Dynamic aid allocation

Patrick Carter, Fabien Postel-Vinay, Jonathan Temple

Journal of International Economics 95 (2015) 291–304

Dynamic aid allocation☆

Patrick Carter [a,b], Fabien Postel-Vinay [c,d,e], Jonathan Temple [b,*]

[a] Overseas Development Institute, 203 Blackfriars Road, London SE1 8NJ, UK
[b] University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK
[c] University College London, 30 Gordon Street, London WC1H 0AX, UK
[d] Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE, UK
[e] Sciences Po, 27 rue Saint-Guillaume, 75337 Paris Cedex 07, France

ARTICLE INFO

Article history:
Received 5 July 2013
Received in revised form 22 October 2014
Accepted 4 November 2014
Available online 14 November 2014

JEL classification:
F35
O41

Keywords:
Foreign aid
Aid allocation
Global redistribution

ABSTRACT

This paper introduces a framework for studying the optimal dynamic allocation of foreign aid among multiple recipients. We pose the problem as one of weighted global welfare maximization. A donor in the North chooses an optimal path for international transfers, anticipating that consumption and investment decisions will be made by optimizing households in the South, and accounting for limits in the extent to which recipients can effectively absorb aid. We present quantitative results on optimal aid policy by applying our approach to a neoclassical growth model, where the scope for aid-funded growth is determined by the recipients' distance from steady-state.

1. Introduction

Since the 1969 Pearson Commission, there has been a standard benchmark for the generosity of foreign aid programs: developed countries should donate at least seven-tenths of one per cent of their GDP. In practice, however, relatively few countries have achieved this level of generosity. Moreover, its foundations are easy to question. The target was calculated more than forty years ago using a combination of the Harrod–Domar model and financial programming. This rather mechanical approach would not command much support today.

In this paper, we introduce a new framework for studying optimal aid policies. We model foreign aid as a form of global redistribution. A utilitarian, forward-looking social planner seeks to maximize a weighted average of welfare in the global North and the global South. The planner decides on an optimal path of international transfers, anticipating that consumption and investment decisions in the North and South will be made by optimizing households within each region. These households cannot borrow or lend internationally, and the scope for global redistribution is limited by diminishing returns to aid. This framework can be used to study the optimal generosity of aid, and its relationship to absorptive capacity. It can also be used to inform the timing of aid, and its allocation across countries that differ in development levels. We describe circumstances in which donors should seek to increase generosity over time, and examine whether middle-income countries should ever be a priority for aid. Importantly, the framework is tractable and could be extended in many directions; for example, it could be adapted to study the implications of capital mobility for optimal aid policies.

The starting point is to model both the global North and South as neoclassical Ramsey economies. These economies differ in their levels of TFP and income, and in their distances from steady-state. Obstfeld (1999) analyzed the effect of exogenous transfers on a Southern economy; in this paper, that problem will be nested within the problem facing a Northern donor, so that the level and timing of transfers will be endogenously determined. At first glance, the Ramsey model may seem too stylized for this purpose, since it neglects political economy forces that will often be central to aid effectiveness. But the Ramsey model casts sharp light on a direct consequence of aid flows, which is to relax intertemporal resource constraints. Studying this in isolation should enhance our understanding of the choices facing donors.

A further motivation is the growing interest in cash transfers to households as a form of poverty alleviation. Hanlon et al. (2010) argue that transfers direct to households would be more beneficial than

☆ This paper was presented as a keynote address at the June 2013 conference of the Global Political Economy Network. We thank the editor, two anonymous referees, conference participants, Mario Crucini, Nigel Duck, Huw Edwards, Ben Zissimos, and seminar audiences at Bristol and Brunel for comments and suggestions. We are especially grateful to Francesco Caselli for detailed comments on an earlier version, but we are responsible for the remaining shortcomings of the paper. Carter thanks the ESRC and British Academy for financial support under awards ES/I901876/1 and PF110079 respectively, and the Royal Economic Society for a Junior Fellowship. Temple thanks the British Academy for support from a Mid-Career Fellowship under award MD120067.
* Corresponding author.
E-mail address: jon.temple@bristol.ac.uk (J. Temple).

http://dx.doi.org/10.1016/j.jinteco.2014.11.005

more traditional forms of aid. In 2013, the Indian government launched an ambitious Direct Benefit Transfer scheme, intended eventually to replace multiple welfare programs with cash transfers to households. Using evidence from randomized trials, Gertler et al. (2012) and Haushofer and Shapiro (2013) find that cash transfers to poor households are partly invested. It is therefore interesting to ask: what happens when aid is used to relax household budget constraints, and what are the implications for optimal aid policies?

In our framework, utility functions are concave and some global redistribution is optimal, but its extent will be constrained by diminishing returns to aid. Obstfeld (1999) concluded that the welfare benefit of (exogenous) aid was modest even without an aid absorption constraint, but our analysis turns that logic on its head. If the donor is also seen as a Ramsey economy, the opportunity cost of aid for the donor is similarly modest. Hence, we sometimes find that donors should be generous, especially for recipient economies that are close to subsistence. The welfare impact of aid may be substantial, and aid can be justified even when a substantial fraction of it is wasted.

To investigate the quantitative implications, we use simulations. Donor and recipient initial conditions are based on data from the Penn World Table, combined with assumptions on structural parameters. We consider isoelastic (CRRA) preferences and Stone–Geary 'subsistence consumption' preferences. Under Stone-Geary preferences, the effects of aid on investment and growth are especially strong, and the donor maintains a higher level of aid generosity for a longer period.

The limits to recipient absorptive capacity play an important role throughout. When recipients are capable of absorbing relatively high levels of aid, the optimal path of aid is typically front-loaded. But if the South lacks absorptive capacity, the North may want to increase the generosity of its aid over time, relative to Northern GDP. This reflects our assumption that, as the South develops, it can use aid more effectively. We examine what happens if the donor nevertheless maintains aid as a fixed share of its GDP, as in the Pearson Commission benchmark, and compare this to the fully-flexible optimal path. We show that, perhaps surprisingly, the welfare costs of restricting aid to a fixed proportion of donor GDP are often borne in equilibrium by the North, rather than the South.

Another issue of recent interest has been whether donors should make transfers to middle-income countries (see, for example, Kanbur and Sumner, 2012). We therefore study what happens in the case of two aid recipients, where one is a low-income country and one a middle-income country with a larger population. Our simulations of optimal policies indicate large changes over time in aid generosity and in the division of aid between recipients. Given diminishing returns to aid intensity, most of the aid may be directed at the middle-income recipient initially, with the allocation later switching towards the smaller, poorer recipient.

Throughout the analysis, we assume that the capital account is closed. This assumption is common in the literature, but some commentators argue that poor countries do not need aid when investment can be financed by capital inflows. In practice, capital has tended to flow to middle-income countries rather than the poorest countries. And even when the capital account is open, there is still a role for aid, to finance both consumption and the accumulation of assets owned domestically. It is likely that the welfare gains from aid would be less in this case, since initial consumption is higher; nevertheless, financing higher consumption would remain valuable. The quantitative investigation of this would be an interesting topic for future research.[1]

Since the model is stylized, it is worth clarifying the intended contribution. The paper provides a new way of framing the decision problems facing aid donors. The model could be extended, and made more realistic, in many directions; it is a first step towards richer quantitative frameworks that could inform future aid policies. The current simulations serve two more limited aims. The first is to investigate some qualitative results about aid policies when absorption constraints matter: these results are not special cases, but can emerge as important under reasonable parameter assumptions. The second aim is to illustrate what might be learnt by future research using more complicated models. For example, the current analysis is too preliminary to suggest a replacement for the Pearson Commission benchmark, but it draws attention to some relevant considerations. A companion paper, Carter (2014), uses related ideas to study aid allocation rules of the type sometimes implemented by donors.

The existing literature has generally considered simpler models of North–South interactions, often in models with just one or two periods. The majority of this work is theoretical, and studies the effects of exogenous transfers rather than deriving their optimal time path. Some papers explore the transfer problem, and especially the possibility of transfer paradoxes driven by terms-of-trade effects; Eaton (1989) surveys the older literature. An alternative approach is taken by Chamon and Kremer (2009), who construct a multi-country model in which developing countries gradually integrate with the world economy. They discuss the potential role of aid in accelerating this process, but aid is not included in the version of the model they calibrate.

The paper is structured as follows. The next section provides the formal description and analysis of the donor's decision problem. Section 3 describes the assumptions used in our simulations. Section 4 presents initial analysis of the optimal path of aid and Section 5 extends the analysis to subsistence economies. Section 6 considers two recipients which differ in population size and income. Section 7 carries out a sensitivity analysis and discusses potential extensions, before Section 8 concludes.

2. The model

Our paper deliberately takes a narrow view of the donor's problem, seen exclusively in terms of international resource transfers. We hope to show that even this narrow view could inform the design of aid programs. Since the 1960s, cumulative spending on foreign aid has exceeded three trillion dollars in nominal terms, a figure that would be even higher in today's prices. Yet basic issues remain contested and, in some cases, rarely studied. How generous should aid flows be, and should donors choose aid targets relative to donor resources, or to recipient GDP? To what degree should this generosity be greater, early in the development process? When allocating aid across multiple recipients, how sensitive should aid flows be to recipient income levels? These are all debates that can be informed by the approach that we develop here.

We use the framework to revisit some long-standing questions. Some of these questions concern the Pearson Commission benchmark, which appeared in United Nations Resolution number 2626, from October 1970. The 0.70% target was originally justified using a financing gap calculation, based on assessed capital needs for a Harrod–Domar economy (see Clemens and Moss, 2007, for a full history). It seems unsatisfactory to base aid policies on a model and form of analysis that are clearly outdated. Even the basic assumption that aid should be a fixed share of donor GDP might not be supported by a more rigorous approach. We agree with Clemens and Moss (2007) that it seems backwards to determine aid levels based on the size of donors, rather than conditions in recipient economies. The optimal degree of redistribution should be dependent on the extent of between-country differences, the nature of preferences, and the extent of absorption constraints, all of which play a role in our analysis.

