

On the Potential of Foreign Aid as Insurance

Stéphane Pallage

Michel A. Robe

Catherine Bérubé*

June 2004

Abstract

In this paper, we quantify foreign aid's potential as an insurance mechanism against macroeconomic shocks. Within a dynamic model of aid flows between two endowment economies, we show that at least three fourths of the large welfare costs of macroeconomic fluctuations in poor countries could be alleviated by a simple reallocation of aid flows across time. In developing countries subject to persistent macroeconomic shocks, the welfare gain is of first-order magnitude.

Keywords: Macroeconomic fluctuations, Welfare, Foreign aid, Consumption smoothing

JEL Classification: F35, E32, E60

*We thank Kevin Carey for helpful discussions, and Şebnem Akkaya, Oya Celasun, Ravi Kanbur, Andrei Kirilenko, Robert L. Losey, İnci Ötker-Robe, Christian Zimmermann, and participants to the conference celebrating “75 Years of Development Research” at Cornell University for useful comments. This research was partially supported by FCAR/FQRSC and a Kogod endowed fellowship. The usual disclaimer applies. **Pallage:** Department of Economics, Université du Québec à Montréal, C.P. 8888 Succursale Centre-Ville, Montreal, QC, H3C 3P8, Canada. Tel: 514-987-3000 (ext.8370). Email: pallage.stephane@uqam.ca. **Robe (corresponding author):** Kogod School of Business at American University, 4400 Massachusetts Avenue NW, Washington, DC 20016, United States. Tel: 202-885-1880. Email: mrobe@american.edu. **Bérubé:** Finance Canada (Canadian Ministry of Finance), 140 O'Connor Street, Ottawa, Ontario K1A 0G5, Canada. Tel: 613-996-9963. Email: Berube.Catherine@fn.gc.ca.

On the Potential of Foreign Aid as Insurance

Abstract

In this paper, we quantify foreign aid's potential as an insurance mechanism against macroeconomic shocks. Within a dynamic model of aid flows between two endowment economies, we show that at least three fourths of the large welfare costs of macroeconomic fluctuations in poor countries could be alleviated by a simple reallocation of aid flows across time. In developing countries subject to persistent macroeconomic shocks, the welfare gain is of first-order magnitude.

Keywords: Macroeconomic fluctuations, Welfare, Foreign aid, Consumption smoothing

JEL Classification: F35, E32, E60

1 Introduction

Developing countries are subject to strong macroeconomic shocks. Recent evidence suggests that these shocks are mostly exogenous (Kose, 2002) and that they result in consumption volatility levels deleterious to their residents' welfare (Pallage & Robe, 2003). Unfortunately, developing countries' access to private international financial markets typically dries up precisely when their economies hit the doldrums. Therefore, because aid inflows are a major source of external capital for these nations, foreign aid has the potential to play a key role in smoothing out the impact of output fluctuations on their domestic consumption.¹ In practice, though, foreign aid does not play this role. Indeed, in most developing countries, the bulk of net aid inflows are positively correlated with the contemporaneous level of the local gross domestic product (GDP). That is, foreign aid is procyclical.² In this paper, we assess foreign aid's potential as an insurance device.

Precisely, we quantify the possible welfare gains from changing the timing of foreign aid disbursements while keeping their average level constant. To keep the analysis transparent, we build a simple dynamic model of aid flows between two endowment economies: a large altruistic donor country and a much smaller recipient country, with symmetric information sets. We approximate the fact that much of the aid to poor countries comes as outright grants by positing that aid is given out with no expectation of repayment. Finally, we assume that the donor country adjusts aid flows in each period to maximize a weighted average of its own expected utility and that of the recipient country.

We first solve for the optimal aid policy in this environment. Not surprisingly, it is countercyclical: that is, aid flows should be negatively correlated with the recipient's endowment stream. Such a policy, however, does not merely dampen the variability of the recipient's consumption – it massively reduces it. For a poor recipient country, the optimal aid policy ends up bringing the very large percentage volatility of per capita consumption down to the much lower one in the donor country. The effect of this policy on volatility in the donor country itself is negligible.

Next, we obtain quantitative estimates of the effects of this policy on the recipient's welfare. To do so, we calibrate the donor and the recipient's consumption processes to match key moments of the actual local-currency figures from the World Bank or extant studies. We set the donor

¹In the last fifteen years, official development assistance has made up between 10% and 60% of net capital flows to all developing nations; for the poorest countries, the average ratio has ranged from 50% to 90% (World Bank, 2002).

²See Pallage & Robe (2001) and Bulír & Hamann (2003). Gupta, Clements, & Tiongson (2003) document that food aid is countercyclical, but only in countries with the greatest need for such aid. In most countries that receive food aid, food aid flows are acyclical.

country's altruism parameter to replicate the magnitude of foreign aid flows observed in practice (8% of GDP for the median developing country). We consider typical ranges for the risk aversion levels and discount factors, and carry out robustness checks.

We find that altering the timing of aid disbursements would be worth at least 0.25%, and quite possibly more than 8%, of permanent consumption in the recipient country. Predictably, the magnitude of the recipient's welfare gain varies with the assumed level of risk aversion, the mean growth rate of consumption, and the magnitude and (especially) persistence of aggregate shocks. Strikingly, however, the fraction of the welfare costs of macroeconomic fluctuations that could be alleviated by merely changing the timing of aid flows always exceeds 75%. In short, we find that changing the intertemporal pattern of foreign aid inflows has the potential to create first-order welfare improvements through better risk sharing and increased consumption smoothing.

We choose to abstract from the possibility of informational asymmetries or agency conflicts between donor and recipient or between various constituencies in the donor or recipient countries.³ This approach allows us to keep the analysis simple despite the dynamic nature of foreign aid relationships. More importantly, it allows us to provide a quantitative sense of the full potential of using foreign aid as an insurance device, which has recently been advocated in both academic and policy circles (Collier & Dehn, 2001; Caballero, 2003; *The Economist*, 2003, for example).

We are fully aware that moral hazard may reduce the extent to which it is desirable to offer insurance against macroeconomic shocks. We are also fully aware that implementing a counter-cyclical aid policy may be fraught with practical difficulties. Issues related to informational imperfections, however, arise with any insurance policy. What matters is that the insurance benefits be large enough. Our analysis shows that, in the presence of persistent macroeconomic shocks, these benefits would be of first-order magnitude.⁴ In other words, our results indicate that it is worth revisiting foreign aid's potential as an insurance device.

Our analysis is related to Arellano, Bulir, Lane, & Lipschitz (2002). Those authors use a computable general equilibrium model to quantify the impact of foreign aid flows on the production

³Several recent papers study how heterogeneous information sets or conflicts between parties can influence foreign aid contracts. Those papers' main focus is on the use of foreign aid to promote specific policies in the recipient country (Murshed & Sen, 1995; Casella & Eichengreen, 1996; Orphanides, 1996; Lahiri & Raimondos-Moller, 1997a, 1997b; Asiedu & Villamil, 2002; Azam & Laffont, 2003; Svensson, 2003) or on the optimal allocation of a given amount of foreign aid between potential recipients (Lahiri & Raimondos-Moller, 2000; Svensson, 2000).

