random variable $X(\omega)$,

$$\begin{aligned}\hat{\Pr}(X(\omega) = x|\omega \in S) &= \hat{\mu}_S(X^{-1}(x) \cap S) \\ &= \hat{\mu}_S(X^{-1}(x))\end{aligned}$$

where the last equality is intended to highlight that $\hat{\mu}_S$, just like $\hat{\mu}$, allows for the measurement of any $\mathcal{W}$-measurable random variable[22] such as X. However, lacking the measure μ, note that a scenario is generally silent regarding the true likelihood of its own occurrence $\mu(S)$, the corresponding best estimate being given by $\hat{\mu}(S)$, and thus on the random variable's unconditional density $f_X(x) \equiv \Pr(X(\omega) = x)$,

$$\begin{aligned}\Pr(X(\omega) = x) &= \Pr(X(\omega) = x|\omega \in \Omega) \\ &= \mu(X^{-1}(x)) \\ &= \int_{\mathcal{W}} \mu_S(X^{-1}(x))\mu(S)dS\end{aligned}$$

Note that the lack of μ not only prevents us from verifying whether $\hat{f}_{X|S}(x) = f_{X|S}(x)$ is true $\forall x \in \mathcal{B}(\mathbb{R})$, but also from evaluating the above integral, that is unless we decide to rely on our best estimate $\hat{\mu}_S = \hat{\mu}$. Of course, accuracy of the resulting unconditional density estimate $\hat{f}_X$ is inherently sensitive to the potentially error prone measure $\hat{\mu}$.

[22] Conversely, the restriction of μ to $\mathcal{S}$, $\mu_{|}\mathcal{S}$, only allows for the evaluation of $\mathcal{S}$-measurable random variables.

B.2 Application

Conditional simulation

Conditional on the occurrence of a scenario S, as defined above, we are interested in the following density given the estimated version of our model,

$$
\begin{aligned}
f_{Y|S} &\equiv \Pr\left(Y_t = Y | Z_{i1t-1}, Z_{i2t} = Z^S\right) \\
&= \Pr\left(Y_t = Y | c, d, u_t, v_t, Z_{2t}, Z_{t-1}\right) \Pr(c, d, u_t, v_t, Z_{i2t} | Z_{it-1}; S)
\end{aligned}
$$

where it is postulated that $S \subset \Omega$ with $Z_{i2t}(\omega) = Z^S$ for all $\omega \in S$ will occur with probability one. Of course, $f_{Y|S}$ may be degenerate as illustrated by a deterministic case in which $|S| = 1$. Conversely, as long as the density is not degenerate, computing $f_{Y|S}$ requires a measure $\hat{\mu}_S$ as captured by the conditional densities $\hat{\Pr}(c_i, d_i, u_{it}, v_{it}|S)$ and $\hat{\Pr}(Z_{i2t}|Z_{it-1}; S)$. In order to recover aggregate moments of interest, e.g. quantiles, we sample a vector,

$$
\begin{aligned}
\{\tilde{y}_{jt}\}_{j=1}^{J} &= \left\{ \sum_{i=1}^{n} \tilde{s}_{ijt} \tilde{y}_{ijt} \right\}_{j=1}^{J} \\
&= \left\{ \sum_{i=1}^{n} 1(\tilde{v}_{ijt} > -\tilde{z}_{ijt}\hat{\gamma} - \tilde{d}_{ij}) \exp(\tilde{x}_{ijt}\hat{\beta} + \tilde{c}_{ij} + \tilde{u}_{ijt}) \right\}_{j=1}^{J}
\end{aligned}
$$

from $\hat{f}_{Y|S}$ by drawing $\{\tilde{c}_{ij}, \tilde{d}_{ij}, \tilde{v}_{ijt}, \tilde{u}_{ijt}, \tilde{x}_{ijt}, \tilde{z}_{ijt},\}_{j=1}^{J}$ from the estimated conditional densities $\hat{\Pr}(c_i, d_i, u_{it}, v_{it}|S)$ and $\hat{\Pr}(Z_{i2t}|Z_{it-1}; S)$.

Unconditional simulation

Consider the particular case in which we sample from the entire sample space $S = \Omega$,

$$
\begin{aligned}
f_Y &\equiv f_{Y|\Omega} \\
&= \Pr\left(Y_t = Y | Z_{it-1}\right) \\
&= \Pr\left(Y_t = Y | c_i, d_i, u_{it}, v_{it}, Z_{i2t}, Z_{it-1}\right) \Pr(c_i, d_i, u_{it}, v_{it}) \Pr(Z_{i2t} | Z_{it-1})
\end{aligned}
$$

$\hat{\Pr}(c_i, d_i, v_{it}, u_{it})$ being given by the model, the unconditional approach proceeds as follows,

1. Estimate a density $\hat{\Pr}(Z_{i2t}|Z_{it-1})$, thus yielding a measure $\hat{\mu}_S = \hat{\mu}$ on the entire space Ω
2. Draw a matrix of J independent samples $\{\tilde{c}_{ij}, \tilde{d}_{ij}, \tilde{v}_{ijt}, \tilde{u}_{ijt}, \tilde{x}_{ijt}, \tilde{z}_{ijt},\}_{j=1}^{J}$

3. Calculate the resulting selection indicator and outcome variable $\{\tilde{s}_{ijt}, \tilde{y}_{ijt},\}_{j=1}^{J}$

4. Aggregate $\{\tilde{Y}_{jt}\}_{j=1}^{J} = \{\sum_{i=1}^{n} \tilde{s}_{ijt}\tilde{y}_{ijt}\}_{j=1}^{J}$

where each element of the resulting vector $\{\tilde{Y}_{jt}\}_{j=1}^{J}$ signifies an independent sample from the unconditional estimate density $\hat{f}_Y$.

C Data

Table C1 - Data - Variables and Sources

Variable	Source	Explanation
New arrangement	IMF	Indicator equals one when new arrangement was put in place
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 [14]
GDP Growth	WEO	In percent
GDP per capita	WEO	log of level in USD
GDP	WEO	log of level in USD billion
Credit-to-GDP gap	BIS; WDI; WEO	Deviation of credit-to-GDP from its 5-year moving average
Exchange rate variaton	WEO	Variation of bilateral nominal exchange rate against the USD
Government Stability	ICRG	Government unity, legislative strength, popular support
Interconnectedness	WEO; DOTS	See Poulain and Reynaud (2017)
3-month interest rate	WEO	In percent
VIX	CBOE	Measure of risk aversion
Oil price	WEO	Deviation from 5-year moving average

Table C2 - Countries in the sample

Albania	Ecudador	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	