We formulate the North's decisions in terms of an explicit dynamic optimization problem. The North decides on an optimal path of transfers to the South. As in Obstfeld (1999), aid will benefit Southern economies in two ways. Aid accelerates the rate at which the South converges to its steady-state, and allows the South to sustain a higher level of consumption than otherwise. Stripping the aid problem down to

[1] It is surprising that few papers on international transfers have considered them alongside capital mobility; exceptions include Galor and Polemarchakis (1987). Growth models under capital mobility were discussed by Rebelo (1992). The study by Caselli and Feyrer (2007) finds that returns to capital are broadly comparable in a sample of 53 countries, but their sample contains relatively few low-income countries, and only seven countries are in sub-Saharan Africa. In the later study by Lowe et al. (2013), even their extended sample has only ten low-income countries.

these two roles will clarify their implications for aid policies and allocation decisions. In our simulations, we find that accelerating convergence to steady-state plays only a minor role. This does not mean the aid is wasted, however. In a model of this type, aid is effective to the extent that it raises consumption immediately, later, or both.[2]

In our framework, the Northern donor seeks to maximize the following objective function:

$$\max_{\{c_{Nt},a_t\}} \int_0^\infty u(c_{Nt})e^{-\rho t}dt + \omega \cdot L_S \int_0^\infty u(c_{St})e^{-\rho t}dt$$

where c_{Nt} is per capita consumption in the North, c_{St} is per capita consumption in the South, $u(\cdot)$ is the instantaneous utility function, a_t is aid, ω is the relative weight that the North places on the utility of the South, L_S is the relative population of the South, and ρ is the discount rate. The choice of aid will matter through intertemporal resource constraints: aid transfers will reduce the resources available to the North, and increase those available to the South. Taking this into account, we derive the path for aid which maximizes the objective function when all households are optimizing over time. We achieve this by nesting the South's optimal control problem within the optimal control problem of the North.

Framing the donor's problem as one of weighted global welfare maximization has several advantages. First, the opportunity costs of aid arise endogenously from the structure of the model. Second, our approach provides a mapping between an explicit weight ω on Southern utility and optimal aid policies.[3] We can readily study normative questions: how generous should aid be, and what does the optimal time profile look like?

We now set out the decision problems formally. The North and South are both characterized as Ramsey economies. We study the South's problem first, and then the optimal control problem facing the Northern donor, where the optimizing behavior of the South represents a constraint in the North's decision problem. To simplify the presentation, initially we set the rates of technical progress and population growth to zero, but the necessary extensions are straightforward.

We assume that aid is ultimately distributed to Southern households, and these are individually too small to internalize the effects of their actions on donor policies. Hence, we can consider the donor's problem without needing to allow strategic interactions between donors and (multiple) recipients. That would require analysis in terms of a dynamic game, and would not be straightforward. For the same reason, we do not model political economy forces explicitly, but these simplifications allow an analysis that is richer in other ways.

2.1. The decision problem for the South

Decision problems for the South have previously been considered by Chatterjee et al. (2003), Obstfeld (1999) and Turnovsky (2009), among others. In our case, the representative Southern household is assumed to solve:

$$\begin{aligned} &\max_{\{c_{St}\}} \int_0^\infty u(c_{St})e^{-\rho t}dt \\ \text{subject to}: \; &\dot{k}_{St} = A_S f_S(k_{St}) + g(a_t, \overline{\varkappa}_{St}) - c_{St} - \delta k_{St} \\ &k_{S0} \;\text{given}, \end{aligned}$$

where k_{St} is household capital per worker and $\overline{\varkappa}_{St}$ is aggregate capital per worker in the South, and a_t is aid. The non-standard aspect of the South's decision problem is the role for absorption constraints, in reduced form. Aid enters the South's intertemporal resource constraint through the 'aid impact function' $g(\cdot)$. This function depends on endogenous variables, aid and capital, and also on exogenous variables such as the relative population and TFP of the South; we typically suppress the dependence on exogenous variables for notational simplicity. We assume that $g(0, \overline{\varkappa}_{St}) = 0$, and that $g(\cdot)$ is non-decreasing in $\overline{\varkappa}_{St}$ and has strictly diminishing marginal returns to a_t.

Political economy approaches often imply that some aid is wasted or diverted. Rather than develop a specific structural model, our aid impact function translates Northern aid donations into the actual transfers received by Southern households, allowing for wastage or diversion.[4] Our formulation nests the case where the marginal benefit of aid is declining in aid intensity, the ratio of aid to recipient GDP. This idea has often been investigated empirically, as in Burnside and Dollar (2000) and Clemens et al. (2012).[5] One possible story is that, as aid intensity increases, so does the proportion of aid that is wasted or appropriated by a local elite. High aid intensity may also lead to Dutch Disease effects. It could have adverse effects on a domestic political equilibrium, partly by influencing rents to sovereignty, and perhaps by undermining long-term accountability and state capacity. A proliferation of aid projects and programs could overwhelm the capacity of the recipient government. These mechanisms have been widely discussed (see, for example, Temple (2010)). As a consequence, it seems essential to allow for limits to absorptive capacity. In our framework, these limits are eased by growth: as the GDP of an aid recipient increases, the recipient can use a given amount of aid more effectively.

The representative Southern household takes the paths of aid a_t and aggregate capital per worker $\overline{\varkappa}_{St}$ as given (hence, the Southern households do not internalize the effect of their investment decisions on future aid absorption). The above problem leads to the Euler equation:

$$\frac{\dot{c}_{St}}{c_{St}} = -\frac{1}{\varepsilon_{u'(c)}}\left\{A_S f'_S(k_{St}) - \rho - \delta\right\}. \tag{1}$$

where $\varepsilon_{u'(c)}$ is the elasticity of marginal utility with respect to consumption. Note that this Euler equation is standard, precisely because the Southern households do not internalize the economy-wide benefits of their investment decisions. Hence, the Southern households invest too little; a Southern planner would front-load investment to improve future aid absorption. The South's Euler equation would then be given by:

$$\frac{\dot{c}_{St}}{c_{St}} = -\frac{1}{\varepsilon_{u'(c)}}\left\{A_S f'_S(k_{St}) + g'_2(a_t, \overline{\varkappa}_{St}) - \rho - \delta\right\}. \tag{2}$$

where $k_{St} = \overline{\varkappa}_{St}$ at each instant.

In the first version of the South's problem, the donor's optimal policy is not time-consistent, due to the externality. We follow much of the literature on conditionality, and assume that the donor has access to a commitment technology. This could take the form of explicit commitments to aid made through domestic legislation or international agreements that would be costly to reverse. This simplifies the analysis, and is relatively natural given our interest in normative questions. Further, in the cases considered below, the quantitative effects of the externality will be modest.

2.2. The decision problem for the North

We now consider the North's decision problem, first for the case with a Southern planner, so that Eq. (2) is the relevant Euler equation. The North's decisions must respect $k_{St} = \overline{\varkappa}_{St}$. The Northern planner

[2] For poorer countries, higher consumption should hardly be considered wasteful: it is likely to mean improvements in nutrition, shelter, and personal health, among other welfare benefits. Moreover, these forms of consumption may influence current or future productivity, as in Steger (2000, 2002).

[3] These two advantages contrast with, for example, Scholl (2009). In several respects her analysis is richer than ours, but it is less well-suited to studying normative questions.

[4] Lahiri and Raimondos-Møller (1997) used a similar idea in their study of aid allocation across economies that vary in their trade policies. The terminology 'aid impact function' is taken from Dudley and Montmarquette (1976), which used a related idea. Kemp and Wong (1993) considered the transfer problem when a fixed proportion of aid is wasted.

[5] See Sumner and Mallett (2013) and Temple (2010) for further discussion and references.

then solves:

$$\max_{\{c_{Nt},a_t\}} \int_0^\infty [u(c_{Nt}) + \omega \cdot L_S \cdot u(c_{St})] \cdot e^{-\rho t} dt$$

$$\text{subject to}: \quad \dot{k}_{Nt} = f_N(k_{Nt}) - c_{Nt} - a_t - \delta k_{Nt}$$
$$\dot{k}_{St} = A_S f_S(k_{St}) + g(a_t, k_{St}) - c_{St} - \delta k_{St}$$
$$\frac{\dot{c}_{St}}{c_{St}} = -\frac{1}{\varepsilon_{u'(c)}} \{A_S f'_S(k_{St}) + g'_2(a_t, k_{St}) - \rho - \delta\},$$
$$k_{S0},\ k_{N0}\ \text{given}, c_{S0}\ \text{free}$$

where we have normalized the TFP of the North to unity.[6] This formulation rules out time-zero transfers of part of the North's capital stock: see Kemp et al. (1990). We can now derive some results. We denote the current-value costate variables associated with k_{Nt}, k_{St} and c_{St} as x_{Nt}, x_{St} and z_{St}, respectively. In this context, each costate variable can be interpreted as the (current-value) shadow value the North would attach to a one-unit increment in the corresponding state variable. The associated optimality conditions are as follows.[7] The first-order condition for c_{Nt} is the usual one: $u'(c_{Nt}) = x_{Nt}$. The first-order condition for a_t is:

$$x_{Nt} = g'_1(a_t, k_{St}) x_{St} - \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} g''_{12}(a_t, k_{St}). \tag{3}$$

The three Euler equations are:

$$\dot{x}_{Nt} = \{\rho + \delta - f'_N(k_{Nt})\} x_{Nt} \tag{4}$$

$$\dot{x}_{St} = \{\rho + \delta - A_S f'_S(k_{St}) - g'_2(a_t, k_{St})\} x_{St} + \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} \left\{A_S f''_S(k_{St}) + g''_{22}(a_t, k_{St})\right\} \tag{5}$$

$$\dot{z}_{St} = \left\{\rho + \frac{1}{\varepsilon_{u'(c)}} [A_S f'_S(k_{St}) + g'_2(a_t, k_{St}) - \rho - \delta]\right\} z_{St} + x_{St} - \omega L_S u'(c_{St}) \tag{6}$$

Optimality at the initial date further requires $z_{S0} = 0$.[8] In fact, the Northern planner's problem has a solution with $z_{St} \equiv 0$ at all dates. The Southern planner internalizes the capital externality, and the shadow value of Southern capital is the same for both planners. The level of aid is pinned down by the first-order condition (3). The solution is time-consistent, as the South values capital in the same way as the Northern planner.

We now consider what happens if the South is not run by a planner, and investment decisions are made by households. Then Eq. (1) replaces Eq. (2) as the last dynamic constraint in the Northern planner's problem. Again requiring $k_{St} = \overline{\varkappa}_{St}$ at each instant, the first-order condition for aid becomes:

$$x_{Nt} = g'_1(a_t, k_{St}) x_{St}, \tag{7}$$

The Euler equation for x_{St} is:

$$\dot{x}_{St} = \{\rho + \delta - A_S f'_S(k_{St}) - g'_2(a_t, k_{St})\} x_{St} + \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} A_S f''_S(k_{St}), \tag{8}$$

and the Euler equation for z_{St} finally becomes:

$$\dot{z}_{St} = \left\{\rho + \frac{1}{\varepsilon_{u'(c)}} [A_S f'_S(k_{St}) - \rho - \delta]\right\} z_{St} + x_{St} - \omega L_S u'(c_{St}) \tag{9}$$

Again optimality requires $z_{S0} = 0$, but no solution exists with $z_{St} \equiv 0$. The Northern planner's Euler Eq. (8) would then be incompatible with the South's Euler Eq. (1) because of the capital externality. (For example, one can see that $z_S > 0$ in the steady-state.) The externality leads to time inconsistency. The social returns to capital are higher than perceived by Southern households, given the benefits of higher output for aid absorption. To maintain the returns to capital and encourage capital accumulation, the Northern planner is less generous to the decentralized South than it would be to a social planner in the South. But at later dates, in the wake of past investment, the Northern planner would like to transfer additional aid in order to increase Southern consumption. Hence, the North's aid decision is time inconsistent.

