DISCUSSION PAPER SERIES

No. 9596

DYNAMIC AID ALLOCATION

Patrick Carter, Fabien Postel-Vinay and Jonathan Temple

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP9596.php**

ISSN 0265-8003

DYNAMIC AID ALLOCATION

Patrick Carter, Bristol University
Fabien Postel-Vinay, University of Bristol and CEPR
Jonathan Temple, University of Bristol and CEPR

Discussion Paper No. 9596
August 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9596

August 2013

ABSTRACT

Dynamic Aid Allocation*

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.

JEL Classification: F35 and O41
Keywords: aid allocation, economic growth and foreign aid

Patrick Carter
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Email: paddy.carter@bristol.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=171234

Fabien Postel-Vinay
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Email: fabien.postel-vinay@bristol.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=155651

Jonathan Temple
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Email: jon.temple@bristol.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=147806

*This paper was presented as a keynote address at the June 2013 conference of the Global Political Economy Network. We thank conference participants, Mario Crucini, Nigel Duck, Huw Edwards, Ben Zissimos, and seminar audiences at Bristol 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 under the Junior Fellowship scheme.

Submitted 31 July 2013

1 Introduction

Since the 1969 Pearson Commission, there has been a standard benchmark for the generosity of foreign aid programmes: 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. The scope for global redistribution is limited by diminishing returns to aid. The framework is relatively tractable and could be extended in many directions. Even in a simple form, it gives rise to a range of new results on the optimal timing and generosity of foreign aid.

In more detail, we model both the global North and South as neoclassical Ramsey economies that 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 budget constraints; studying this effect in isolation should contribute to an enhanced understanding of aid policies.

A further motivation is the growing interest in the use of 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 traditional government-to-government aid. The Indian government recently announced an ambitious Direct Benefit Transfer scheme, intended eventually to replace multiple welfare programs with cash transfers to households. Using evidence from a randomized experiment, Gertler et al. (2012) find that cash transfers to poor households are partly invested. Given this recent work, there is a case for understanding this particular aspect of aid programmes: what happens when aid is used primarily to relax household budget constraints, and what are the implications for optimal aid policies?

In our framework, the assumption of concave utility functions implies that some degree of 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. Once the donor is also modeled as a Ramsey economy, the opportunity cost of aid for the donor is similarly modest. Hence, our approach sometimes implies that donors should be relatively 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.

In our calibration of the model, donor and recipient initial conditions are based on data from the Penn World Table, combined with assumptions on structural parameters. As well as considering isoelastic (CRRA) preferences, we also analyze the implications of supposing that donor and recipient utility functions take the Stone-Geary 'subsistence consumption' form. When the donor and recipients have Stone-Geary preferences, the effects of aid on investment and growth are stronger than in the CRRA case. The donor maintains a higher level of aid generosity for a longer period. The recipient's desire to escape its subsistence constraint means that the initial change in investment is much higher than under CRRA preferences, and the peak impact on consumption occurs later.

The limits to recipient absorptive capacity play an important role throughout the analysis. When recipients are capable of absorbing relatively high levels of aid, the optimal path of aid is typically front-loaded. But if the absorptive capacity of the South is low, we show that the North's 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 will be able to 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, perhaps surprisingly, that 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 (2011)). 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 then indicate large changes in aid generosity over time and a division of aid between recipients that swings from one to the other. Given diminishing returns to aid intensity, most of the aid may be directed at the middle-income recipient initially, with the allocation switching towards the smaller, poorer recipient as the two recipients grow.

The remainder of the paper will expand on these points, and is structured as follows. The next section sketches the building-blocks of our approach, and relates it to the literature. Section 3 provides the formal description and analysis of the donor's decision problem. Section 4 describes the assumptions used in our simulations. Section 5 presents initial analysis of the optimal path of aid and Section 6 extends the analysis to subsistence economies. Section 7 considers the case of two recipients which differ in population size and income. Section 8 carries out a sensitivity analysis and discusses some possible extensions, before section 9 concludes.

2 Discussion

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 programmes. 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, and to absorption capacity? These are all debates that can be informed by the approach that we develop here.

The most common justification for aid programmes starts from the wide disparities in living standards across countries, providing a case for some form of global redistribution (for example, Tinbergen (1976)). It remains the case that global interpersonal inequality in income arises mainly from between-country differences in mean income, rather than income variation within countries (for example, Milanovic (2005), chapter 9). The optimal degree of redistribution will then depend on the extent of between-country differences, the nature of preferences, and the extent of absorption constraints.

Based on these considerations, we can revisit some long-standing questions. As we noted in the introduction, the Pearson Commission benchmark for donors has been in place for more than four decades; it appears in United Nations Resolution number 2626, from October 1970. It 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. Clemens and Moss (2007) observe that it seems backwards to determine aid levels based on the size of donors, rather than conditions in recipient economies.

Instead, 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 these two roles will help to understand their implications for aid policies and allocation decisions. In our calibration of the model, 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. 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.

A brief sketch helps to convey the main ideas of the paper. The Northern donor seeks to maximize the following objective function:

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

where c_{Nt} is per capita consumption in the North, c_{St} is per capita consumption in the South, $u\left(\cdot\right)$ is the instantaneous utility function, a_t is aid, ω is the relative weight that the Northern donor places on the utility of the South, L_S is the relative population of the South, and other notation is standard. The choice of aid matters through the budget constraints of the North and South: if positive, aid 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. The framework can be used to study normative questions: how generous should aid be, and what does the optimal time profile look like?

We consider two different specifications for the donor's problem, one in which the optimal policy is time-consistent, and one in which it is not. In the latter case, we will follow much of the literature on conditionality, and assume that the donor has access to a commitment technology. In practice, 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. The consideration of time-consistent policies within our framework would be an interesting area for further work.[1]

Again for simplicity, 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. This means that we can consider the donor's problem without having to allow for 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.