⁴In that sense, our results are related to the large literature on the welfare benefits of unemployment insurance programs which has shown that, despite the possibility of moral hazard, it is optimal to offer generous unemployment benefits in a variety of model economies (Pallage & Zimmermann, 2001; Wang & Williamson, 2002).

of tradable goods in developing countries. They also measure the welfare implications of observed aid patterns, by contrasting welfare levels when aid is volatile and procyclical versus when aid is kept constant at the mean. Our focus is different. We do not seek to identify the impact of foreign aid on a particular sector of the economy. Instead, our goal is to highlight the consumption-smoothing potential of foreign aid. Therefore, rather than take the aid policy as exogenous and then compute its impact on the recipient country's economy, we take the recipient country's aggregate consumption process as given and then quantify the welfare gain from fine-tuning aid disbursements to smooth out that consumption stream.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 derives the optimal aid policy. Section 4 puts our theoretical results in perspective. Section 5 discusses computational aspects and provides our quantitative results. Section 6 concludes.

2 Model

To study the insurance potential of foreign aid in a recipient country, we build a model of aid flows between two endowment economies – a rich donor and a poor recipient. We abstract from modelling any impact that aid might have on GDP or on GDP growth. This approach reflects the fact that, after decades of empirical research, there is little consensus on whether foreign aid affects growth in practice (Easterly, 2003).⁵

Ours, then, is a simple infinite horizon model of a donor country and a recipient country (respectively $i = D, R$) with time-varying endowments of a single, non-storable consumption good. Each country's harvest of the good at time t , $y_{i,t}$, is a random variable that follows a known stochastic process Φ_i . The donor country is strictly richer than the recipient country. We assume that Φ_R and Φ_D are independent processes, which reflects empirical evidence that there has been very little co-movement between consumption levels in rich and developing countries in any of the last four

⁵Hansen & Tarp's (2000) review of the early literature, for example, concludes that most empirical studies show aid to benefit growth. Yet, according to the widely cited studies by Boone (1994, 1996), foreign aid is not invested but is instead fully consumed. In the same vein, Easterly (1999) and Dollar & Easterly (1999) argue that the links between aid and investment, and between investment and growth, are both tenuous. On the question of whether good economic stewardship by recipient governments might be key to foreign aid's fostering growth (Burnside & Dollar, 2000), several recent studies based on very similar data end up with conflicting conclusions: See, in particular, Hansen & Tarp (2001) and Easterly, Levine, & Roodman (2003) and the discussion and references in Easterly (2003). In sum, nothing is as sure as the fact that the aid-growth relation is at best unclear. We therefore focus on another direct effect of aid flows on welfare, through consumption smoothing.

decades (Kose, Otrok, & Whiteman, 2003; Kose, Prasad, & Terrones, 2003).

The representative agents in both countries are infinitely lived. The recipient country's representative agent has preferences over his own total consumption c_R over time:

$$U_R = \sum_{t=0}^{\infty} \beta_R^t u_R(c_{R,t}) \quad (1)$$

where $\beta_R \in [0, 1]$ is that agent's discount factor. In contrast, the donor country's representative agent not only cares about his own consumption c_D but also displays a form of altruism, in that he puts some weight on the well-being of the recipient country's residents:

$$U_D = \sum_{t=0}^{\infty} \beta_D^t [A u_D(c_{D,t}) + (1 - A) u_R(c_{R,t})] \quad (2)$$

where $\beta_D \in [0, 1]$ is the donor agent's discount factor and $A \in (0, 1)$ is the weight that he puts on his own utility.⁶

The donor country can allocate part of its harvest $y_{D,t}$ at time t to foreign aid a_t . We approximate the fact that much of the aid to poor countries comes as outright grants by positing that aid is given out with no expectation of repayment.⁷ We assume, as is the case in practice (Prati, Sahay, & Tressel, 2003), that the recipient country cannot save any part of these aid flows.⁸ With no possibility of storage, the periodic consumption levels in each country are respectively $c_{D,t} = y_{D,t} - a_t$ and $c_{R,t} = y_{R,t} + a_t$.

The donor's generosity is financially constrained in two ways. First, each period $0 \leq a_t \leq y_{D,t}$. Second, we shall calibrate the altruism parameter A in Section 5 so that the lifetime average $a_t/y_{R,t}$ ratio is in line with its empirically documented counterpart. This second restriction leads to the following remarks. One, by calibrating A to past data, we are in effect positing that the

⁶Our analysis implicitly assumes that the two countries have equal populations or, alternatively, are abstractions for two groups of donors and recipient nations whose respective total populations have equal sizes. At first sight, this assumption might seem restrictive. In fact, it covers many individual donor-recipient pairs because the model is readily generalized to cases in which the recipient country's population is smaller than the donor country's. It also applies to an aggregate of the Organization for Economic Co-operation and Development's donor economies (i.e., the 22 nations that belong to the OECD's Development Assistance Committee), on the one hand, and an aggregate of all African recipients, on the other hand.

⁷For example, 1975 to 1995 data from Chang, Fernandez-Arias, & Servén (1999) indicate that the median grant element in official development assistance loans to non-oil-producing African countries exceeded 90%.

⁸This assumption is also in line with anecdotal evidence that recipients who do not spend all of the aid they receive in any given period can expect to see future aid flows reduced, because they have "just demonstrated" a lack of absorption capacity.

donor's preferences are time-invariant. Two, our analysis may not apply in the long run to recipient countries that are growing at a much faster clip than the donor is. Accordingly, we shall assume that $y_{R,t}$ grows at the same rate as does $y_{D,t}$. This simplification is without much loss of generality, in that our qualitative results are even stronger if we assume instead that the recipient's growth rate is lower than the donor's (which is empirically true for most of the world's poorest countries).