The model based on decentralized investment decisions is arguably more realistic, and will be the focus of our simulations. When we later examine the contrasting case of a Southern planner, the optimal generosity of aid is similar, but with higher investment in the early stages of the transition. These effects are quantitatively modest, suggesting that time-consistent aid paths would not look greatly different from those under full commitment. As noted earlier, we assume that the donor can commit to an entire path for aid flows, and leave the investigation of time-consistent aid policies to future work, perhaps using the approach of Cohen and Michel (1988).

Contributions related to ours include Kopczuk et al. (2005), Weyl (2014) and especially Kemp et al. (1990). The latter paper uses optimal control methods to study the timing and generosity of aid needed to maximize weighted global welfare, just as we do. A crucial difference with our paper is that, although they briefly acknowledge the possibility of absorption constraints, they do not analyze them. Instead, they concentrate on the case where the South's resource constraint is linear in aid. They show that, at time zero, the optimal policy involves a transfer of part of the North's capital stock to the South.[9] In our work, we rule out stock transfers, a step which is a natural counterpart to the assumption of absorption constraints, as they note. A further difference is that Kemp et al. (1990) is purely theoretical, whereas we present quantitative results based on simulations.

2.3. The steady-state

We briefly discuss some properties of the steady-state, considering first the case of a social planner in the South. The first-order conditions and Euler equations for the dynamic problem, solved for the steady-state, lead to equations which implicitly define the steady-state level of aid, a^*, and consumption and capital for both North and South:

$$u'(c^*_N) = \omega \cdot L_S \cdot g'_1(a^*, k^*_S) u'(c^*_S)$$
$$f'(k^*_N) = \rho + \delta$$
$$A_S f'(k^*_S) = \rho + \delta - g'_2(a^*, k^*_S)$$
$$c^*_N = f(k^*_N) - a^* - \delta k^*_N$$
$$c^*_S = f(k^*_S) + g(a^*, k^*_S) - \delta k^*_S$$

The first equation has a natural interpretation: the Northern donor balances the marginal cost of aid – represented by the marginal utility

[6] Note that, for simplicity, we posit a Northern planner that chooses Northern consumption as well as aid. We could instead allow individual Northern households to choose their consumption paths (but not aid); that formulation would ultimately be equivalent in the current setting.

[7] We assume throughout that the necessary conditions for optimality will indeed characterize an optimal solution. It is not straightforward to establish that these conditions are sufficient in this particular case, given the free initial condition of one of the state variables, discussed below.

[8] This is the extra transversality condition required for optimality, given the free initial condition for one of the state variables: c_{S0} is not fixed because it partly depends on the North's choices. Note, however, that the feasible outcomes for c_{S0} remain constrained by the intertemporal resource constraints and the Euler equation for Southern consumption. See (Léonard and Long, 1992), chapter 7, especially section 7.7.

[9] For a textbook presentation of their results, see Brakman and van Marrewijk (1998).

of forgone Northern consumption – against the marginal benefit of aid, namely the marginal utility of Southern consumption weighted by ω and L_S, and multiplied by the derivative of the aid impact function with respect to aid.[10] Under our assumptions, the equation has some interesting implications when parameters are such that $c_N^* > c_S^*$. First, under absorption constraints, it is not optimal to equalize steady-state marginal utilities even when $\omega = 1$. Second, it can be shown that steady-state aid is increasing in ω, as expected. Third, another simple result applies under isoelastic preferences: steady-state aid is decreasing in the intertemporal elasticity of substitution. All of these results are in line with intuition.

For the alternative case where Southern consumption decisions are decentralized to households, the steady-state solution takes a more complicated form:

$$\begin{aligned} u'(c_N^*)\left[1+\frac{\rho\varepsilon_{u'(c_S)}g_2'(a^*,k_S^*)}{c_S^* A_S f''(k_S^*)}\right] &= \omega \cdot L_S \cdot g_1'(a^*,k_S^*)u'(c_S^*) \\ f'(k_N^*) &= \rho+\delta \\ A_S f'(k_S^*) &= \rho+\delta \\ c_N^* &= f(k_N^*) - a^* - \delta k_N^* \\ c_S^* &= f(k_S^*) + g(a^*,k_S^*) - \delta k_S^* \end{aligned}$$

Here the externality, the effect of capital accumulation in the South on aid absorption, leads to a wedge term in the equilibrium ratio of marginal utilities, represented by the term in square brackets. This term is greater than one, and keeps the marginal utilities further apart than in the Southern planning case. This is because a Southern planner would tend to choose a higher steady-state capital stock, allowing improved aid absorption and higher Southern consumption in steady-state. The effect becomes unimportant, and the model approaches the planning case, when the derivative $g_2'(a^*, k_S^*)$ is close to zero. Otherwise, the importance of the externality is increasing in ρ and the absolute value of $\varepsilon_{u'(c_S)}$ and decreasing in A_S and the absolute value of $f''(k_S^*)$.

In this version of the model, both k_S^* and k_N^* are determined by exogenous parameters and independent of aid flows. If we make explicit the dependence of c_N^* and c_S^* on aid, we have:

$$u'(c_N^*(a^*))\left[1+\frac{\rho\varepsilon_{u'(c_S)}g_2'(a^*,k_S^*)}{c_S^*(a^*)A_S f''(k_S^*)}\right] = \omega \cdot L_S \cdot g_1'(a^*,k_S^*)u'(c_S^*(a^*))$$

which is a single equation in a single unknown, the equilibrium aid flow a^*. The right-hand-side of the expression is decreasing in aid. Further, the left-hand-side is increasing in aid for the case we consider in our simulations, with isoelastic preferences, a specific $g(a, \overline{\varkappa}_S)$ function and strictly positive steady-state aid. These results imply that, when the Northern donor places a higher weight ω on Southern welfare, this must lead to higher aid in steady-state. It can also be shown that a larger relative population in the South (higher L_S) and higher total factor productivity in the North lead to higher steady-state aid.[11]

In other respects, the externality complicates the comparative statics. We saw earlier that, with isoelastic preferences, increasing the intertemporal elasticity of substitution reduces the extent of global redistribution, as expected. But this is not always true in the version of the model with an externality. In principle, there could be reasonable functional forms and parameter values for which steady-state aid is increasing in the intertemporal elasticity of substitution over some range. This counter-intuitive possibility arises because the externality-related wedge term depends on the elasticity. Since the effects of the externality are modest in our simulations, we do not pursue this question further.

The structure of the model does not rule out negative transfers. In fact, for a sufficiently low weight ω on Southern welfare, the North may choose to transfer income from the South to the North ('negative aid') at some points in time. This never happens in our simulations for the values of ω that we consider, but would emerge for weights close to zero. A richer analysis, at the expense of greater complexity, would either constrain aid to be non-negative — this is what Kemp et al. (1990) call the 'non-cooperative' case — or modify the aid impact function so that negative transfers are costly for the North to implement.

Note that, in our simulations, both the South and North will start below their respective balanced growth paths. We could assume that the North is on its balanced growth path and decides to donate a fixed share of its GDP, so that aid to the South grows steadily over time. This would remove the need to model the North explicitly, and Carter (2014) and Carter and Temple (2014) use this simplification. But if we want to study aid generosity, and allow the share of aid in Northern GDP to change over time, then the more general approach of this paper is useful. Aid generosity has implications for Northern consumption, and modeling the North explicitly is a natural way to capture the time-varying opportunity cost of aid.

3. Simulation assumptions

We now describe the assumptions used in our simulations. As we noted earlier, the simulations that we present are best seen as preliminary. This is partly because the model is stylized, and partly because the appropriate parameter choices are uncertain. The simulations illustrate what might be learnt from richer quantitative analyses in the future. In the past, the broad principles of aid programs, and some of their details, have often drawn on rather simple and unsatisfactory models, as Easterly (1999) emphasizes.

For the instantaneous utility function, we first adopt isoelastic (CRRA) preferences:

$$u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$$

where the elasticity of marginal utility with respect to consumption is $-\sigma$, and σ is the reciprocal of the intertemporal elasticity of substitution.

The parameter choices for preferences are relatively straightforward. We assume that σ is equal to two, a common choice in the literature.[12] We set the discount rate ρ equal to 0.03; this corresponds to the choices of Obstfeld (1999) and Gourinchas and Jeanne (2006). Note that both σ and ρ will influence the optimal generosity of aid.

For simplicity, we assume that donors and recipients have access to a Cobb–Douglas technology with a common exponent on capital, but different levels of TFP. We assume the exponent on capital is 0.50. This is higher than most estimates of physical capital's share of income, but some authors argue that a broader notion of capital is needed for neoclassical growth models to be consistent with the data: see Mankiw et al. (1992) and, especially, Barro and Sala-i-Martin (2004). Our choice of 0.50 has been used in related contexts, such as in Kraay and Raddatz (2007). The choice matters because it will influence the speed at which the South converges to its steady-state, and the rate at which the marginal product of capital falls along the transition.

To explore the predictions of the model, we also need to make assumptions about the long-run growth rates of GDP per capita and population, the size of the South relative to the North, and their initial capital–output ratios. Data on output, investment and population are taken from the Penn World Table version 6.3. We estimate the capital stock for 110 countries using the perpetual inventory method over the

[10] Although it may seem odd that the marginal utility in the South is weighted by L_S, note that typically the aid impact function will deflate aid by the population of the South, and hence the L_S term in the first equation will be offset by the implicit dependence of the derivative $g_1'(a^*, k_S^*)$ on $1/L_S$.

[11] We phrase this result in terms of Northern TFP, because results for Southern TFP are harder to derive.

[12] See Kraay and Raddatz (2007) for some references to empirical estimates. Obstfeld (1999) uses $\sigma = 2.5$.

Table 1
Calibration.