To see this, we can contrast our work with Scholl (2009). In her study, Scholl analyzes aid conditionality in terms of a dynamic contracting problem, where the contract is required to be self-enforcing. In many respects this analysis is much richer than ours, but it is not readily adapted to studying normative questions. Scholl defines the objective function of the donor as increasing in the utility of the South and decreasing in a convex function of aid, rather than including the utility of Northern consumption directly. This approach is less suitable for the questions that we are interested in, because the relevant parameters and functional forms correspond less readily to the structure of the North's decision problem. Scholl pins them down by assuming the cost of aid to the North is quadratic, and then, in her calibration, forcing the model's equilibrium to match the observed share of aid in Southern GDP. Since we are primarily interested in normative questions, we do not want to constrain the share of aid in southern GDP to match the data. Further, since we model the North's decision as that of a utilitarian social planner in which utility from Northern consumption enters

[1] As we discuss later in the paper, the time inconsistency arises due to an externality. In the cases considered in our simulations, the effects of this externality are relatively modest. See Section 8 below.

directly, the opportunity costs of aid arise endogenously from the structure of the model. This eliminates the arbitrariness in the donor's objective function that would otherwise be present. A further advantage of our approach is that it yields a mapping between an explicit weight ω on Southern utility and optimal aid policies.

One simplification is that we model absorption constraints as a reduced form. Political economy approaches often imply that some aid is wasted or diverted. Rather than develop a specific structural model, we adopt a general 'aid impact function' which translates a given Northern aid flow into the actual transfers received by Southern households, allowing for some degree of wastage or diversion.[2] For simplicity of exposition, we start with the case with a fixed population and no technical progress. Capital accumulation in the South is given by

$$\dot{k}_{St} = A_S f_S\left(k_{St}\right) + g\left(a_t, \bar{k}_{St}; L_S, A_S\right) - c_{St} - \delta k_{St}$$

where k_{St} is household capital per worker in the South, $\bar{k}_{St}$ is the South's aggregate capital-labour ratio, A_S is TFP in the South, and a_t is aid. Hence aid enters the budget constraint not in a linear way, as in Obstfeld (1999), but through the aid impact function $g\left(\cdot\right)$. The semi-colon in the function separates the endogenous variables, aid and capital, from the exogenous variables such as the relative population and TFP of the South. For much of the paper, the dependence on these exogenous variables will be suppressed for notational simplicity.

This 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).[3] One possible story is that, as aid intensity increases, so does the proportion of those aid flows that are wasted or appropriated by a local elite. If the proportion wasted is increasing in aid intensity, this has an immediate corollary: as countries develop, they will be better placed to absorb aid effectively.

Although wastage is the simplest motivation for diminishing returns to aid intensity, others are possible. High aid intensity may be associated with an increase in the relative price of non-traded goods, leading to a relative contraction of the traded sector (for example, Rajan and Subramanian (2011) and van der Ploeg and Venables (2013)); with adverse effects on a domestic political equilibrium, partly by influencing rents to sovereignty, and perhaps by undermin-

[2]The terminology 'aid impact function' is taken from the early paper by Dudley and Montmarquette (1976), which used a related idea in the study of aid allocation.

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

ing long-term accountability and state capacity (for example, Moore (2004)); and with a proliferation of aid projects and programmes which could overwhelm the bureaucracy of the domestic government (for example, Kanbur (2006)). These mechanisms have been widely discussed in the literature. As a consequence, it seems essential to allow for limits to the effective absorption of aid.

A less central assumption is that the capital account is closed. This can be seen as an extreme case of international credit frictions, which lead the marginal product of capital to be higher in poorer countries. It might be objected that poorer countries do not need aid, because capital accumulation can be financed by capital inflows. In practice, capital has tended to flow to middle-income countries rather than the poorest countries. And even in the presence of an open capital account, there would still be a role for aid in our framework, to finance consumption. The consideration of an open capital account would be an interesting extension to our analysis.[4]

We now briefly discuss the relevant literature. The overall aims of our paper are similar to those of Kopczuk et al. (2005), who also provide a formal and quantitative analysis of global redistribution. They consider optimal taxation problems in which taxes can be used to finance international transfers, but have disincentive effects. This leads to an account of the opportunity cost of foreign aid which is much richer than ours, but our analysis lends itself more easily to dynamic considerations and the derivation of optimal aid paths over time. Moreover, our findings will apply to a Northern donor acting in isolation, whereas implementation of the Kopczuk et al. (2005) tax-and-transfer policies would require something close to a world government.

The closest precursor to our paper is the optimal control approach of Kemp et al. (1990). They consider 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 budget constraint is linear in aid. They show that the optimal policy of the North involves an immediate 'stock transfer', which transfers part of the North's capital stock to the South at time zero.[5] In our work, we rule out stock transfers of this form, a step which is a natural counterpart to the assumption of

[4]Its relevance is suggested by the empirical findings of Caselli and Feyrer (2007), who estimate returns to capital directly and show that returns are broadly comparable in a sample of 53 countries. Note, however, that their sample is dominated by rich countries and middle-income countries; only 7 of the 53 countries are in sub-Saharan Africa.

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

absorption constraints, as they note. A second difference between the two papers is that Kemp et al. (1990) is purely theoretical, whereas we present quantitative results based on simulations.

Other models of North-South interactions have been studied. Some of these explore the transfer problem, and especially the possibility of transfer paradoxes driven by terms-of-trade effects; see Brakman and van Marrewijk (1995), Brakman and van Marrewijk (1998), Lahiri and Raimondos (1995) and Shimomura (2007) for examples and references, and Eaton (1989) for discussion of the older literature. Shimomura considers a North-South model with two goods and shows that an exogenous transfer can even be Pareto improving, leaving neither region worse off. These analyses tend to be static. Fisher (1995) considers interactions between two open economies with the potential for endogenous growth, and studies the dynamic effects of lump-sum transfers of fiat assets. Transfers to a relatively patient recipient will make current generations in the donor country worse off, but increase the wages of infinitely many future generations in both countries. The majority of this work is theoretical, and studies the effects of exogenous transfers rather than deriving their optimal time path. A paper which takes a more quantitative approach is that by Chamon and Kremer (2009). They construct a multi-country model in which developing countries gradually integrate with the world economy, and discuss the potential role of aid in accelerating this process. However, aid is not included in the version of the model that they calibrate.