Throughout the paper, we assume that both representative agents have constant relative risk averse, time-separable preferences over their own consumption. This assumption makes analytical solutions to the donor's problem possible. Furthermore, because one would expect macroeconomic fluctuations to bring about massive welfare costs under the alternative assumption of time-non-separable preferences (Otrok, 2001), this assumption guarantees that the large welfare gain estimates we obtain in Section 5 are not merely due to our choices of utility functions. There is no evidence suggesting that residents of developing, emerging, and developed economies exhibit different intensities of relative risk aversion (Ostry & Reinhart, 1992), so we posit that the donor and recipient agents have the same constant level of risk aversion γ . That is:

$$u_i(c_i) \equiv \begin{cases} \frac{c_i^{1-\gamma}-1}{1-\gamma} & (\gamma \neq 1) \\ \ln(c_i) & (\gamma = 1) \end{cases} \quad (i = D, R)$$

3 Optimal aid policy

The donor's problem is recursive and can be written as:

$$V(y_{D,t}, y_{R,t}) = \max_{a_t} (Au_D(y_{D,t} - a_t) + (1 - A)u_R(y_{R,t} + a_t) + \beta_D E[V(y_{D,t+1}, y_{R,t+1})]) \quad (3)$$

$$\text{s.t. } 0 \leq a_t \leq y_{D,t}$$

Aid does not affect future realizations of $y_{D,t}$ and $y_{R,t}$, so the optimal aid rule is straightforward:

Proposition 1 *It is optimal for the donor to transfer to the recipient, each period, an amount:*

$$a_t^* \equiv a_t^*(y_{D,t}, y_{R,t}) = \max\left[0, \frac{(1 - A)^{\frac{1}{\gamma}} y_{D,t} - A^{\frac{1}{\gamma}} y_{R,t}}{(1 - A)^{\frac{1}{\gamma}} + A^{\frac{1}{\gamma}}}\right] \quad (4)$$

Proof. The first order condition for an interior optimum to the donor's problem (3) is:

$$\frac{u'_R(y_{R,t} + a_t)}{u'_D(y_{D,t} - a_t)} = \frac{A}{1 - A} \quad (5)$$

With logarithmic utility ($\gamma = 1$), condition (5) simplifies to $\frac{y_{D,t}-a_t}{y_{R,t}+a_t} = \frac{A}{1-A}$, which implies the following optimal rule for foreign aid disbursements: $a_t^* = (1 - A)y_{D,t} - Ay_{R,t}$. This is a special case of (4) with $\gamma = 1$. If $\gamma \neq 1$, then the first-order condition (5) becomes: $\left(\frac{y_{D,t}-a_t}{y_{R,t}+a_t}\right)^\gamma = \frac{A}{1-A}$ which, given the non-negativity condition on a_t , again yields the predicted optimal aid rule (4). ■

In the remainder of the paper, we focus on aid-dependent countries, i.e., countries for which aid flows are always positive ($a_t^* > 0 \forall t$). From (4), aid is positive as long as a simple condition is met: the donor's endowment must be sufficiently bigger than the recipient's. Accordingly:

Definition 1 (Aid Dependency) *We say that the recipient country is aid-dependent if:*

$$\frac{y_{D,t}}{y_{R,t}} > \left(\frac{A}{1-A}\right)^{1/\gamma} \forall t. \quad (6)$$

Corollary 1 : *For aid-dependent countries, the optimal aid policy is countercyclical from the recipient's viewpoint.*

Proof. By definition, an aid-dependent country has $a_t^* > 0 \forall t$. Now, recall that Φ_R and Φ_D are independent processes by assumption. Hence, $\frac{\partial a_t^*}{\partial y_{R,t}} = -\frac{A^{\frac{1}{\gamma}}}{(1-A)^{\frac{1}{\gamma}} + A^{\frac{1}{\gamma}}} < 0$. ■

Of particular relevance to our study is the extent to which the aid policy $a^* \equiv \{a_t^*\}_{t=1,2,\dots}$ can reduce consumption volatility in the recipient country. Intuitively, the donor should provide insurance to the recipient to an extent that reflects his own income (relative to the recipient's), altruism, and risk aversion. Proposition 1 and Corollary 1 together confirm this intuition. Our next proposition goes further. It shows that, if the recipient is aid-dependent, then the optimal aid policy does not merely reduce the volatility of the recipient's consumption stream, it goes as far as equalizing the variability of per capita consumption in the donor and recipient countries.

Precisely, Proposition 2 focuses on a logical volatility benchmark that is familiar from the literature on business cycles (Kydland & Prescott, 1982, for example): the standard deviation of per capita consumption expressed as a percentage of its mean. Let $\mu[X_t|\Omega_t]$ and $\sigma[X_t|\Omega_t]$ denote, respectively, the mean and standard deviation of any random variable X_t conditional on the donor and recipient's common information set at time t , Ω_t . Then:

Proposition 2 *In aid-dependent countries, the optimal aid policy drives the conditional volatility of per capita consumption down to that in the donor country.*

Proof. Again, by definition of an aid-dependent country, $a_t^* > 0 \forall t$. Given an optimal policy of strictly positive aid flows a_t^* , let $\frac{\sigma[y_{R,t}+a_t^*|\Omega_t]}{\mu[y_{R,t}+a_t^*|\Omega_t]}$ and $\frac{\sigma[y_{D,t}-a_t^*|\Omega_t]}{\mu[y_{D,t}-a_t^*|\Omega_t]}$ denote, respectively, the percentage

volatility of per capita consumption in the recipient and donor countries in period t conditional on the information set Ω_t . Now notice that any strictly positive optimal aid policy can be written as $a_t^* = (1 - \delta)y_{D,t} - \delta y_{R,t}$, where $\delta \equiv A$ ($\gamma = 1$) or $\delta \equiv \frac{A^{1/\gamma}}{(1-A)^{1/\gamma} + A^{1/\gamma}}$ ($\gamma \neq 1$). Hence:

$$\frac{\sigma[y_{R,t} + a_t^* | \Omega_t]}{\mu[y_{R,t} + a_t^* | \Omega_t]} = \frac{\sigma[(1 - \delta)(y_{D,t} + y_{R,t}) | \Omega_t]}{\mu[(1 - \delta)(y_{D,t} + y_{R,t}) | \Omega_t]} = \frac{\sigma[\delta(y_{D,t} + y_{R,t}) | \Omega_t]}{\mu[\delta(y_{D,t} + y_{R,t}) | \Omega_t]} = \frac{\sigma[y_{D,t} - a_t^* | \Omega_t]}{\mu[y_{D,t} - a_t^* | \Omega_t]}$$

The last equality proves the claim. ■

4 Discussion

Although very intuitive and simple, Proposition 2 sends a powerful message to policy makers that does not seem to have reached their ears yet. It basically states that a simple reallocation of aid flows across time can in fact bring the massive macroeconomic volatility in developing countries, which has been shown to have a deleterious effect on their welfare (Pallage & Robe, 2003), down to the much smoother ride enjoyed by rich countries. In essence, without changing the actual amounts given on average, it is possible to provide a substantial amount of insurance to poor countries. Given that the percentage volatility of consumption in poor countries is at least three to four times that in their developed counterparts, it should be possible to cut the former by at least two thirds.