	N	S	S'	$S1$	$S2$
Population (L)	901	4844	2393	680	4164
Capital Stock (K)	103,600	56,030	25,900	1320	54,710
Output (Y)	31,410	28,050	13,330	1070	26,990
K/Y	3.30	2.00	1.94	1.23	2.03
α	0.5	0.5	0.5	0.5	0.5
k	10.88	3.99	3.78	1.52	4.11
k^*	11.891	11.891	11.891	11.891	11.891
k/k^*	0.915	0.336	0.317	0.128	0.346
ρ	0.03	0.03	0.03	0.03	0.03
σ	2	2	2	2	2
g	0.02	0.02	0.02	0.02	0.02
n	0.015	0.015	0.015	0.015	0.015
δ	0.06	0.06	0.06	0.06	0.06
A	1	0.274	0.271	0.121	0.325
$(Y/L)/(Y_N/L_N)$	1	0.17	0.16	0.05	0.19
L/L_N	1	5.376	2.656	0.755	4.622

period 1960–2007, following Bernanke and Gürkaynak (2002) and adopting their depreciation rate of 0.06.

We then aggregate countries into two units, the global North (the donor) and the global South (the recipient). The 33 countries aggregated into the Northern economy are those with output per capita above 20,000 in 2007 international (PWT) dollars, while the remainder are classed as the South. For our initial investigations, we sometimes exclude China and India, with their large populations.[13] This keeps the quantitative analysis comparable with the aid decisions made in practice. As is well known, aid receipts per capita are low for China and India, partly reflecting the 'small country bias' in aid allocation; excluding these two countries helps to keep the model close to the data. This leaves us with an aid recipient whose population size is 2.7 times that of the donor, denoted S' in Table 1. The final two columns of Table 1, S1 and S2, show the South sub-divided into low and middle income recipients, for use in Section 6.

The cut-off of 20,000 international dollars roughly corresponds to the upper quartile of the GDP per capita distribution. We then calculate total output, capital stock and population for the two regions, for the most recent year in the PWT 6.3 data (2007), making no distinction between population and labor force. This procedure yields a capital–output ratio for each region. Since we assume Cobb–Douglas production functions for both North and South, the capital–output ratios imply the initial levels of capital per effective worker. We can then infer the relative TFP and GDP per capita of the South.

We assume that rates of technical progress and population growth are the same for donor and recipient, helping to ensure a balanced growth path. The first assumption is common in the empirical growth literature. We adopt a rate of technical progress of 2% a year, as in Mankiw et al. (1992). This is also approximately the average growth rate in our Northern group of countries for the most recent decade in the data. The assumption that population grows at the same rate in donor and recipient can be justified as a long-run outcome, given that population growth rates are falling in the developing world. We assume that the long-run population growth rate is 1.5% a year. This is approximately the average rate over the last decade in the Southern group of countries, but somewhat higher than in the North over the same period. Under these assumptions, the North begins the simulation with capital per effective worker about 10% below its steady-state value (see the 'N' column of Table 1).

One of the main uncertainties relates to the severity of absorption constraints. As noted earlier, there are several reasons why the marginal benefit of aid might be declining in aid intensity, the ratio of aid to recipient GDP. We capture this idea by multiplying aid by a simple 'wastage' function: a proportion of aid is wasted, and this proportion is increasing in aid intensity. This drives a (potentially time-varying) wedge between the transfer made by the North, and the aid received by households in the South. Our chosen functional form, mapping aid a_t into transfers to households, is:

$$g(a_t, \overline{\varkappa}_{St}) = \frac{a_t}{L_{St}}\left(1 - \frac{a_t}{\upsilon Y_S(\overline{\varkappa}_{St})}\right) \tag{10}$$

where $Y_S(\overline{\varkappa}_{St})$ is output in the South, and higher values of υ correspond to greater aid effectiveness for any given capital–labor ratio. The growth rate of capital is then a quadratic function of aid intensity, which is parallel to the empirical literature, where growth is sometimes a quadratic function of aid intensity; see Clemens et al. (2012) for a recent example.

A convenient way of interpreting this function is to ask when aid intensity a_t/Y_{St} is sufficiently high that the marginal benefit of aid is zero for the South. This happens when $a_t/Y_{St} = \upsilon/2$. The appropriate severity of this absorption constraint is a matter of debate. In our baseline, we set υ so that an extra dollar of aid has zero impact when the aid/GDP ratio is 25%.[14]

A final choice relates to ω, the relative weight of Southern utility in the objective function of the North. The use of such a weight seems essential to any normative study of transfers between the North and South. We are not seeking to estimate the weight that Northern citizens currently adopt, nor to defend a particular choice of ω on prior grounds. Instead, our framework provides a mapping between assumptions on ω and optimal aid generosity. A possible parallel would be with Kopczuk et al. (2005), who build on the earlier work of Mirrlees (1971): these papers allow a mapping between alternative degrees of inequality aversion and tax rate schedules. In our simulations, we choose a baseline weight $\omega = 0.1$ to represent a donor that is genuinely altruistic, but imperfectly so. We will investigate later how optimal aid changes when ω takes higher values. As we will see, even the choice of $\omega = 0.1$ — so that Northern citizens have ten times the welfare weight of Southern citizens — can lead to shares of aid in Northern GDP substantially higher than in the data.

3.1. Welfare analysis

To compare the lifetime utilities associated with different consumption paths, we use Hicksian Equivalent Variation (HEV) as in Gourinchas and Jeanne (2006). Hence our welfare results, in moving from one lifetime utility to another, are expressed in terms of the constant proportional change in consumption at each instant that would generate the new level of utility. Under isoelastic preferences, this is given by:

$$\lambda = \left(\frac{U_{aid}}{U_{aid=0}}\right)^{\frac{1}{1-\sigma}} - 1$$

if $\sigma \neq 1$, and $\lambda = e^{\rho(U_{aid} - U_{aid=0})} - 1$ otherwise. Depending on the simulation experiment, we sometimes calculate this measure for the South in isolation, and sometimes for the Northern donor's objective function, including the weight ω.

4. Optimal aid policies

In this section, we study the optimal time path for aid using simulations. To do this, we use the relaxation algorithm of Trimborn et al. (2008) to solve the system of differential equations implied by the optimal control solution. In our baseline case, aid generosity should be highest at the beginning, so that the South accelerates towards its steady-state. But this front-loading result is sensitive to absorption

[13] From the set of possible recipients, we also drop a handful of Eastern European countries and some smaller countries for which data are unavailable.

[14] Our chosen functional form implies that, beyond this level of aid intensity, the marginal benefit of aid is negative. This would not be plausible for some interpretations of the aid impact function. Under a pure 'wastage' interpretation, a functional form in which the marginal benefit declines to zero would be preferable; see Carter (2014).

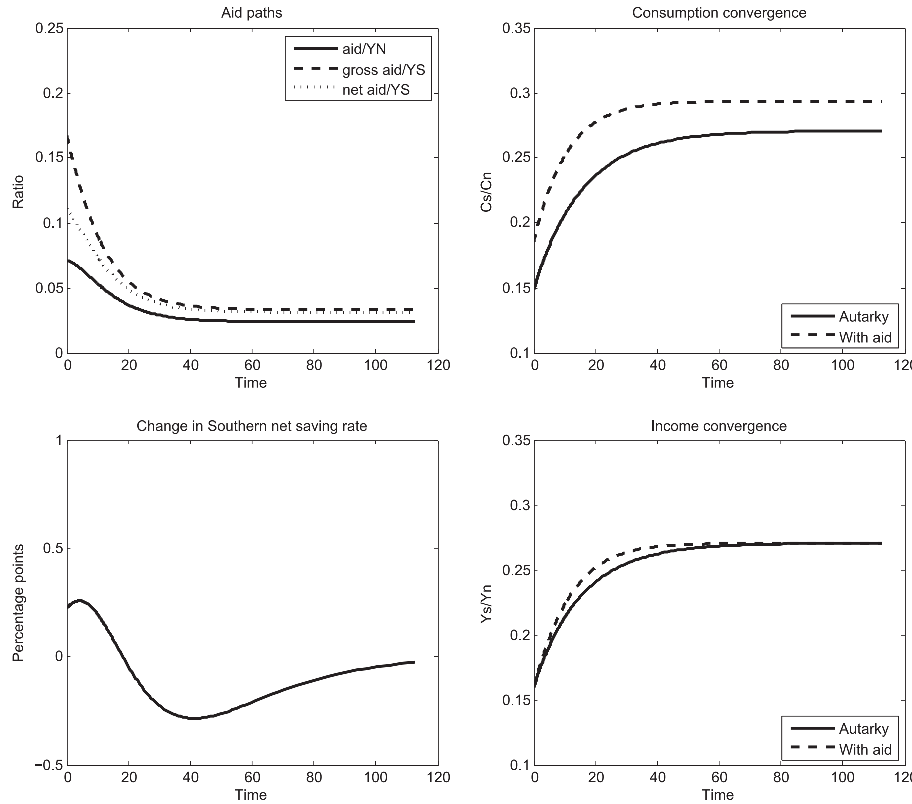

Fig. 1. Baseline case; decentralized version. The panel 'Aid paths' shows the optimal path of aid over time, relative to Northern and Southern GDP. The gap between gross and net aid in the South represents wastage. The two right-hand panels show consumption and output in the South relative to the North, an upward slope signifying convergence. In both panels convergence under optimal aid (dashed line) is compared to a zero-aid counterfactual (autarky, solid line). The bottom-left panel shows the difference aid makes to the Southern net saving rate compared to autarky.

constraints. There are scenarios in which the North should increase the generosity of its aid (relative to its GDP) as the South grows, since the South can then absorb a given level of aid more effectively.

The first panel of Fig. 1 shows the optimal path of aid in our baseline case. Aid relative to Northern GDP is initially 7.1%, falling to 4% after 17 years and to 2.4% in steady-state. Aid flows on this scale are clearly too high to be realistic, and were not seen even under the Marshall Plan.[15] Importantly, however, long-run generosity is highly sensitive to the assumed curvature of the utility function. If we reduce σ by just 10%, to $\sigma = 1.8$, the North will continue to be generous early on, but much less so at longer horizons. Hence, diminishing returns to consumption can provide a powerful motivation for international transfers, but at longer horizons, parameter assumptions matter a great deal.

Moreover, it is likely that the North places even less weight on the utility of the South than we are currently assuming. There are other possible reasons for the divergence between the model and the flows observed in practice. We do not model the marginal cost of public funds: aid is financed by lump-sum taxation rather than distortionary taxes. Nor do we incorporate any political economy constraints on the donor, such as taxpayer resistance to large international transfers. Hence, natural extensions to our model would yield smaller ratios of aid to Northern GDP for given values of ω and σ.