A companion paper to ours, Carter (2012), uses some of the ideas of this paper to study the properties of aid allocation rules. In that paper, donors allocating aid across multiple recipients commit to a rule in which transfers are a simple function of recipient characteristics. For particular parameter values, Carter establishes the optimal weights on recipient characteristics within a given rule, and compares them with allocation rules that have been adopted in practice. He shows that, contrary to the usual assumption in the literature, the poorest countries should not always be the first priority. This is a consequence of absorption constraints: if there are diminishing returns to aid, concentrating aid heavily on the poorest recipients will sometimes be sub-optimal.

3 The model

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 technological progress and population growth to zero, but the necessary extensions are straightforward.

3.1 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:

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

where k_{St} is the Southern household's capital-labour ratio, and $\bar{k}_{St}$ the South's aggregate capital-labour ratio. We assume that $g\left(0, \bar{k}_{St}\right) = 0$, and that $g(\cdot)$ is non-decreasing in $\bar{k}_{St}$ and has strictly diminishing marginal returns to a_t. As indicated previously, the aid impact function also depends on exogenous variables such as the population size of the South, we suppress this dependence for notational simplicity.

The representative Southern household takes the paths of aid a_t and aggregate capital as given. The above problem leads to the standard 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, and $k_S = \bar{k}_S$. Due to the aid impact function, capital accumulation in the South involves a positive externality: it improves the future effectiveness of aid. In the decentralized model, Southern households do not internalize this effect of aggregate capital, and invest too little from society's point of view. In contrast, a Southern planner would choose a higher saving rate in the early stages of the transitional dynamics, to improve future aid absorption. In this alternative version of the model, the South's Euler equation is given by:

$$\frac{\dot{c}_{St}}{c_{St}} = -\frac{1}{\varepsilon_{u'(c)}} \left\{ A_S f_S' \left(k_{St} \right) + g_2' \left(a_t, k_{St} \right) - \rho - \delta \right\}. \tag{2}$$

3.2 The decision problem for the North

Suppose the South is ruled by a benevolent dictator, so that (2) is the relevant Euler equation. The North's decisions must respect $k_S = \bar{k}_S$. The Northern planner then solves:

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

$$\begin{aligned}
\text{subject to:} \quad & \dot{k}_{Nt} = f_N\left(k_{Nt}\right) - c_{Nt} - a_t - \delta k_{Nt} \\
& \dot{k}_{St} = A_S f_S\left(k_{St}\right) + g\left(a_t, k_{St}\right) - c_{St} - \delta k_{St} \\
& \frac{\dot{c}_{St}}{c_{St}} = -\frac{1}{\varepsilon_{u'(c)}} \left\{ A_S f_S'\left(k_{St}\right) + g_2'\left(a_t, k_{St}\right) - \rho - \delta \right\}, \\
& k_{S0},\ k_{N0} \text{ given}, c_{S0} \text{ free}
\end{aligned}$$

where we have normalized the TFP of the North to unity. This formulation rules out time-zero stock transfers of the Kemp et al. (1990) type. We can now derive some results. Denoting the costate variables associated with k_{Nt}, k_{St} and c_{St} as x_{Nt}, x_{St} and z_{St}, respectively, we have the following set of optimality conditions.[6] The first-order condition for c_{Nt} is the usual one: $u'\left(c_{Nt}\right) = x_{Nt}$. The first-order condition for a_t is:

$$x_{Nt} = g_1'\left(a_t, k_{St}\right) x_{St} - \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} g_{12}''\left(a_t, k_{St}\right). \tag{3}$$

The three Euler equations are:

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

$$\dot{x}_{St} = \left\{ \rho + \delta - A_S f_S'\left(k_{St}\right) - g_2'\left(a_t, k_{St}\right) \right\} x_{St} + \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} \left\{ A_S f_S''\left(k_{St}\right) + g_{22}''\left(a_t, k_{St}\right) \right\} \tag{5}$$

$$\dot{z}_{St} = \left\{ \rho + \frac{1}{\varepsilon_{u'(c)}} \left[A_S f_S'\left(k_{St}\right) + g_2'\left(a_t, k_{St}\right) - \rho - \delta \right] \right\} z_{St} + x_{St} - \omega L_S u'\left(c_{St}\right) \tag{6}$$

[6] We assume throughout that the necessary conditions for optimality will indeed characterize an optimal solution. In future work, we will examine whether these conditions are sufficient in the context of our particular problem.

Optimality at the initial date further requires $z_{S0} = 0$.[7] In fact, the Northern planner's problem has a consistent solution with $z_{St} \equiv 0$ at all dates. If the South is run by a planner, the capital externality is internalized by the South 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 still 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 (1) replaces (2) as the last dynamic constraint in the Northern planner's problem. Again requiring $k_S = \overline{k}_S$ in the North's problem, the first-order condition for aid becomes:

$$x_{Nt} = g_1' \left(a_t, k_{St} \right) x_{St}, \tag{7}$$

The Euler equation for x_{St} is:

$$\dot{x}_{St} = \left\{ \rho + \delta - A_S f_S' \left(k_{St} \right) - g_2' \left(a_t, k_{St} \right) \right\} x_{St} + \frac{z_{St} c_{St}}{\varepsilon_{u'(c)}} A_S f_S'' \left(k_{St} \right), \tag{8}$$

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

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

Again optimality requires $z_{S0} = 0$, but no solution exists with $z_{St} \equiv 0$. The Northern planner's Euler equation (8) would then be incompatible with the South's Euler equation (1) because of the capital externality. (For example, one can see that $z_S > 0$ in the steady state.) This externality is also the source of time inconsistency. The social returns to capital are higher than perceived by Southern households, given the impact of higher income on aid effectiveness. Hence, the Northern planner is less generous to the decentralized South than it would be to a social planner in the South: this maintains the returns to capital and encourages capital accumulation. 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, and so the North's aid decision is time inconsistent.

[7] 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 budget constraints and the Euler equation for Southern consumption. See Léonard and Long (1992), chapter 7, especially section 7.7.

The model based on decentralized investment decisions is arguably more realistic, and that model is the focus of our later simulations. In those simulations, the quantitative importance of the externality is relatively modest. The externality does not greatly affect the optimal generosity of aid, but leads to higher investment in the early stages of the transition. These effects are not large, 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).