Because this volatility reduction is so massive, it is worth discussing the conditions under which Proposition 2 holds. On the one hand, the analysis assumes that the donor's economy is larger than the recipient's. Our results do not suggest that a small donor country such as Luxembourg (population 454,000) should provide consumption smoothing services to China (population 1.3 billion). This is because a donor's ability to insure the recipient country is a function not only of the per-capita magnitude of macroeconomic shocks but also of how many people reside in that country. Likewise, it might not apply to developing countries that have a significantly faster rate of growth than the donor nations. This is because, if the former expect to catch up with the latter in a reasonable time frame, then the aid-dependency restriction in the proposition would be violated. On the other hand, despite these apparent limitations, Proposition 2 is quite general: as long as the recipient country is poor, its population is not huge, and its economy is growing at most at the same rate as the donor country's, it is easy to show that our results carry through. Much of Africa, plus most poor countries in other regions of the Southern Hemisphere, meet these three conditions.

Our results show that donors can use foreign aid to drastically reduce consumption volatility in developing countries. Are we remotely suggesting that residents in donor countries should subsidize policy mistakes by recipient governments? Whether aid should be conditioned on good recipient

policies is beyond the scope of this paper, and our model is not appropriate to discuss this issue. In the poor countries for which our analysis is most relevant, however, our abstracting from aid conditionality is without much loss of generality in that most of the observed macroeconomic fluctuations are due to exogenous shocks such as terms-of-trade or world-price changes.⁹

Finally, what is the impact of the optimal aid policy on the donor's welfare? In our model, the donor bears no cost of adopting the optimal aid policy a^* . In fact, he strictly benefits from it since this policy is his optimal choice. Yet, a^* is arguably not the observed aid policy: in contrast to the predictions of Corollary 1 and Proposition 2, the actual net foreign aid flows to developing countries are positively correlated with each recipient country's GDP (Pallage & Robe, 2001; Bulír & Hamann, 2003) and fiscal position (Bulír & Lane, 2002). In sub-Saharan Africa, for example, the contemporaneous correlation between real net aid inflows and real GDP per capita has been positive in all but three countries since 1970, and exceeds 0.3 in most cases. This observation begs the question: Are donors acting irrationally?

Not necessarily. Notwithstanding the seemingly poor record of foreign aid in boosting recipients' growth rates (Easterly, 2003), much of the aid to developing countries has long helped finance projects in the hope of improving their growth prospects and thereby reducing poverty. When aid is meant for growth, it is important to the donor that the most productive projects be the ones that end up being financed. However, the recipient government's own incentives might differ. To align these possibly divergent preferences, donor agencies often require that recipient governments contribute a matching grant that makes it costly for them to propose bad projects or to misuse the aid. Because aid recipients have fewer domestic resources in lean years, though, conditioning aid on matching-fund disbursements should lead to a positive correlation between aid inflows and local GDP. That is, the observed aid procyclicality could simply be the consequence of aid contracts meant to solve agency conflicts when growth is paramount and consumption volatility is, at best, viewed as a secondary issue. In contrast, our goal in this paper is to show that the specific timing of foreign aid flows has the potential to massively improve poor countries's welfare, i.e., aid volatility is of first-order importance. Our model therefore focuses on insurance considerations.

For the skeptic reader who might question the relevance of the donor's preferences in our model, and would rather see the altruistic parameter $1 - A$ driven down to 0, we could recast the question

⁹Kose & Riezman (2001) show that world price shocks account for about 82% of total consumption volatility in African countries. In the average developing country, that figure reaches 92% (Kose, 2002). These findings are consistent with earlier evidence that terms-of-trade shocks fluctuations are the main source of macroeconomic fluctuations in these countries (Mendoza, 1995).

of the welfare cost of the proposed policy in terms of its effect on the volatility of the donor's own consumption stream. If we focus on aid to African countries, to which our analysis plainly applies, the median ratio of official development assistance to donor GDP is less than 0.1% over the period 1969-1995 (Pallage & Robe, 2001). Clearly, any change in the sequencing of aid disbursements while maintaining that ratio in the long run will have only minimal impact on the volatility of consumption in the donor countries. Now, recall from the literature on the welfare cost of business cycles that, in industrialized countries, changing aggregate consumption volatility by more significant amounts has only trivial consequences on local residents' (selfish) welfare, regardless of the model economy used to assess these welfare effects (Lucas, 1987; Dolmas, 1998). Thus, whether one accepts our model preferences or not, it remains that, for donors, adopting policy a^* or sticking to their current policy makes no difference quantitatively.¹⁰

5 Computing the welfare implications of aid patterns

Sections 3 and 4 establish that, at no or negligible cost to the donor, the optimal timing of aid flows can reduce the very large percentage conditional volatility of per capita consumption in the recipient country down to the much lower level in the donor country. In this section, we quantify the welfare implications of this strong theoretical result for the recipient country.

Formally, we measure the welfare gain that would accrue to the recipient country's representative resident if the actual aid process, say \tilde{a} , were replaced by the optimal aid rule a^* derived in Propositions 1 and 2. In contrast to the latter, the actual aid flows to most developing nations are positively correlated with each recipient country's GDP. In the process of assessing the potential of foreign aid as an insurance device, we therefore quantify the relevance of this procyclicality for the welfare of the recipient country's residents.

We follow the usual approach and measure this welfare gain as the percentage consumption increase, in all states of the world and at all dates, that would make the recipient country's representative agent indifferent between the two reference environments (actual *vs.* optimal aid). In the context of our model, this computation amounts to computing the welfare effects of providing the agent with two alternative consumption series that have a similar mean but different volatilities: $c_{R,t}^* = y_{R,t} + a_t^*(y_{D,t}, y_{R,t})$ at the optimum, versus $\tilde{c}_{R,t} = y_{R,t} + \tilde{a}_t$ in practice.

In order to carry out these computations, we must make some assumption on the stochastic

¹⁰These welfare effects, already minor if the donor's entire aid goes to a single recipient country, should be even smaller if the donor gives aid to many countries – each subject to imperfectly correlated macroeconomic shocks.

processes that govern each country's consumption stream per capita, $\tilde{c}_{i,t}$ ($i = D, R$). We consider in turn two processes that are familiar from the literature on the welfare costs of business cycles. Since we are interested in percentage deviations from trend, we work throughout with the logarithms of the relevant series.

In the first process, the natural logarithm of real per capita consumption fluctuates randomly and shocks to consumption levels are temporary:

$$\ln \tilde{c}_{i,t} = \ln c_{i,0} + z_{i,t} \quad \text{with } z_{i,t} \rightsquigarrow \mathcal{N}(0, \sigma_{z_i}^2) \quad (i = D, R) \quad (7)$$

We parametrize σ_{z_i} in process (7) to the standard deviation of the cyclical component of the Hodrick-Prescott (HP)-filtered logarithms of real per capita private consumption. This approach is the same as that taken by Lucas (1987) and many others, and provides a floor for the welfare cost of macroeconomic fluctuations in an economy (Dolmas, 1998). This first process, thus parameterized, likewise provides a lower-bound estimate of the welfare gains from using foreign aid as insurance.