Aid intensity in the South is initially 16.8%, falling to 3.3% in the long run. Fig. 1 also shows the effect of aid on the rate at which Southern output and consumption converge towards Northern levels, and the effect of aid on the Southern saving rate (expressed as the deviation from a zero-aid counterfactual). The front-loading of aid enables the South to raise consumption by 14.4% immediately.

The effect of aid on growth is familiar from Obstfeld (1999): there is an initial, but modest, acceleration which is eventually followed by slower growth relative to the autarky counterfactual. Growth is ultimately slower because aid brings capital accumulation forward, so that capital is accumulated quickly early on and more slowly later. Even though aid intensity is high in this baseline calibration, the effect on output growth is small. In annual terms, the initial growth rate is 0.57 percentage points higher, with the difference eliminated after 14 years; roughly 25 years later, the growth rate is 0.08 percentage points lower than it would have been without aid.

It may seem surprising that the growth and welfare effects of aid are not larger. It is clear that, with isoelastic utility, exogenous transfers of this magnitude have only limited effects on optimal investment. The explanation is that, when the benefits of investment are high, it will be undertaken even in the absence of transfers, because forgoing consumption is not costly under these preferences. Hence Obstfeld (1999) finds modest growth effects, and he notes (p. 136) that the result is 'likely to be a robust feature of any plausible model in which aid is funneled through the private sector'.[16]

More broadly, our results on welfare effects are also consistent with the work of Gourinchas and Jeanne (2006) on capital mobility. They showed that the welfare benefits of capital inflows in a calibrated Ramsey model are unexpectedly modest. This is partly because convergence will be rapid even in the absence of foreign capital, and partly because accumulating capital more rapidly brings forward a reduction in its marginal product. This intuition is helpful in understanding why the effects of aid on productivity are also relatively modest in our setting. But in the case of aid, we can take this logic further: if relaxing

[15] For comparison, the Marshall Plan entailed aid generosity of just over one per cent of US GDP on average, over the years it operated (Crafts, 2013).

[16] Larger effects on investment and growth can obtain under Stone–Geary preferences, when relaxing the intertemporal resource constraint becomes more valuable. We examine this case later in the paper.

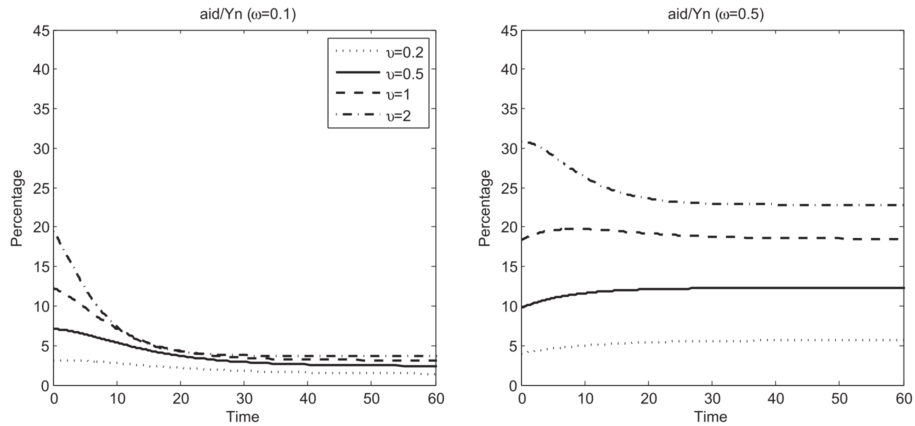

Fig. 2. Decentralized a/Yn, varying absorptive capacity (υ). Optimal paths of aid relative to Northern GDP, as Southern absorptive capacity varies, for two degrees of Northern altruism. Smaller values of υ correspond to a tighter absorption constraint, with higher levels of waste at low levels of aid intensity. In the left-hand panel the Northern planner places a weight of $\omega = 0.1$ on Southern utility, $\omega = 0.5$ in the right-hand panel.

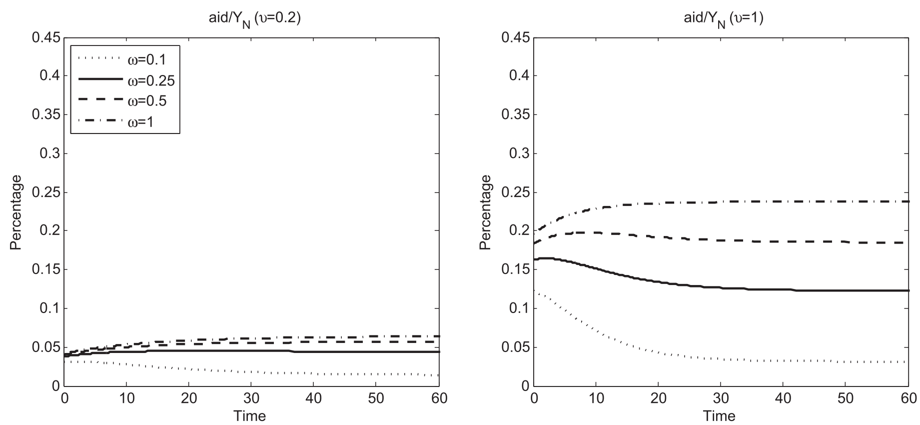

Fig. 3. Decentralized a/Yn, varying Northern altruism ω. Optimal paths of aid relative to Northern GDP, with Northern altruism increasing in ω, for two degrees of Southern absorptive capacity. In the left-hand panel Southern absorptive capacity is low, $\upsilon = 0.2$, and higher in the right-hand panel $\upsilon = 1$.

the South's resource constraint is not all that valuable to the South, it is likely that tightening the North's resource constraint is not all that costly to the North. This is why we find large transfers to be optimal, even though the productivity benefits of a given transfer are limited.

Scenarios in which aid increases over time are shown in Figs. 2 and 3. For some combinations of ω (the weight on the South's utility) and υ (the aid absorption parameter) it may be optimal for the North to back-load aid: as the South develops, it becomes better able to use aid effectively, and the North should increase the generosity of aid relative to its GDP.

What happens if the North gives a fixed share of its GDP, as in the Pearson Commission 0.70% benchmark? When the fixed share is chosen optimally, the welfare losses for the South relative to the fully-flexible optimal aid policy are modest. Perhaps counter-intuitively, if the North has to donate a fixed share of GDP, the costs of this restriction are sometimes borne by the North in equilibrium: the North becomes more generous in the long run than it otherwise would, and the South sometimes does better than it otherwise would.[17] We also find that, if aid is fixed as a share of Northern GDP, departures from the optimal choice of this share are not particularly costly: the donor's objective function is relatively flat as a function of the aid share.

Fig. 4 compares outcomes when the North can vary aid as a share of its GDP, and when the share is fixed. By construction, the donor must always do at least as well when free to choose a flexible path for aid, but the effect of fixing the share on the donor's objective is surprisingly modest. In our baseline case, the associated utility loss is equivalent to an ongoing reduction in consumption of 0.2% for the North and South. Looking at the two economies individually, the restriction to a fixed share leaves Northern households worse off (− 0.8 %) and Southern households better off (+ 0.6 %). But this is not always the case. For some cases with higher ω, the restriction to a fixed share leaves the North better off and the South worse off. The role of ω hints that these results arise from the interplay of the benefits of aid at short horizons, the benefits at long horizons, and whether fixing the aid share is wasteful due to absorption constraints. Finally, we limit aid to the Pearson Commission benchmark, 0.70 % of Northern GDP. Compared to our optimal flexible path, the much lower generosity means that the HEV welfare gain in the South falls from 11% to 2%.

5. Optimal aid for subsistence economies

We now consider preferences with a subsistence level of consumption. It is well known that the single-sector Ramsey model with isoelastic preferences makes some unrealistic predictions: fast convergence to the steady-state, and a sharp decline in investment rates and in the marginal product of capital as capital is accumulated. Nor can it readily accommodate periods of international divergence in GDP per

[17] This result could be eliminated by alternative policy rules; for example, the share of aid in Northern GDP could be assumed to be a decreasing function of the GDP per capita of the South.

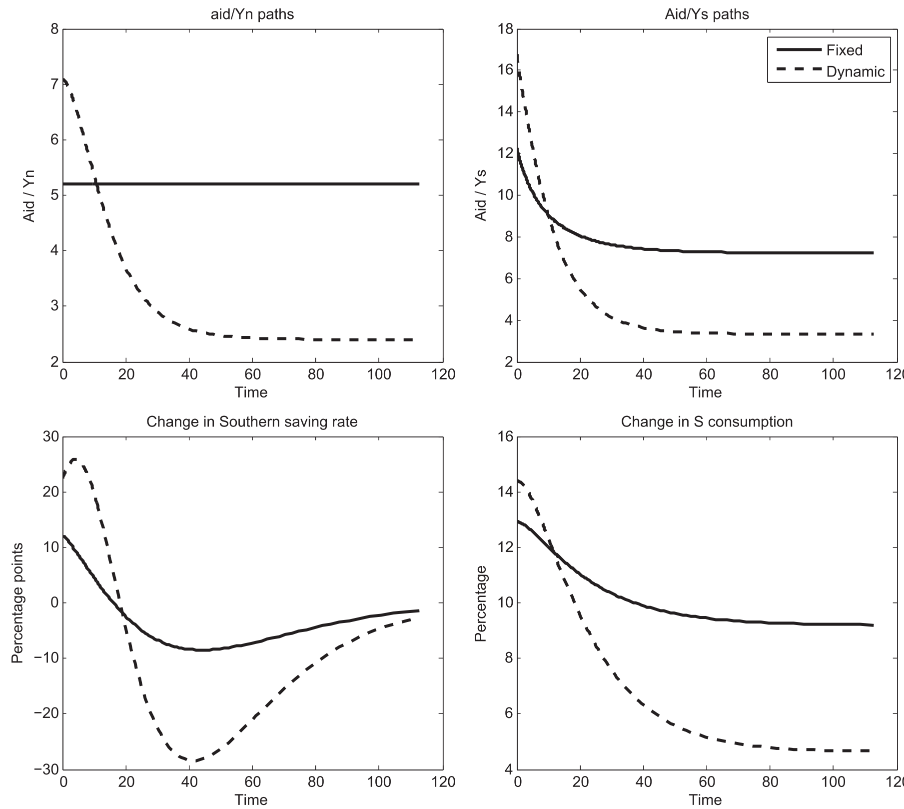

Fig. 4. Fixed v. dynamic; baseline case. Optimal aid in our baseline parameterization, comparing the unconstrained dynamic solution (dashed line) to the case in which the donor gives a fixed share of its GDP (solid line). The upper two panels show the optimal paths of aid, relative to Northern and Southern GDP. The lower-left panel shows the change in the Southern net saving rate (the rate with aid minus the rate without) in both cases. The lower-right panel shows the percentage gain in Southern consumption achieved by aid relative to autarky in both cases.

capita, or explain why investment rates are sometimes low in poorer countries. Ben-David (1998) and Steger (2009) argue that Stone–Geary (SG) preferences overcome these problems. Low investment can co-exist with high returns to investment, because the opportunity cost of investment is high when households are close to subsistence. We examine the implications for the generosity and timing of aid, including whether donors will want to front-load aid.