3.3 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 a set of equations which implicitly define the steady-state level of aid, a^*, and consumption and capital for both North and South:

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

The first equation has a natural interpretation: the Northern donor balances the marginal cost of aid - represented by the marginal utility 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.[8]

For the alternative case where Southern consumption decisions are decent-

[8]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 $g_1' \left(a^*, k_S^*\right)$ on L_S.

ralized to households, the steady-state solution takes a more complicated form:

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

where 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. It plays a minor role if $g_2'\left(a^*,k_S^*\right)$ is close to zero, in which case the externality is modest. Otherwise, its importance is increasing in ρ and the absolute value of $\varepsilon_{u'(c_S)}$ and decreasing in A_S and the absolute value of $f''(k_S^*)$.

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:

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

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, it can be shown that the left-hand-side is increasing in aid for the case we consider in our simulations, which has CRRA utility, a specific $g\left(a,\bar{k}_S\right)$ 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 population in the South (higher L_S) and higher total factor productivity in the North will lead to higher steady-state aid.[9] Given the externality, other comparative static results are harder to derive, and some are unexpectedly ambiguous: for example, under CRRA utility, equilibrium aid is not necessarily decreasing in the intertemporal elasticity of substitution.

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.

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

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 absorption function so that negative transfers are costly for the North to implement.

4 Calibration assumptions

We will use the above model to derive some quantitative findings, and qualitative results that could inform the broad principles of aid programmes. This section describes the assumptions we will use in our simulations. Although the approach is stylized, we note that formal analysis of foreign aid has often relied on simpler and less satisfactory models; see Easterly (1999) for more discussion.

For the instantaneous utility function, we adopt the CRRA form:

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

where the elasticity of marginal utility with respect to consumption is $-\sigma$, and σ is also the inverse 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.[10] We assume the time preference parameter ρ is equal to 0.03; this is the continuous-time equivalent of that used by Obstfeld (1999) and Gourinchas and Jeanne (2006). Note that the choices of 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, the discussion of the Ramsey model in Barro and Sala-i-Martin (2004). Our choice of 0.50 has been employed in related contexts, such as Kraay and Raddatz (2007). This parameter choice is important because it will influence the speed at which the South converges to its steady-state, and the extent to which the marginal product of capital will fall along the transition.

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

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 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 income 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.[11] This keeps the quantitative analysis comparable with the aid decisions made in practice. As is well known, aid receipts per capita are low for India and China, partly reflecting the 'small country bias' in aid allocation; excluding these two countries is the simplest way 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 whole South sub-divided into low and middle income recipients, for use in section 7.)

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 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 the capital-labor ratio in efficiency units. We can then infer the relative TFP and GDP per capita of the South.

In order to ensure the existence of a steady-state balanced growth path, we assume that rates of technical progress and population growth are the same for donor and recipient. 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

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

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.

Table 1: Calibration

	N	S	S′	S1	S2
Population (L)	901	4844	2393	680	4164
Capital Stock (K)	103600	56030	25900	1320	54710
Output (Y)	31410	28050	13330	1070	26990
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

One of the main uncertainties in the calibration 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 wedge between the transfer made by the North, and the aid received by the representative household in the South. Our chosen functional form, mapping total aid a_t into transfers to households, is:

$$g\left(a_t, \bar{k}_{St}\right) = \frac{a_t}{L_{St}}\left(1 - \frac{a_t}{vY_S(\bar{k}_S)}\right) \quad (10)$$

where $Y_S(\bar{k}_S)$ is output in the South, and higher values of v correspond to greater aid effectiveness for any given capital-labor ratio. The ratio of investment to GDP is then a quadratic function of aid intensity, just as the empirical literature sometimes relates growth to 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 $a_t/Y_{St} = v/2$. The appropriate severity of this absorption constraint is a matter of debate. In our baseline, we set v so that an extra dollar of aid has zero impact once an aid intensity of 25% is reached.[12]

A final choice relates to ω, the relative weight of Southern utility in the objective function of the North. 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.

4.1 Welfare analysis

To compare the lifetime utilities associated with different consumption paths, we use the 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 CRRA 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.

5 Optimal aid policies

In this section, we study the optimal path for aid. To do this, we use the relaxation algorithm of Trimborn et al. (2008) to simulate the system of differential

[12]Our chosen functional form implies that aid impact becomes negative beyond this level, which may seem unrealistic given a 'wastage' interpretation of the aid impact function, and a functional form in which aid impact tends to zero may seem preferable. See Carter (2012) for an alternative functional form.

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 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 aid more effectively.

The first panel of Figure 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.[13] 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. Figure 1 also shows the impact of aid on the rate at which Southern income and consumption converge towards Northern levels, and the impact of aid upon the Southern savings rate (expressed as deviation from a zero-aid counterfactual). The front-loading of aid enables the South to raise consumption by 14.4% immediately.[14]

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

[13] 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, 2011).

[14] This contrasts with Obstfeld (1999), who studied the effect of a fixed quantity of aid, and found the greatest impact on consumption would take place several decades after aid had commenced.

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.

Scenarios in which aid increases over time are shown in figures 2 and 3. For some combinations of ω (the weight on the South's utility) and v (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.

We now discuss the effect of requiring the Northern donor to maintain aid as a fixed share of its GDP, as in the Pearson Commission 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 often 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.[15] We also find that, if aid is fixed as a share of Northern GDP, departures from the optimal share are not costly: the donor's objective function is relatively flat as a function of the aid share.

Figure 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 utility loss associated with fixing the aid share is equivalent to a 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 the 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.

What happens if the aid share in Northern GDP is fixed at the Pearson Commission benchmark, the 0.70% figure that is often discussed? Compared to our optimal flexible path, the lower generosity means that the HEV gain in the South falls from 11% to 2%.

[15] This result could be eliminated by a more realistic aid rule, such as one in which the share of aid in Northern GDP was assumed to be declining in the GDP per capita of the South.

6 Optimal Aid for Subsistence Economies

We now consider preferences with a subsistence level of consumption. It is well known that the transitional dynamics of the single-sector Ramsey model with CRRA preferences cannot account for a number of empirical regularities. That model typically predicts a rate of convergence to steady-state higher than observed, and often, too sharp a decline in saving rates and in the marginal product of capital as capital is accumulated; nor can it readily accommodate periods of international divergence in GDP per capita. Rebelo (1992), Ben-David (1998) and Steger (2009) have all argued that a Ramsey model with Stone-Geary (SG) preferences is an attractive way to overcome these deficiencies. We examine the implications of these preferences for the generosity and timing of aid, including whether donors will want to front-load aid, and whether imposing a fixed aid share for the North is relatively more costly in this case.