Let λ_{RtoD} denote the welfare gain to the recipient agent of a change in the timing of aid flows that reduces the percentage volatility of his periodic consumption, σ_{z_R} , down to the donor's consumption volatility, σ_{z_D} . Likewise, let λ denote the welfare gain to the same recipient of eliminating all macroeconomic fluctuation, i.e., of getting σ_{z_R} down to 0. Given that our economies are characterized by isoelastic preferences and consumption process (7), and given our parametrization of σ_{z_D} and σ_{z_R} , both λ_{RtoD} and λ have closed-form solutions (Lucas, 1987):

$$\begin{cases} \lambda_{RtoD} = e^{\frac{\gamma(\sigma_{z_R}^2 - \sigma_{z_D}^2)}{2}} - 1 \\ \lambda = e^{\frac{\gamma\sigma_{z_R}^2}{2}} - 1 \end{cases} \quad (8)$$

To the extent that the exogenous shocks that affect developing economies are persistent (Cashin, McDermott, & Pattillo, 2004) and that their impact on domestic consumption is not easily smoothed out (Kose, Prasad, & Terrones, 2003), process (7) will significantly underestimate the welfare gain that could be achieved by using foreign aid to reduce consumption volatility in poor countries. A simple way to account for shock persistence is to consider an alternative consumption process, in which both the level and the growth rate of per capita consumption are stochastic:

$$\ln \tilde{c}_{i,t} = \ln \tilde{c}_{i,t-1} + (1 + g_i - \frac{1}{2}\sigma_{\epsilon_i}^2) + \epsilon_{i,t} \quad \text{with } \epsilon_{i,t} \rightsquigarrow \mathcal{N}(0, \sigma_{\epsilon_i}^2) \quad (i = D, R) \quad (9)$$

where g_i is the mean growth rate of the series in country i . This second process assumes that shocks to consumption levels are persistent but shocks to the growth rate of consumption are not.

In a recipient economy characterized by CRRA utility and consumption process (9), the optimal aid policy identified in Proposition 1 effectively reduces the conditional standard deviation of shocks to per capita consumption from σ_{ϵ_R} to σ_{ϵ_D} . Let η_{RtoD} and η denote the welfare gains to the recipient of reducing σ_{ϵ_R} to σ_{ϵ_D} or to 0, respectively. Both η_{RtoD} and η have closed-form expressions (Obstfeld, 1994):

$$\eta_{RtoD} = \begin{cases} \left(\frac{1-\beta e^{(1-\gamma)(g_R - \gamma\sigma_{\epsilon_R}^2/2)}}{1-\beta e^{(1-\gamma)(g_R - \gamma\sigma_{\epsilon_D}^2/2)}} \right)^{\frac{1}{1-\gamma}} - 1 & (\gamma \neq 1) \\ \left(e^{\frac{\gamma(\sigma_{\epsilon_D}^2 - \sigma_{\epsilon_R}^2)}{2}} \right)^{\frac{\beta}{1-\beta}} & (\gamma = 1) \end{cases} \quad (10)$$

$$\eta = \begin{cases} \left(\frac{1-\beta e^{(1-\gamma)(g_R - \gamma\sigma_{\epsilon_R}^2/2)}}{1-\beta e^{(1-\gamma)g_R}} \right)^{\frac{1}{1-\gamma}} - 1 & (\gamma \neq 1) \\ \left(e^{-\frac{\gamma\sigma_{\epsilon_R}^2}{2}} \right)^{\frac{\beta}{1-\beta}} & (\gamma = 1) \end{cases} \quad (11)$$

5.1 Calibration

To further calibrate the two representative agents' preferences, we need values for the discount factor, β , and the coefficient of relative risk aversion, γ . Only annual consumption data are available for most developing countries, so we set $\beta = 0.96$ (which implies an annual discount rate of 4%). Our main results are qualitatively similar using the alternative values $\beta = 0.98$ and $\beta = 0.945$. For γ , we take values between 1 and 10, which correspond to levels of risk aversion within the ranges suggested for countries in both the Northern and the Southern hemispheres (Mehra & Prescott, 1985; Ostry & Reinhart, 1992; Reinhart & Végh, 1995).

We choose an intensity for the donor's altruism parameter, A , such that foreign aid's certainty-equivalent consumption equals 5% of the recipient's permanent consumption. This figure is consistent with three premises. (i) Including grants, the value of technical assistance, and concessional loans, the average aid-to-GDP ratio is 8%. That percentage is itself consistent with both the median aid-to-GDP ratio for 63 developing countries over the period 1969-1995, as computed in Pallage & Robe (2001) using annual Official Development Assistance (ODA) figures published by the Organization for Economic Co-operation and Development (OECD), and with the 1998 mean aid-to-GDP ratio for all aid-recipient countries, as reported by Bulír & Lane (2002) based on a survey of IMF desk economists. (ii) Between 1965 and 1994, the gross saving rate in aid-dependent countries has seldom exceeded 20% (Loayza, Lopez, Schmidt-Hebbel, & Servén, 1998), so that the average aid-to-consumption ratio is approximately 10%. (iii) In practice, aid flows are procyclical. Hence, expressed as a fraction of the recipient's permanent consumption, the certainty-equivalent of the

actual foreign aid inflows is much less than the average aid-to-consumption ratio. For simplicity, we set this certainty equivalent at half the latter, or 5%, which is a conservative estimate.

To calibrate the parameters of consumption processes (7) and (9), we use estimates from Pallage & Robe (2003). These figures are based on annual, 1968-1996, local-currency real private consumption data for 36 countries from the World Bank's *World Development Indicators* database. For process (7), the σ_{z_i} 's are parametrized to the standard deviation of the cyclical component of HP-filtered logarithms of real per capita consumption. For the donor country, we use $\sigma_{z_D} = 1.36\%$, the estimate for the United States. We want results for a representative recipient country, so we set $\sigma_{z_R} = 5\%$, which is the mean estimate for the developing-country sample, and check robustness to values from $\sigma_{z_R} = 4\%$ to $\sigma_{z_R} = 10\%$. The latter figure might seem very high but is typical for a third of the African countries in the sample.¹¹ For process (9), we need not only consumption volatility estimates for the donor and the recipient but also a mean growth rate estimate for the recipient. Pallage & Robe have estimates ranging from 0% to 3% for g_R and from 3% to 6% for σ_{ϵ_R} . We therefore pick $g_R = 1\%$ and $\sigma_{\epsilon_R} = 5\%$ as central values, and carry out robustness checks with σ_{ϵ_R} ranging from 3% to 6%. In contrast, estimates of σ_{ϵ_D} do not vary much across donors. Consequently, we take only one value: the U.S. estimate, $\sigma_{\epsilon_D} = 1.35\%$.