Formally, the framework is identical to that used earlier, but with instantaneous utility now given by $\left((c_{jt}-\bar{c})^{1-\sigma}-1\right)/(1-\sigma)$ for $j = N, S$ where $\bar{c}$ corresponds to a subsistence level of consumption. With labor-augmenting technical progress, the fixed subsistence level $\bar{c}$ influences the transitional dynamics, but the model will have an asymptotic balanced growth path; see, for example, Ohanian et al. (2008). The qualitative nature of the dynamic path will depend on initial conditions. Given technical progress, there will ultimately be convergence for a wide range of initial capital stocks, but the growth and convergence process will often be slow. In particular, if the South begins close to subsistence, the growth rates of capital per worker and GDP per worker will be slower than the rate of technical progress for a transitional period, and living standards will temporarily diverge.

In our experiments, we set the subsistence level so that ten-year growth in output per capita in a simulation of the South (without aid) is close to that observed over 1997–2007, assuming labor-augmenting technical progress of 2 % a year. Table 2 lists the growth rates obtained from simulations, for alternative choices of $\bar{c}$. For our baseline version of the South, the annual growth rate over 1997–2007 was 3.71 % and hence we set $\bar{c} = 0.25$ for our baseline case. This will turn out to be roughly 50 % of the initial consumption level for the South in autarky.

As expected, under SG preferences, aid generosity should be high early in the development process. We present some results in Fig. 5, which could be compared with the earlier CRRA results in Fig. 1. The optimal level of aid is initially higher under SG preferences, at 9.6% of Northern GDP. It declines to the same asymptotic level of 2.4% (recall that the CRRA and SG growth paths are asymptotically equivalent) but converges to this level more slowly in the SG case. Under SG preferences, Southern output converges more slowly towards the Northern level. This is a natural result: when the Southern economy is close to subsistence, the opportunity cost of investment is high. But the optimal time path for aid helps to close the gap between North and South to a greater extent under SG preferences than under CRRA preferences. This is another natural result, since relaxing the intertemporal resource constraint is more valuable in the SG case.

The effect of SG preferences on the generosity of aid means that aid has a greater impact on consumption in the SG case. The peak change in consumption, relative to autarky, is similar (17% under SG, 14% under CRRA) but the duration of the effect on consumption differs greatly. Under SG preferences, Southern consumption is 11% higher than the zero-aid counterfactual 50 years after the commencement of aid. Under CRRA preferences, the equivalent figure is just 6%. This difference reflects the greater generosity of aid in the SG case: aid is 5.5% of Northern GDP at $t = 50$ under SG preferences, but just 2.5% under CRRA. Although the absolute change in the initial annual growth rate, relative to the zero-aid counterfactual, is similar to the CRRA case at 0.59 percentage points, autarky output growth is slower under SG preferences,

Table 2
Output growth in South under alternative subsistence thresholds.

$\bar{c}_S$	0	0.1	0.25	0.375
Growth	0.0527	0.0473	0.0362	0.0212

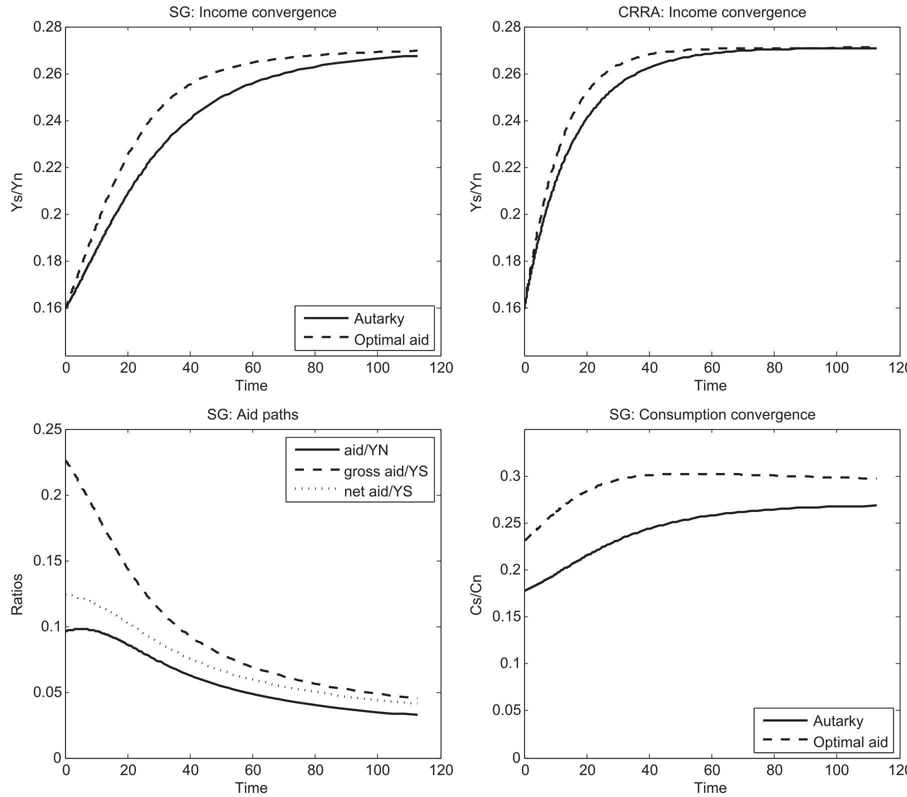

Fig. 5. Baseline case: subsistence version. Optimal aid under Stone–Geary preferences. For ease of comparison, the income convergence panel from Fig. 1, under CRRA preferences, is shown again here.

at 3.5% per year.[18] So the impact of aid on growth with SG preferences is more substantial, multiplying the initial growth rate by 1.17, compared with a multiple of 1.08 under CRRA preferences.

To compare the SG and CRRA cases in more detail, Fig. 6 shows outcomes based on the same fixed quantity of aid.[19] The higher aid intensity in the SG case reflects the fact that Southern output grows more slowly under these preferences. The effects of aid on net investment, consumption and growth are markedly greater in the SG case. Since our assumptions on preferences vary across the two cases, more direct welfare comparisons are not meaningful.

6. Optimal aid with two recipients

In this section, we study the donor's problem when there are two aid recipients. Additional recipients can be integrated into the Northern planner's optimal control problem in the obvious way. With multiple recipients, allocation decisions are connected because aid to one country reduces consumption in the North, increasing the opportunity cost of aid to other countries.

We construct a low-income recipient, $S1$, using the countries in the lowest quarter of the GDP per capita distribution. The other recipient, $S2$, is an aggregate of 'middle income' countries, defined as those between the lower and upper quartile of the GDP per capita distribution. Table 1 lists some of the relevant numbers. We now include China and India, unlike in our baseline case. As a result, the middle-income country is 'large': the population of $S2$ is 6.1 times greater than the population of $S1$.

Output per capita in $S2$ is about 19% that of the North, whereas the poor recipient, $S1$, has output per capita only 5% that of the North. Hence, in this first experiment, output per capita in $S2$ is about four times greater than in $S1$. The long-run ratio of output per capita between the two is smaller, however, at 2.5. This is because the initial capital stock in $S1$ is further beneath its long-run level than $S2$. The ratio k_0/k^* is 0.13 for $S1$ and 0.35 for $S2$. Hence, without aid, $S1$ initially grows faster than $S2$: annualized growth is 13% for $t = 0$ in $S1$, 6% in $S2$.

This calibration can address an important policy issue: does low consumption help to motivate aid even for middle-income countries? Fig. 7 illustrates the optimal aid allocation under CRRA preferences. The contrast between the two recipients is striking. A substantial amount of aid is allocated to the middle-income recipient, and its response is familiar from our baseline case. But despite the poor recipient $S1$ starting further beneath its balanced growth path, and with greater scope for aid-induced growth, the effect on investment (which mirrors the effect on output convergence) is negligible. Instead of a front-loaded aid path that induces accelerated investment, the donor keeps aid intensity high throughout and increases generosity in absolute terms over time. This reflects the aid absorption constraint: the donor cannot be too generous to $S1$ initially, because the optimal quantity of aid already takes this economy close to the point where the marginal benefit of aid is zero. The recipient $S1$ uses aid to fund a higher level of consumption, with little change in investment behavior, because this recipient correctly anticipates that aid will increase over time. As for the middle-income recipient $S2$, where the population is larger and the absorption constraint is less binding, the donor is generous early on but then scales aid back over time as the recipient grows. This is reflected in the shares of the two recipients in total aid: the middle-income recipient receives the bulk of the aid at shorter horizons, and the low-income recipient $S1$ at longer horizons.

To explore further, we change the population of $S2$ so that it matches that of $S1$. Fig. 8 shows what happens: there is a sharp reduction in the

[18] With Stone–Geary preferences and initial conditions close to subsistence consumption, growth accelerates initially in the autarky case. In our baseline SG case, Southern growth peaks after four years.

[19] The fixed level chosen is that which is optimal in the CRRA case. The third panel shows the percentage point change in the growth rate.