Under Stone-Geary preferences, the qualitative nature of the dynamic path will depend on initial conditions. In a model with technical progress, there will ultimately be convergence for a wide range of initial capital stocks, but depending on initial conditions, the South's capital per worker and GDP per worker may initially grow more slowly than the rate of technical progress, and living standards in the North and South will temporarily diverge (Ben-David (1998)). We will study the effect of aid on the extent of this divergence.

The framework now becomes:

$$\max_{\{c_{Nt},a_t\}} \int_0^\infty \left[\frac{(c_{Nt}-\bar{c})^{1-\sigma}-1}{1-\sigma} + \omega \cdot L_S \cdot \frac{(c_{St}-\bar{c})^{1-\sigma}-1}{1-\sigma} \right] \cdot e^{-\rho t} dt$$

$$\begin{aligned}
\text{subject to:} \quad & \dot{k}_{Nt} = f_N(k_{Nt}) - c_{Nt} - a_t - \delta k_{Nt} \\
& \dot{k}_{St} = A_S f_S(k_{St}) + g\left(a_t, \bar{k}_{St}\right) - c_{St} - \delta k_{St} \\
& \dot{c}_{St} = \frac{(c_{St}-\bar{c})}{\sigma} \cdot \left\{ A_S f_S'(k_{St}) - \rho - \delta \right\}, \\
& k_{S0},\ k_{N0} \text{ given},\ c_{S0} \text{ free.}
\end{aligned}$$

where $\bar{c}$ corresponds to a subsistence level of consumption. The extension to the case of labor-augmenting technical progress is straightforward. In that case, although the threshold $\bar{c}$ matters at the start of a transition, it is ultimately irrelevant, because we hold it fixed in absolute terms even as technology improves. Hence, this model will have an asymptotic balanced growth path as in Ohanian et al. (2008).

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.

Table 2: Output Growth in South and Subsistence Consumption

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

As expected, under SG preferences, aid generosity should be especially high early in the development process. We present some results in Figure 5, which could be compared with the earlier CRRA results in 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 income 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 dynamic flow of 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 aid reduces the otherwise high opportunity cost of investment when close to subsistence.

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 rate of output growth, 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, at 3.5%.[16] So the impact of aid upon 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, Figure 6 shows outcomes based on the same fixed quantity of aid.[17] The higher aid intensity in the SG case reflects the fact that Southern output grows more slowly under these preferences. The effect of aid on net investment, consumption and growth is markedly greater in the SG case. Since our assumptions on preferences vary across the two cases, more direct welfare comparisons are not meaningful.

Figure 7 compares the optimal fixed-share policy with fully flexible aid under SG preferences. The fixed-share constraint on the donor leads to greater generosity at longer horizons. Although we do not present the welfare results in detail, we continue to find that a fixed-share rule can leave the South better off because of the greater generosity at long horizons. There is little evidence to support the idea that a fixed-share rule would be more costly under SG preferences.

7 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 another country.

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. For this calibration, we use the full sample of countries including 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

[16]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.

[17]The fixed level chosen is that which is optimal in the CRRA case. The third panel shows the percentage change in the growth rate as opposed to the absolute difference.

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 in the case of $S2$. 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 output growth is 13% for $t = 0$ in $S1$, 6% in $S2$.

This calibration allows us to study an important policy issue: does low consumption help to motivate aid even for middle-income countries? Figure 8 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 impact upon saving and investment (which mirrors the impact upon income 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 saving behaviour, 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 less binding, the donor is generous early on but then scales back aid 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 now consider three variants. First, we change the population of $S2$ so that it matches that of $S1$. Next, we change the levels of TFP and of the initial capital stock in $S2$, such that initial per capita income is unchanged but the economy is now on its balanced growth path. Finally, we relax the constraint of the aid impact function for $S1$ by using $v = 1$.

Figure 9 shows what happens when $S1$ and $S2$ have the same population sizes. This leads to a sharp reduction in the quantity of aid given to $S2$, but qualitatively the paths are unchanged. It is worth noting that the level of aid given to $S1$ is virtually unchanged, despite the large reduction in aid given to $S2$

with an associated reduction in the opportunity cost of aid to the North. This result arises because $S1$ is already being given as much aid as it can effectively absorb.[18]

The next experiment retains equal population sizes, but the initial capital stock and the level of TFP for recipient 2 are adjusted so that this recipient is now on its balanced growth path (retaining the initial output per capita seen in the data). Hence, we are now examining the sensitivity of aid allocation to the distance from steady-state. In effect, the recipient $S2$ has no prospect of aid-funded growth, because it is already on its balanced growth path. Figure 10 shows that the result is a constant flow of aid to $S2$, expressed as a proportion of donor GDP, and a time-path for total aid that is back-loaded. This reflects the back-loading of aid to the low-income recipient $S1$. Again, the quantity of aid given to $S1$ is barely affected by the change in circumstances of $S2$; this relies on treating the total aid budget as endogenous at each instant, rather than considering the division of a fixed budget.

The donor allocates a flow of aid to $S2$ that is almost as high as that given in the former case, when $S2$ was below its steady-state and aid contributed to faster convergence. Aid intensity in $S2$ is sustained at 13% of Southern GDP, in contrast to the former experiment, in which aid intensity is initially 15% but falls below 13% after two years and falls to just 0.5% in the long run. This mirrors a finding in Carter (2012): aid is less valuable to economies with strong growth prospects, because these economies will see rapid increases in their output and consumption even in the absence of aid.[19]

In our final experiment, we relax the aid absorption constraint in the low-income recipient $S1$. The results are shown in Figure 11. The low-income recipient is sufficiently impoverished that the donor again pushes aid up against the absorption constraint, and increases its generosity over time. The impact on investment is relatively muted. Increased generosity to $S1$ has a noticeable impact on $S2$, where aid intensity falls from 0.5% in year 50 to zero, and aid is given solely to $S1$.

[18]Initial aid intensity is close to 25% of Southern GDP, the point at which the marginal benefit of aid to the South is zero.