5.2 Quantitative results

As other papers have shown (Obstfeld, 1994; Pemberton, 1996; Dolmas, 1998), the welfare costs of consumption volatility vary greatly depending on whether consumption shocks are temporary, as in process (7), or persistent, as in process (9). It is therefore useful to provide a benchmark for the maximum possible welfare gains from changing aid flow patterns in a typical recipient country by first quantifying, for each consumption process, the welfare gain that would accrue to the recipient country's representative agent if *all* consumption volatility were eliminated. Figures 1.a and 1.b and the first row of Table 1 provide these benchmarks, respectively under processes (7) and (9). Precisely, Figure 1.a plots λ and Figure 1.b plots η , for various values of the recipient agent's risk aversion and of consumption volatility (respectively, σ_{z_R} and σ_{ϵ_R}). The cost estimates in Figure 1 are very large. For example, suppose that the coefficient of risk aversion is $\gamma = 2.5$ ($\gamma = 5$), which is well within the accepted range, and that $\sigma_{z_R} = 5\%$ or $\sigma_{\epsilon_R} = 5\%$. Then, we see from Table 1 that the cost of macroeconomic volatility is $\lambda = 0.32\%$ ($\lambda = 0.63\%$) of permanent consumption in the first model economy, in which shocks to consumption levels are transitory – but that same cost

¹¹All these volatility estimates were computed with a weight of 10 for the HP filter and, as in other papers on the cost of business cycles, under the assumption that the shocks in (7) are i.i.d.

skyrockets to $\eta = 5.9\%$ ($\eta = 9.4\%$) of permanent consumption in the second model economy, in which shocks are persistent.

A natural question, then, is what fraction of these large volatility costs can be eliminated by reallocating aid flows through time. For our central parameter values, the third row of Table 1 shows that this ratio is about 92%. Across the entire range of parameter values considered, Figure 2 shows that typically more than four fifths, and always at least three fourths, of the welfare costs of macroeconomic volatility could be eliminated by replacing the actual foreign aid stream \bar{a}_t by the optimal aid process a_t^* of Propositions 1 and 2 (i.e., by a foreign aid stream of the same average magnitude but different intertemporal properties).

To provide a second benchmark for the potential costs of ignoring the insurance implications of foreign aid, we express the welfare gains from replacing the actual foreign aid stream \bar{a} by the optimal aid process a^* (i.e., λ_{RtoD} in Figure 3.a or η_{RtoD} in Figure 3.b) as a fraction of the gain from providing \bar{a} in the first place. We have conservatively assumed that \bar{a} is worth 5% of permanent real per capita consumption in the recipient country. Thus, a value of 1 on the vertical axis means that changing the timing of aid is worth as much as increasing the recipient's per capita consumption by 5% at all dates and in all possible states of the world.

Figure 3.b shows that, if consumption shocks are persistent, then improving the timing of foreign aid flows is often worth as much to the recipient as receiving aid to start with. For example, suppose that $\sigma_{\epsilon_R} = 5\%$ – a moderate level of consumption volatility in a developing country. Then Table 1's second row shows that, even if the recipient's risk aversion is only $\gamma = 2.5$, the welfare gain from improving the timing of foreign aid flows would be worth 5.5% of his permanent consumption!

6 Conclusions

Some policy makers and commentators have argued that:

“Rich countries must (...) pay more attention to the sudden economic shocks to which poor countries are vulnerable. Sudden plunges in the price of coffee or cotton pluck fewer heartstrings than floods or earthquakes, but can be much more destabilising and (make civil war) more likely. Aid should take such shocks into account.”

The Economist, editorial of May 24, 2003.

We argue in this paper that, even ignoring the possibility that macroeconomic volatility may bring about political mayhem or worse, it is a mistake to disregard foreign aid's potential as an

insurance mechanism against macroeconomic fluctuations. We show that merely changing the timing of aid flows can provide recipient countries with substantial insurance against aggregate consumption shocks. Indeed, we find that at least three fourths of the large welfare costs of macroeconomic fluctuations in developing countries could be eliminated by optimally timing aid disbursements to smooth out aggregate consumption in these countries.

Because we focus on the effects of macroeconomic shocks on consumption levels but ignore the link between economic stability and political stability, our results likely underestimate the benefits from countercyclical aid flows. Unfortunately, foreign aid flows *are* procyclical in practice. Our results therefore have two alternative interpretations. First, if aid procyclicality is the outcome of aid contracts designed to alleviate conflicts between donors and recipients under asymmetric information, then our results provide a lower bound estimate of the size of the inefficiencies brought about by these agency conflicts. Second, if the type of aid contracts that bring about aid procyclicality could be replaced by others that do not (for example, through better donor coordination), then our results suggest that the resulting welfare gains might be as large as the gains from receiving aid in the first place.

In sum, our point is simple. We know that the welfare costs of consumption volatility are very large in developing countries. Most of that volatility is the direct consequence of exogenous macroeconomic shocks. Lucas (2003) rightly argues that, regardless of the size of its costs, one should care about consumption volatility only to the extent that it can be reduced by feasible policies. In this paper, we focus on a policy that is feasible and widely used: foreign aid. We show that altering foreign aid flows can, at no or negligible cost to the donors, eliminate between 75% and 95% of the welfare cost of the massive macroeconomic fluctuations that affect developing countries. In light of these results, we submit that the role of foreign aid as an insurance device deserves much greater attention than it has received in the past.

References

- Arellano, C., Bulír, A., Lane, T., & Lipschitz, L. (2002). Aid and tradable goods in aid-dependent countries. Working Paper, Duke University, September.
- Asiedu, E. & Villamil, A. P. (2002). Imperfect enforcement, foreign investment, and foreign aid. *Macroeconomics Dynamics*, 6(4), 476–95.
- Azam, J. & Laffont, J. (2003). Contracting for aid. *Journal of Development Economics*, 70, 25–58.
- Boone, P. (1994). The impact of foreign aid on savings and growth. Center for Economic Performance Working Paper 1265, London, October.
- Boone, P. (1996). Politics and the effectiveness of aid. *European Economic Review*, 40(2), 289–329.
- Bulír, A. & Hamann, A. J. (2003). Aid volatility: An empirical assessment. *IMF Staff Papers*, 50(1), 64–89.