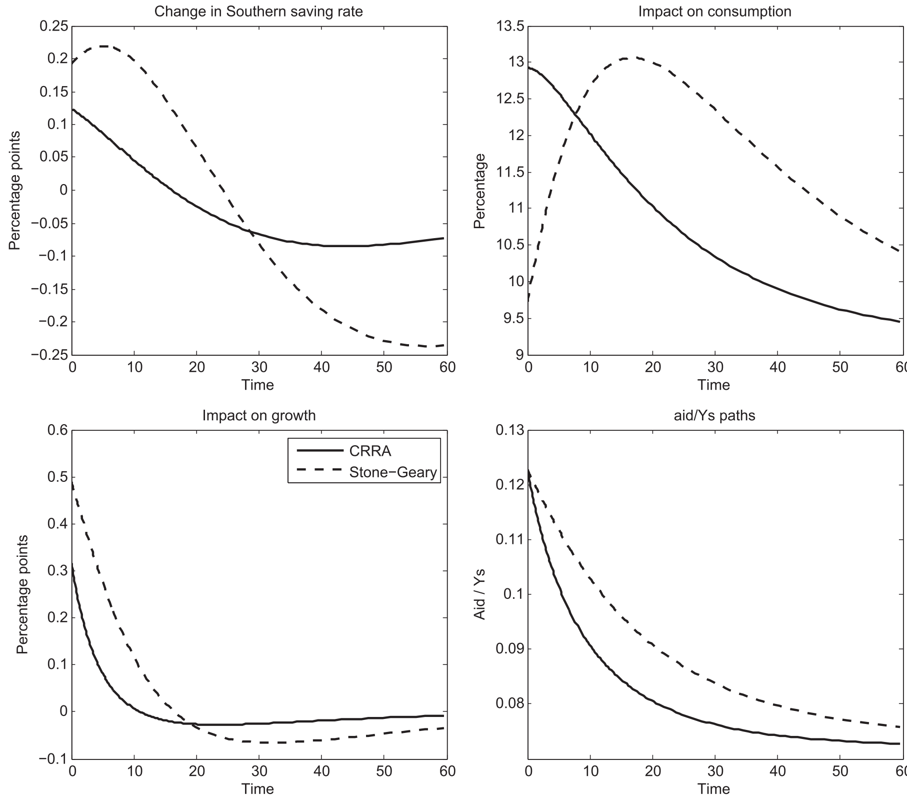

Fig. 6. Fixed aid: CRRA versus SG. This figure compares outcomes under Stone–Geary (dashed line) and CRRA (solid line) preferences, for the same fixed flow of aid, equal to 5% of Northern GDP. The lower-left panel shows the difference made to the growth rate by aid compared to autarky. Descriptions of other panels as before.

quantity of aid given to *S*2, but qualitatively the paths are unchanged. The level of aid given to *S*1 is virtually unchanged, despite the large reduction in aid given to *S*2 with an associated reduction in the opportunity cost of aid to the North. This result arises because *S*1 is already being given as much aid as it can effectively absorb.[20]

7. Sensitivity analysis

In this section, we carry out some sensitivity analyses, and then discuss possible extensions. First, recall the externality that arises in our setup: capital accumulation by Southern households increases the ability of the South to absorb aid. What happens if we take our baseline calibration and assume Southern decisions are made by a social planner? The optimal policy for the Northern planner now leaves Northern households slightly worse off (equilibrium aid is slightly more generous) and the South slightly better off. The effects on the North are modest and the optimal path for aid barely differs, compared to the decentralized version. The peak difference in aid between the two cases, as a percentage of Northern GDP, is only 0.7 percentage points. But the change in Southern consumption is noticeable: the Southern planner anticipates the favorable effect of investment on future aid absorption, and chooses higher investment rates in the first few years of the transitional dynamics. Overall, however, the quantitative importance of the externality is modest. We discuss these results in more detail in an online appendix.

Next, we briefly consider the effects of varying two parameters: σ, the curvature of the utility function, and α, the output–capital elasticity. In modifying these parameters, we still require the Northern and Southern capital–output ratios to match those in the data. Relative to our previous experiments, there is a direct effect of changing each parameter, and an indirect effect which arises from matching the observed capital–output ratios under new parameter assumptions. The alternative route, of only focusing on the direct effect, would lead to an inconsistency between the capital–output ratios in the simulations and those in the data.

As noted earlier, the curvature of the utility function has a major effect on optimal generosity, especially at longer horizons. When the marginal utility of consumption diminishes less rapidly with consumption, the motivation for global redistribution is weakened. In our simulations, with slightly less curvature ($\sigma = 1.8$ rather than $\sigma = 2.0$) aid generosity is reduced and the effect on consumption is smaller. Given technical progress, the level of the balanced growth path is raised by a reduction in σ, and the effect of aid on growth is initially stronger than in our baseline. This effect is modest, however.

Increasing the output–capital elasticity has similarly complicated effects. It makes aid more effective, since capital accumulation becomes more important; it also raises the balanced growth path. We do not present the results in detail, but note that as α increases, the effect of aid on investment is lower but the effect on growth larger. Intuitively, because investment now has greater benefits, the South is able to consume more early on.

We now briefly discuss some possible extensions. Perhaps the most important modification to the analysis would be to introduce political economy considerations within the South. This would introduce a major complication, namely strategic interactions between the donor and recipients. One possibility would be to allow aid to finance public investment, as in Chatterjee et al. (2003) and other work summarized in Turnovsky (2009); see also Lowe et al. (2013). Expenditure on public investment projects could be wasteful to some degree, with the extent of waste increasing in expenditure as in Berg et al. (2013). Politics in the North could also play a role: in principle, the North's objective function could reflect the North seeking political influence, as in the work of Antràs and Padró i Miquel (2011).

[20] Initial aid intensity is close to 25% of the GDP of recipient *S*1, the point at which the marginal benefit of aid to *S*1 is zero.

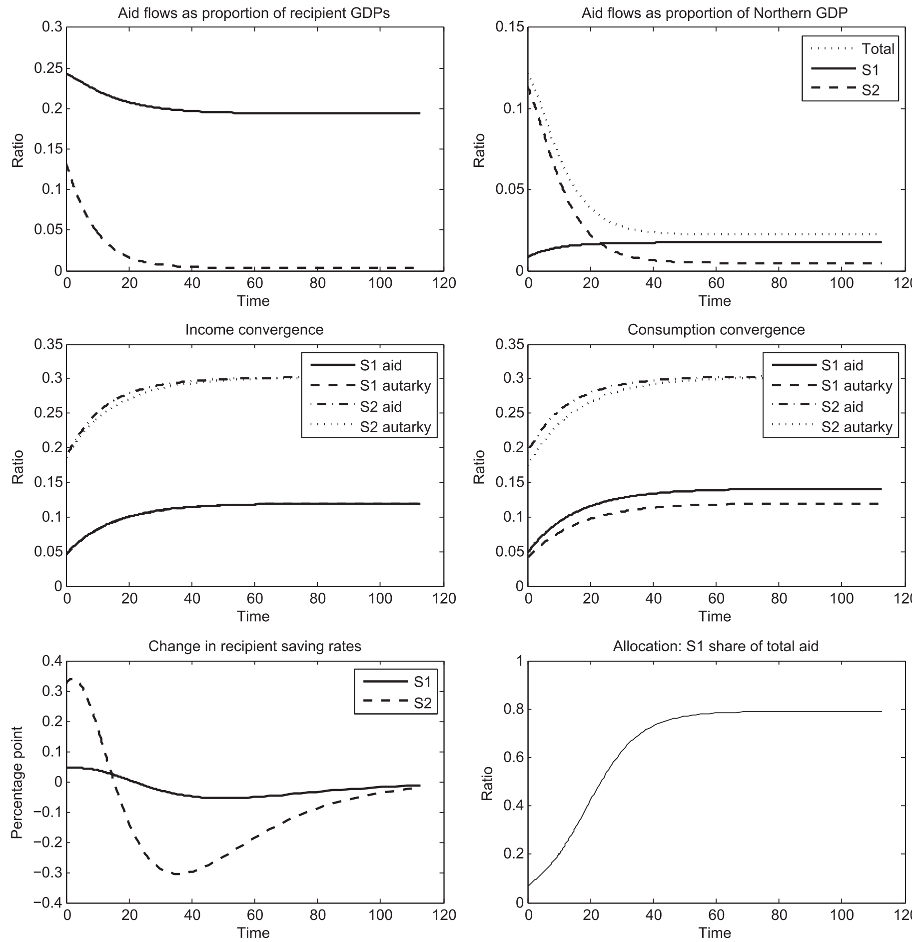

Fig. 7. Two recipients. A two-recipient case: *S*1 small and poor, *S*2 middle-income and populous. The middle panels show the ratios of income and consumption in the two recipients to Northern levels, under optimal aid and in autarky; for income in the case of S2, the effect is too small to be seen.

Another extension could allow aid recipients to borrow internationally, with a risk premium that is increasing in external indebtedness. van der Ploeg and Venables (2013) examine the optimal response of recipient governments to windfall revenues in such an environment, but treat the time path of revenues as exogenous, which seems better suited for natural resource windfalls than aid flows. Perhaps their decision problem could be integrated within that of an altruistic North, using the nested structure we adopt here. A related extension would model recipients as two-sector economies that produce traded and non-traded goods, in which case absorption constraints could arise endogenously through Dutch Disease effects.

Staying closer to our current framework, other possibilities include country-specific population and technical progress dynamics, costs of distortionary taxation in the donor, adjustment costs in the North for aid disbursements, finite-duration aid commitments by the donor, CES production technologies, and the introduction of capital varieties; see Hoxha et al. (2013) on the latter. A more ambitious extension would introduce output volatility in the South, so that aid would have an additional insurance role. This would bring the analysis closer to Arellano et al. (2009), at the expense of greater computational complexity.

8. Conclusions

This paper has introduced a framework for analyzing optimal aid policies using the neoclassical growth model. In contrast to most previous research, the findings are based on a clearly-defined optimization problem for the Northern donor. It takes into account the opportunity cost of aid, optimizing behavior by Southern households, and constraints on the effective absorption of aid. This framing of the problem is a first step in the direction of a richer quantitative analysis, and the simulations indicate what might be learnt from future exercises of this type. Since any tractable model must be stylized, there will always be a need for subjective judgments on the part of donors; but formal models can make visible some considerations and possibilities that would otherwise be obscure.

We find that optimal transfers are influenced by the weight a donor places on recipient welfare, the curvature of the utility function, the recipient's capacity to absorb aid, the relative level of the recipient's balanced growth path, and the recipient's initial distance from that growth path. In our simulations, the scope for aid to raise growth rates plays some part in aid decisions, but the effect on the level of consumption often dominates. The optimal generosity of aid, relative to Northern GDP, should vary over time to reflect conditions in aid recipients. In our baseline case, aid generosity declines over time. But in some scenarios, where donors run up against absorption constraints, there is a time interval over which generosity should be increasing. As for optimal aid intensity — that is, aid relative to Southern GDP — this should generally decline over time. But in cases where the aid absorption constraint is close to binding, the optimal policy may dictate that high aid intensity is maintained for a long time. In the case of multiple recipients, the optimal policy may involve large changes over time in the absolute quantity of aid and the division of aid between recipients.