[19]To arrive at a more conventional result, we would need to assume that aid is used more effectively in countries which also have good growth prospects, as argued in Burnside and Dollar (2000). In our framework, we would need a variable such as institutional quality to enter the aid impact function and, separately, to influence growth prospects.

8 Sensitivity Analysis

8.1 Optimal Aid with a Social Planner in the South

Earlier, we noted the externality that arises in our setup: capital accumulation by Southern households increases the benefits of aid, since growth in the South improves its ability to absorb aid. We also noted that, for our baseline calibration, the quantitative importance of this externality is relatively modest. We now briefly demonstrate this, by contrasting outcomes under social planning with the decentralized case.

To recap, a Southern planner would internalize the aid absorption externality to capital accumulation. Under CRRA preferences, this yields the following decision rule in the South:

$$\dot{c}_{St} = \frac{c_{St}}{\sigma} \cdot \left\{ f'_S\left(k_{St}\right) + g'_2\left(a_t, k_{St}\right) - \rho - \delta \right\}.$$

With a social planner in the South, the optimal policy for the Northern planner now leaves the Northern households slightly worse off (equilibrium aid is slightly more generous) and the South slightly better off. Figure 12 shows that the effects on the North are modest. In the figure, the optimal path for aid with a Southern social planner is indistinguishable from that in the decentralized version. Looking at the results more closely, 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 behavior is noticeable: the Southern planner anticipates the favorable effect of investment on future aid absorption, and hence chooses a path for investment which leads to higher investment and higher growth in the first few years of the transitional dynamics. Aid as a share of Southern GDP is similar in the two cases, because aid and output increase in similar proportions.

8.2 Parameter variation

We now 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 in the paper, the curvature of the utility function has a major effect on optimal generosity, especially at longer horizons. With slightly less curvature ($\sigma = 1.8$ rather than $\sigma = 2.0$) aid generosity is reduced and the effect on consumption levels is smaller. Given the presence of technical progress, 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 impact on growth is greater. Intuitively, because investment now has greater benefits, the South is able to consume more early on.

8.3 Further discussion

We now briefly discuss some possible extensions. Perhaps the most important modification to the analysis would to be introduce a role for government and political economy considerations within the South. As we noted earlier, this introduces 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). Expenditure on public investment projects could be wasteful to some degree, perhaps with the extent of waste increasing in expenditure as in Berg et al. (2013).

A less drastic motivation would be to introduce income heterogeneity within aid recipients: since the world's poor are increasingly located within middle-income countries, it would be interesting to explore the implications for aid policies (Kanbur and Sumner (2011)). Somewhat related, perhaps a version of the North's objective function could be derived from a model in which the North seeks political influence, as in the work of Antràs and Padró i Miquel (2011).

Another extension would be to 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 these revenues as exogenous, which seems better suited for natural resource windfalls than aid flows. The integration of their decision problem within that of an altruistic

North, using the nested structure we adopt here, would be an interesting area for future work. Similarly, another natural extension would be to model recipients as two-sector economies that produce traded and non-traded goods. The integration of Dutch Disease effects could lead to a structural model of absorption constraints, rather than our reduced-form approach. A related idea would be to explore the implications of the analysis in Chatterjee and Naknoi (2010).

Staying closer to our current framework, other possibilities include country-specific population and technological progress dynamics, costs of distortionary taxation in the donor, adjustment costs in the North for aid disbursements, and finite-duration aid commitments by the donor. Other possible extensions could include CES production technologies (Turnovsky (2008)) and the introduction of capital varieties, where the latter could increase the welfare effects of international transfers (Hoxha et al. (2013)). One of the most interesting extensions to the basic framework would be to introduce short-run output volatility in the Southern economies, so that aid would have an insurance role as well as increasing mean consumption. This would bring the analysis closer to that of Arellano et al. (2009), at the expense of greater computational complexity.

9 Conclusions

This paper introduced a framework for analyzing optimal aid policies using the neoclassical growth model. In contrast to most previous research on aid policies, the findings are based on a clearly-defined optimization problem for the Northern donor. The model takes into account the opportunity cost of aid for the Northern donor and optimizing behavior by Southern households. 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 level of the recipient's balanced growth path relative to that of the donor, and the recipient's initial distance from that growth path. In our simulations of the model, 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, the optimal path is one in which aid generosity declines over time. But in some scenarios, where donors run up against aid absorption constraints, there is a time interval over which aid 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 both the absolute quantity of aid and the division of aid between recipients.

We also show that, as might be expected, preferences with subsistence consumption strengthen the case for donor generosity. The impact of aid on Southern consumption, investment and output growth can be dramatic in this case. In general, our numerical results indicate optimal levels of aid generosity that exceed those observed in the data, at least at short horizons. An interesting task for future research is to pin down exactly what features of reality 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 should be seen as 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.

References

Antràs, P. and G. Padró i Miquel (2011). Foreign influence and welfare. *Journal of International Economics 84*(2), 135–148.

Arellano, C., A. Bulir, T. Lane, and L. Lipschitz (2009). The dynamic implications of foreign aid and its variability. *Journal of Development Economics 88*(1), 87–102.

Barro, R. J. and X. Sala-i-Martin (2004). *Economic Growth* (Second ed.). Cambridge, Massachusetts: The MIT Press.

Ben-David, D. (1998). Convergence clubs and subsistence economies. *Journal of Development Economics 55*(1), 155–171.

Berg, A., R. Portillo, S.-C. S. Yang, and L.-F. Zanna (2013). Public investment in resource-abundant developing countries. *IMF Economic Review 61*(1), 92–129.

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

Brakman, S. and C. van Marrewijk (1995). Transfers, returns to scale, tied aid and monopolistic competition. *Journal of Development Economics 47*(2), 333–354.

Brakman, S. and C. van Marrewijk (1998). *The economics of international transfers*. Cambridge, UK: Cambridge University Press.

Burnside, C. and D. Dollar (2000). Aid, policies, and growth. *American Economic Review 90*(4), 847–868.

Carter, P. (2012). Aid allocation rules. Working Paper 12/630, University of Bristol.

Caselli, F. and J. Feyrer (2007). The marginal product of capital. *The Quarterly Journal of Economics 122*(2), 535–568.