- Bulír, A. & Lane, T. (2002). Aid and fiscal management. Working Paper 02/112, International Monetary Fund, June.
- Burnside, C. & Dollar, D. (2000). Aid, policies, and growth. *American Economic Review*, 90(4), 847–68.
- Caballero, R. (2003). The future of the IMF. *American Economic Review – AEA Papers and Proceedings*, 93(2), 31–8.
- Casella, A. & Eichengreen, B. (1996). Can foreign aid accelerate stabilisation. *Economic Journal*, 106(436), 605–19.
- Cashin, P., McDermott, C. J., & Pattillo, C. C. (2004). Terms-of-trade shocks in Africa: Are they long-lived or short-lived. *Journal of Development Economics*, 73(4), 727–44.
- Chang, C. C., Fernandez-Arias, E., & Servén, L. (1999). Measuring aid flows: A new approach. Policy Research Working Paper No. 2050, World Bank.
- Collier, P. & Dehn, J. (2001). Aid, shocks, and growth. Policy Research Working Paper No. 2688, World Bank, October.
- Dollar, D. & Easterly, W. (1999). The search for the key: Aid, investment and policies in Africa. *Journal of African Economies*, 8(4), 546–77.
- Dolmas, J. (1998). Risk preferences and the welfare cost of business cycles. *Review of Economic Dynamics*, 1(3), 646–76.
- Easterly, W. (1999). The ghost of financing gap: Testing the growth model used in international financial institutions. *Journal of Development Economics*, 60(2), 423–38.
- Easterly, W. (2003). Can foreign aid buy growth. *Journal of Economic Perspectives*, 17(3), 23–48.
- Easterly, W., Levine, R., & Roodman, D. (2003). New data, new doubts: A comment on Burnside and Dollar’s “aid, policies, and growth” (2000). *American Economic Review*. Forthcoming.
- Gupta, S., Clements, B., & Tiongson, E. (2003). Foreign aid and consumption smoothing: Evidence from global food aid. Working Paper 03/40, International Monetary Fund, February.
- Hansen, H. & Tarp, F. (2000). Aid effectiveness disputed. *Journal of International Development*, 12, 375–98.
- Hansen, H. & Tarp, F. (2001). Aid and growth regressions. *Journal of Development Economics*, 64(2), 547–70.
- Kose, M. A. (2002). Explaining business cycles in small open economies: How much do world prices matter. *Journal of International Economics*, 56(2), 299–327.
- Kose, M. A., Otrok, C., & Whiteman, C. H. (2003a). International business cycles: World, region, and country-specific factors. *American Economic Review*, 93(4), 1216–39.
- Kose, M., Prasad, E., & Terrones, M. E. (2003b). How does globalization affect the synchronization of business cycles. *American Economic Review – AEA Papers and Proceedings*, 93(2), 57–62.
- Kose, M. A. & Riezman, R. (2001). Trade shocks and macroeconomic fluctuations in Africa. *Journal of Development Economics*, 65(1), 55–80.
- Kydland, F. E. & Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6), 1345–70.
- Lahiri, S. & Raimondos-Moller, P. (1997a). Competition for aid and trade policy. *Journal of International Economics*, 43(3–4), 369–85.
- Lahiri, S. & Raimondos-Moller, P. (1997b). On the tying of aid to tariff reform. *Journal of Development Economics*, 54(2), 479–91.
- Lahiri, S. & Raimondos-Moller, P. (2000). Lobbying by ethnic groups and aid allocation. *Economic Journal*, 110(462), C62–79.

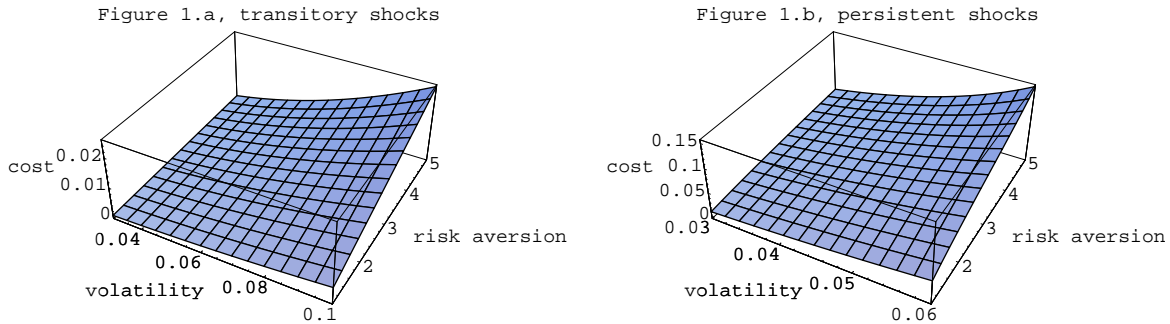
- Loayza, N., Lopez, H., Schmidt-Hebbel, K., & Servén, L. (1998). Saving in the world: the stylized facts. Working Paper, World Bank, November.
- Lucas, R. E. J. (1987). *Models of Business Cycles*. Yrjo Jahnsson Lectures. New York: Blackwell.
- Lucas, R. E. J. (2003). Macroeconomic priorities. *American Economic Review*, 93(2).
- Mehra, R. & Prescott, E. S. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2), 145–61.
- Mendoza, E. (1995). The terms of trade, the real exchange rate, and economic fluctuations. *International Economic Review*, 36(1), 101–37.
- Murshed, S. M. & Sen, S. (1995). Aid conditionality and military expenditure reduction in developing countries: Models of asymmetric information. *Economic Journal*, 105(429), 498–509.
- Obstfeld, M. (1994). Evaluating risky consumption paths: The role of intertemporal substitutability. *European Economic Review*, 38(7), 1471–1486.
- Orphanides, A. (1996). Optimal reform postponement. *Economics Letters*, 52(3), 299–307.
- Ostry, J. D. & Reinhart, C. M. (1992). Private savings and terms of trade shocks: evidence from developing countries. *IMF Staff Papers*, 39(3), 495–517.
- Otrok, C. (2001). On measuring the welfare cost of business cycles. *Journal of Monetary Economics*, 47(1), 61–92.
- Pallage, S. & Robe, M. A. (2001). Foreign aid and the business cycle. *Review of International Economics*, 9(4), 641–72.
- Pallage, S. & Robe, M. A. (2003). On the welfare cost of economic fluctuations in developing countries. *International Economic Review*, 44(2), 677–98.
- Pallage, S. & Zimmermann, C. (2001). Voting on unemployment insurance. *International Economic Review*, 42(4), 903–924.
- Pemberton, J. (1996). Growth trends, cyclical fluctuations, and welfare with non-expected utility preferences. *Economics Letters*, 50, 387–92.
- Prati, A., Sahay, R., & Tressel, T. (2003). Is there a case for sterilizing foreign aid inflows. Mimeo, Research Department, International Monetary Fund, July.
- Reinhart, C. M. & Végh, C. A. (1995). National interest rates, consumption booms, and lack of credibility: A quantitative examination. *Journal of Development Economics*, 46(3), 357–78.
- Svensson, J. (2000). Foreign aid and rent-seeking. *Journal of International Economics*, 51(2), 437–61.
- Svensson, J. (2003). Why conditional aid does not work and what can be done about it. *Journal of Development Economics*, 70(2), 381–402.
- The Economist (2003). The poor man’s curse: Civil wars. *Editorial*, 367(8325), 11.
- Wang, C. & Williamson, S. D. (2002). Moral hazard, optimal unemployment insurance, and experience rating. *Journal of Monetary Economics*, 49(7), 1337–71.
- World Bank (2002). *Global Development Finance*. Washington: World Bank.