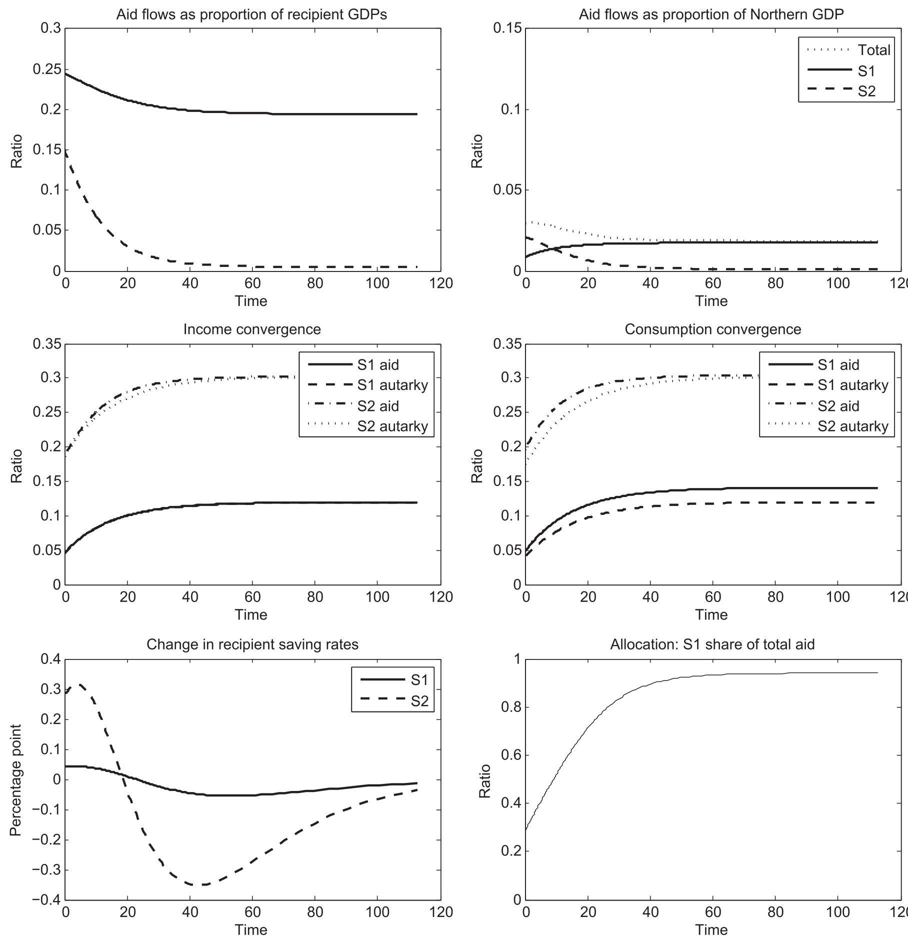

Fig. 8. Two recipients: *S*2 population equal to *S*1. A two-recipient experiment, with the population of *S*2 now reduced to match *S*1.

As expected, the case for donor generosity is strengthened when consumption is close to subsistence. The effects of aid on Southern consumption, investment and growth can be dramatic in this case. In general, our numerical results indicate optimal aid levels that exceed those in the data, at least at short horizons. An interesting task for future research is to pin down the features of reality that could explain this disparity. That citizens of donor economies may place little weight upon the welfare of aid recipients is only one possible explanation; others could include the potential roles of political economy forces and corruption in undermining the effectiveness of aid. Given the obvious importance of these considerations, our approach is a partial view. Nevertheless, it provides some insights and results that could help to inform the design of aid policies, and is simple enough to be extended in many directions.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.jinteco.2014.11.005.

References

Antràs, P., Padró i Miquel, G., 2011. Foreign influence and welfare. J. Int. Econ. 84, 135–148.

Arellano, C., Bulir, A., Lane, T., Lipschitz, L., 2009. The dynamic implications of foreign aid and its variability. J. Dev. Econ. 88, 87–102.

Barro, R.J., Sala-i-Martin, X., 2004. Economic Growth, second edition. The MIT Press, Cambridge, Massachusetts.

Ben-David, D., 1998. Convergence clubs and subsistence economies. J. Dev. Econ. 55, 155–171.

Berg, A., Portillo, R., Yang, S.C.S., Zanna, L.F., 2013. Public investment in resource-abundant developing countries. IMF Econ. Rev. 61, 92–129.

Bernanke, B.S., Gürkaynak, R.S., 2002. Is growth exogenous? Taking Mankiw, Romer, and Weil seriously. NBER Macroeconomics Annual 2001 vol. 16. NBER, pp. 11–72.

Brakman, S., van Marrewijk, C., 1998. The Economics of International Transfers. Cambridge University Press, Cambridge, UK.

Burnside, C., Dollar, D., 2000. Aid, policies, and growth. Am. Econ. Rev. 90, 847–868.

Carter, P., 2014. Aid allocation rules. Eur. Econ. Rev. 71, 132–151.

Carter, P., Temple, J.R.W., 2014. Virtuous circles and the case for aid. Discussion paper 14/636. University of Bristol.

Caselli, F., Feyrer, J., 2007. The marginal product of capital. Q. J. Econ. 122, 535–568.

Chamon, M., Kremer, M., 2009. Economic transformation, population growth and the long-run world income distribution. J. Int. Econ. 79, 20–30.

Chatterjee, S., Sakoulis, G., Turnovsky, S.J., 2003. Unilateral capital transfers, public investment, and economic growth. Eur. Econ. Rev. 47, 1077–1103.

Clemens, M., Moss, T., 2007. The ghost of 0.7 per cent: origins and relevance of the international aid target. Int. J. Dev. Issues 6, 3–25.

Clemens, M.A., Radelet, S., Bhavnani, R.R., Bazzi, S., 2012. Counting chickens when they hatch: timing and the effects of aid on growth. Econ. J. 122, 590–617.

Cohen, D., Michel, P., 1988. How should control theory be used to calculate a time-consistent government policy? Rev. Econ. Stud. 55, 263–274.

Crafts, N., 2013. The Marshall Plan. In: Parker, R.E., Whaples, R.M. (Eds.), Routledge Handbook of Major Events in Economic History, Routledge, pp. 203–213.

Dudley, L., Montmarquette, C., 1976. A model of the supply of bilateral foreign aid. Am. Econ. Rev. 66, 132–142.

Easterly, W., 1999. The ghost of financing gap: testing the growth model used in the international financial institutions. J. Dev. Econ. 60, 423–438.

Eaton, J., 1989. Foreign public capital flows. In: Chenery, H., Srinivasan, T. (Eds.), Handbook of Development Economics vol. 2. Elsevier, pp. 1305–1386.

Galor, O., Polemarchakis, H.M., 1987. Intertemporal equilibrium and the transfer paradox. Rev. Econ. Stud. 54, 147–156.

Gertler, P.J., Martinez, S.W., Rubio-Codina, M., 2012. Investing cash transfers to raise long-term living standards. Am. Econ. J. Appl. Econ. 4, 164–192.

Gourinchas, P.O., Jeanne, O., 2006. The elusive gains from international financial integration. Rev. Econ. Stud. 73, 715–741.

Hanlon, J., Barrientos, A., Hulme, D. (Eds.), 2010. Just Give Money to the Poor. Kumarian Press, Sterling, VA.

Haushofer, J., Shapiro, J., 2013. Household response to income changes: evidence from an unconditional cash transfer program in Kenya. MIT, (Manuscript).
Hoxha, I., Kalemli-Ozcan, S., Vollrath, D., 2013. How big are the gains from international financial integration? J. Dev. Econ. 103, 90–98.
Kanbur, R., Sumner, A., 2012. Poor countries or poor people? Development assistance and the new geography of global poverty. J. Int. Dev. 24, 686–695.
Kemp, M.C., Wong, K., 1993. Paradoxes associated with the administration of foreign aid. J. Dev. Econ. 42, 197–204.
Kemp, M.C., Long, N.V., Shimomura, K., 1990. On the optimal timing of foreign aid. Kobe Bus. Econ. Rev. 35, 31–49.
Kopczuk, W., Slemrod, J., Yitzhaki, S., 2005. The limitations of decentralized world redistribution: an optimal taxation approach. Eur. Econ. Rev. 49, 1051–1079.
Kraay, A., Raddatz, C., 2007. Poverty traps, aid, and growth. J. Dev. Econ. 82, 315–347.
Lahiri, S., Raimondos-Møller, P., 1997. Competition for aid and trade policy. J. Int. Econ. 43, 369–385.
Léonard, D., Long, N.V., 1992. Optimal Control Theory and Static Optimization in Economics. Cambridge University Press, Cambridge, UK.
Lowe, M., Papageorgiou, C., Perez-Sebastian, F., 2013. The Public and Private MPK. MIT, (Manuscript).
Mankiw, N.G., Romer, D., Weil, D.N., 1992. A contribution to the empirics of economic growth. Q. J. Econ. 107, 407–437.
Mirrlees, J.A., 1971. An exploration in the theory of optimum income taxation. Rev. Econ. Stud. 38, 175–208.
Obstfeld, M., 1999. Foreign resource inflows, saving, and growth. In: Schmidt-Hebbel, K., Serven, L. (Eds.), The Economics of Saving and Growth. Cambridge University Press, Cambridge.
Ohanian, L., Raffo, A., Rogerson, R., 2008. Long-term changes in labor supply and taxes: evidence from OECD countries, 1956–2004. J. Monet. Econ. 55, 1353–1362.
Rebelo, S., 1992. Growth in open economies. Carnegie–Rochester Conference Series on Public Policy 36, pp. 5–46.
Scholl, A., 2009. Aid effectiveness and limited enforceable conditionality. Rev. Econ. Dyn. 12, 377–391.
Steger, T.M., 2000. Productive consumption and growth in developing countries. Rev. Dev. Econ. 4, 365–375.
Steger, T.M., 2002. Productive consumption, the intertemporal consumption trade-off and growth. J. Econ. Dyn. Control. 26, 1053–1068.
Steger, T.M., 2009. Economic growth with subsistence consumption. J. Dev. Econ. 62, 343–361.
Sumner, A., Mallett, R., 2013. The Future of Foreign Aid. Palgrave Macmillan, Basingstoke.
Temple, J.R.W., 2010. Aid and conditionality. In: Rodrik, D., Rosenzweig, M.R. (Eds.), Handbook of Development Economics vol. 5. Elsevier, pp. 4415–4523.
Trimborn, T., Koch, K.J., Steger, T.M., 2008. Multidimensional transitional dynamics: a simple numerical procedure. Macroecon. Dyn. 12, 301–319.
Turnovsky, S.J., 2009. Capital Accumulation and Economic Growth in a Small Open Economy. The MIT Press, Cambridge.
van der Ploeg, F., Venables, A.J., 2013. Absorbing a windfall of foreign exchange: Dutch Disease dynamics. J. Dev. Econ. 103, 229–243.
Weyl, E.G., 2014. The Openness-Equality Trade-off in Global Redistribution. University of Chicago, (Manuscript).