Chamon, M. and M. Kremer (2009). Economic transformation, population growth and the long-run world income distribution. *Journal of International Economics 79*(1), 20–30.

Chatterjee, S. and K. Naknoi (2010). The marginal product of capital, capital flows, and convergence. *American Economic Review 100*(2), 73–77.

Chatterjee, S., G. Sakoulis, and S. J. Turnovsky (2003). Unilateral capital transfers, public investment, and economic growth. *European Economic Review 47*, 1077–1103.

Clemens, M. and T. Moss (2007). The ghost of 0.7 per cent: Origins and relevance of the international aid target. *International Journal of Development Issues 6*(1), 3–25.

Clemens, M. A., S. Radelet, R. R. Bhavnani, and S. Bazzi (2012). Counting chickens when they hatch: Timing and the effects of aid on growth. *Economic Journal 122*, 590–617.

Cohen, D. and P. Michel (1988). How should control theory be used to calculate a time-consistent government policy? *Review of Economic Studies 55*(2), 263–274.

Crafts, N. (2011). The Marshall Plan: A reality check. Cage online working paper series, University of Warwick.

Dudley, L. and C. Montmarquette (1976). A model of the supply of bilateral foreign aid. *American Economic Review 66*(1), 132–42.

Eaton, J. (1989). Foreign public capital flows. Volume 2 of *Handbook of Development Economics*, pp. 1305 – 1386. Elsevier.

Fisher, E. O. (1995). Growth, trade, and international transfers. *Journal of International Economics 39*(1-2), 143–158.

Gertler, P. J., S. W. Martinez, and M. Rubio-Codina (2012). Investing cash transfers to raise long-term living standards. *American Economic Journal: Applied Economics 4*(1), 164–192.

Gourinchas, P.-O. and O. Jeanne (2006). The elusive gains from international financial integration. *Review of Economic Studies 73*(3), 715–741.

Hanlon, J., A. Barrientos, and D. Hulme (2010). *Just give money to the poor*. Sterling, VA: Kumarian Press.

Hoxha, I., S. Kalemli-Ozcan, and D. Vollrath (2013). How big are the gains from international financial integration? *Journal of Development Economics 103*, 90–98.

Kanbur, R. (2006). The economics of international aid. In S.-C. Kolm and J. M. Ythier (Eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity*. North-Holland, Amsterdam.

Kanbur, R. and A. Sumner (2011). Poor countries or poor people? development assistance and the new geography of global poverty. CEPR discussion paper 8489, CEPR.

Kemp, M. C., N. V. Long, and K. Shimomura (1990). On the optimal timing of foreign aid. *Kobe Business and Economic Review 35*, 31–49.

Kopczuk, W., J. Slemrod, and S. Yitzhaki (2005). The limitations of decentralized world redistribution: An optimal taxation approach. *European Economic Review 49*(4), 1051–1079.

Kraay, A. and C. Raddatz (2007). Poverty traps, aid, and growth. *Journal of Development Economics 82*(2), 315–347.

Lahiri, S. and P. Raimondos (1995). Welfare effects of aid under quantitative trade restrictions. *Journal of International Economics 39*(3-4), 297–315.

Léonard, D. and N. V. Long (1992). *Optimal control theory and static optimization in economics*. Cambridge, UK: Cambridge University Press.

Mankiw, N. G., D. Romer, and D. N. Weil (1992). A contribution to the empirics of economic growth. *The Quarterly Journal of Economics 107*(2), 407–437.

Milanovic, B. (2005). *Worlds Apart: Measuring International and Global Inequality*. Princeton and Oxford: Princeton University Press.

Moore, M. (2004). Revenues, state formation, and the quality of governance in developing countries. *International Political Science Review 25*(3), 297–319.

Obstfeld, M. (1999). Foreign resource inflows, saving, and growth. In K. Schmidt-Hebbel and L. Serven (Eds.), *The Economics of Saving and Growth*. Cambridge: Cambridge University Press.

Ohanian, L., A. Raffo, and R. Rogerson (2008). Long-term changes in labor supply and taxes: Evidence from OECD countries, 1956-2004. *Journal of Monetary Economics 55*(8), 1353–1362.

Rajan, R. G. and A. Subramanian (2011). Aid, Dutch Disease, and manufacturing growth. *Journal of Development Economics 94*(1), 106–118.

Rebelo, S. (1992). Growth in open economies. *Carnegie-Rochester Conference Series on Public Policy 36*(1), 5–46.

Scholl, A. (2009). Aid effectiveness and limited enforceable conditionality. *Review of Economic Dynamics 12*(2), 377–391.

Shimomura, K. (2007). A Pareto-improving foreign aid in a dynamic North-South model. In S. Lahiri (Ed.), *Theory and Practice of Foreign Aid*. Elsevier, Amsterdam.

Steger, T. M. (2009). Economic growth with subsistence consumption. *Journal of Development Economics 62*(2), 343–361.

Sumner, A. and R. Mallett (2013). *The future of foreign aid*. Basingstoke: Palgrave Macmillan.

Temple, J. (2010). Aid and conditionality. In M. R. Rosenzweig and D. Rodrik (Eds.), *Handbook of Development Economics, Volume 5*, pp. 4415–4523. Elsevier.

Tinbergen, J. (Ed.) (1976). *Reshaping the International Order*. New York: Dutton.

Trimborn, T., K.-J. Koch, and T. M. Steger (2008). Multidimensional transitional dynamics: a simple numerical procedure. *Macroeconomic Dynamics 12*, 301–319.

Turnovsky, S. J. (2008). The role of factor substitution in the theory of economic growth and income distribution: two examples. *Journal of Macroeconomics 30*(2), 604–629.

Turnovsky, S. J. (2009). *Capital Accumulation and Economic Growth in a Small Open Economy*. Cambridge: The MIT Press.

van der Ploeg, F. and A. J. Venables (2013). Absorbing a windfall of foreign exchange: Dutch Disease dynamics. *Journal of Development Economics 103*, 229–243.