Table 1: Consumption volatility, optimal aid timing, and welfare

	Temporary shocks		Persistent shocks	
	$\gamma = 2.5$	$\gamma = 5$	$\gamma = 2.5$	$\gamma = 5$
Consumption volatility costs (% of permanent consumption)	0.32	0.63	5.86	9.36
Gain from timing aid optimally (% of permanent consumption)	0.29	0.58	5.44	8.76
Fraction of volatility costs eliminated by optimal aid timing (%)	92.6	92.6	92.8	93.6

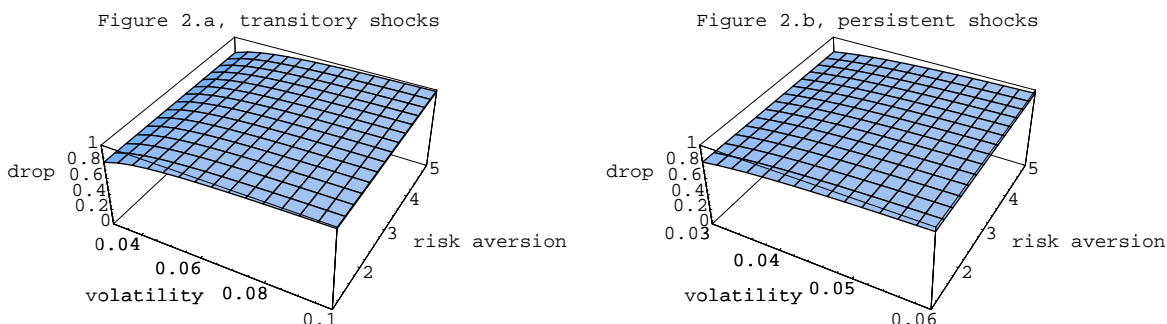
Note: The first row in Table 1 shows the welfare costs of aggregate consumption volatility in a typical recipient country, expressed in terms of permanent consumption. The second row shows the welfare gain from replacing the actual foreign aid stream \bar{a}_t with the optimal aid process a_t^* of Propositions 1 and 2 (i.e., by foreign aid flows with the same trend but different cyclical properties). That gain is also measured in terms of permanent consumption in the recipient country. The third row gives the ratio of the previous two, i.e., it shows the fraction of the welfare costs of aggregate consumption volatility that could be eliminated by replacing the actual foreign aid stream \bar{a}_t with the optimal aid process a_t^* . The figures in all rows are provided for two levels of risk aversion in the recipient country ($\gamma = 2.5$ and $\gamma = 5$), under the assumption that the percentage annual volatility of consumption (σ_{z_i} if shocks are assumed temporary, or σ_{ϵ_i} if they are assumed persistent) is 5% in the recipient country ($i = R$) and 1.35% in the donor country ($i = D$).

Figure 1: Welfare costs of macroeconomic volatility (% of permanent consumption)



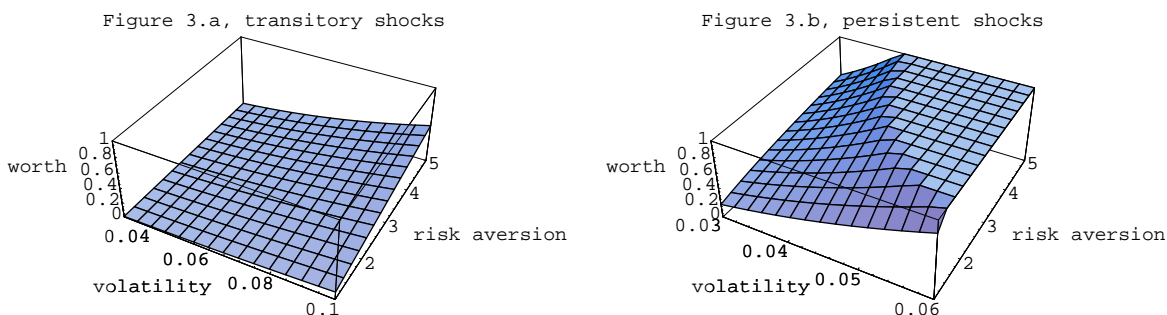
Note: Figure 1 plots the welfare costs of macroeconomic volatility (λ in Figure 1.a and η in Figure 1.b) in the recipient country as a percentage of permanent consumption, for typical levels of risk aversion (γ) and consumption volatility (σ_{z_R} in Figure 1.a and σ_{ϵ_R} in Figure 1.b). A value of 0.01 (1%) on the vertical axis means that the country's representative resident would be indifferent between seeing all consumption volatility removed and getting an extra 1% of consumption in all states and at all dates. Figure 1.a plots the value taken by λ in Equation (8), when consumption shocks are not persistent. Figure 1.b plots the value taken by η in Equation (10), when consumption shocks are persistent, for a mean growth rate of per capita real consumption $g_R = 1\%$.

Figure 2: Fraction of macroeconomic volatility costs eliminated by timing aid flows optimally



Note: Figure 2 plots, for typical levels of risk aversion and consumption volatility in recipient countries, the fraction of the welfare costs of macroeconomic volatility that could be eliminated by replacing the actual foreign aid stream \tilde{a}_t with the optimal aid process a_t^* of Propositions 1 and 2 (i.e., by foreign aid flows with the same mean but different cyclical properties). A drop value of 0.8 on the vertical axis means that 80% of the cost λ [in Figure 2.a, when consumption follows (7)] or η [in Figure 2.b, when consumption follows (9)] could be eliminated by reallocating aid flows optimally through time. Figure 2.a and Figure 2.b are drawn under the assumption that the percentage volatility of consumption in the donor country (respectively, σ_{z_R} or σ_{ϵ_R}) is 1.36%. Figure 2.b assumes a mean growth rate of per capita real consumption in the recipient country of $g_R = 1\%$.

Figure 3: What is improving the timing of aid worth, compared to getting aid in the first place?



Note: Figure 3 displays, for typical levels of risk aversion and consumption volatility in recipient countries, the welfare gain that would accrue residents of a recipient country if the actual foreign aid stream \tilde{a}_t were replaced by the optimal aid process a_t^* of Propositions 1 and 2 (i.e., by foreign aid flows with the same mean but different cyclical properties). A value of 1 on the vertical axis means that changing the timing of aid is worth as much as increasing the recipient's per capita consumption by 5% at all dates and in all possible states of the world. Figure 3.a plots the value taken by $\lambda_{Rt0D}/5\%$ in Equation (8), when consumption shocks are not persistent. Figure 3.b plots the value taken by $\eta_{Rt0D}/5\%$ in Equation (10), when consumption shocks are permanent. Both Figures are drawn under the assumption that the percentage volatility of consumption in the donor country (respectively σ_{z_R} or σ_{ϵ_R}) is 1.36%. Figure 3.b assumes a mean growth rate of per capita real consumption in the recipient country of $g_R = 1\%$.