10 Figures

Figure 1: Baseline Case; Decentralized Version

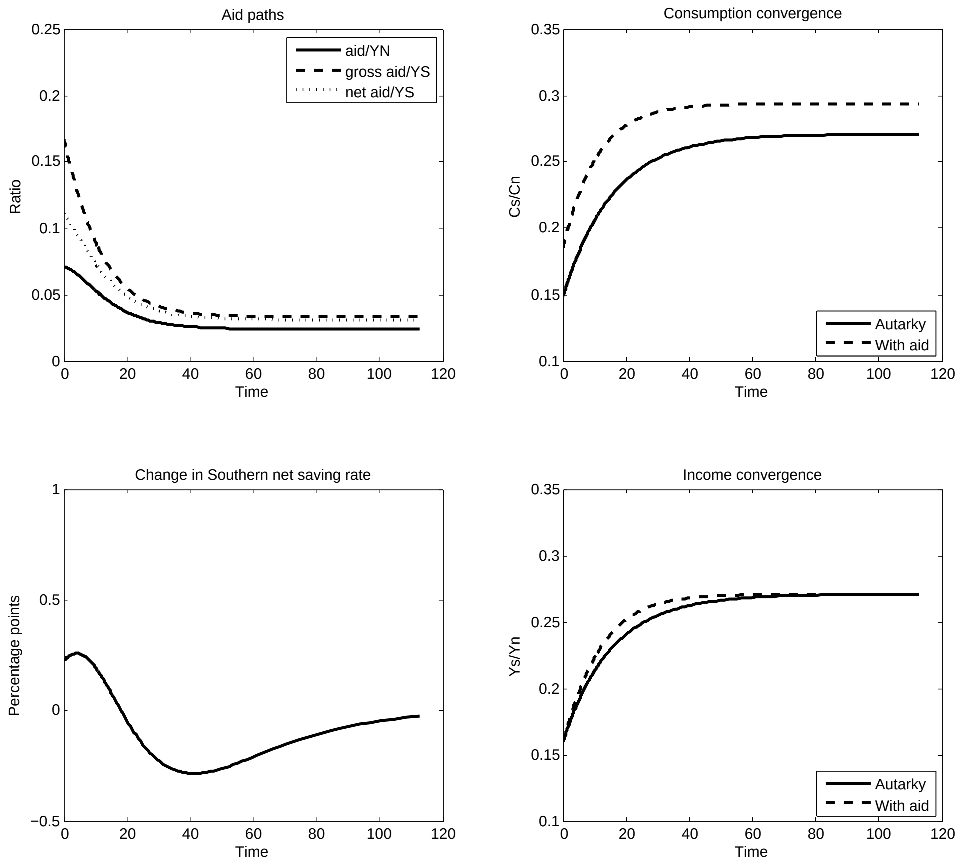

The panel 'Aid paths' shows the optimal path of aid over time, relative to Northern and Southern GDP. The gap between gross and net aid in the South represents wastage. The two right-hand panels show consumption and output in the South relative to the North, an upward slope signifying convergence. In 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.

Figure 2: Decentralized a/Yn, varying absorptive capacity (v)

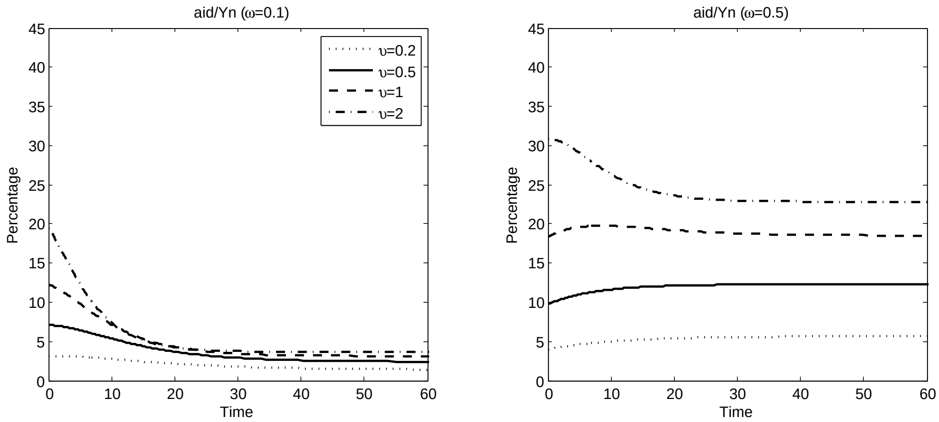

Optimal paths of aid relative to Northern GDP, as Southern absorptive capacity varies, for two degrees of Northern altruism. Smaller values of v correspond to a tighter absorption constraint, with higher levels of waste at low levels of aid intensity. In the left-hand panel the Northern planner places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

Figure 3: Decentralized a/Yn, varying Northern altruism ω

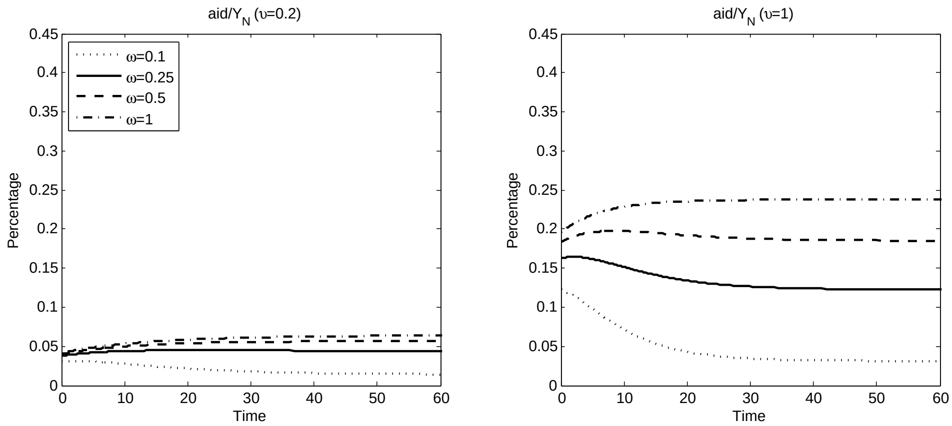

Optimal paths of aid relative to Northern GDP, with Northern altruism increasing in ω, for two degrees of Southern absorptive capacity. In the left-hand panel Southern absorptive capacity is low, $v = 0.2$, and higher in the right-hand panel $v = 1$.

Figure 4: Fixed v. Dynamic; Baseline Case

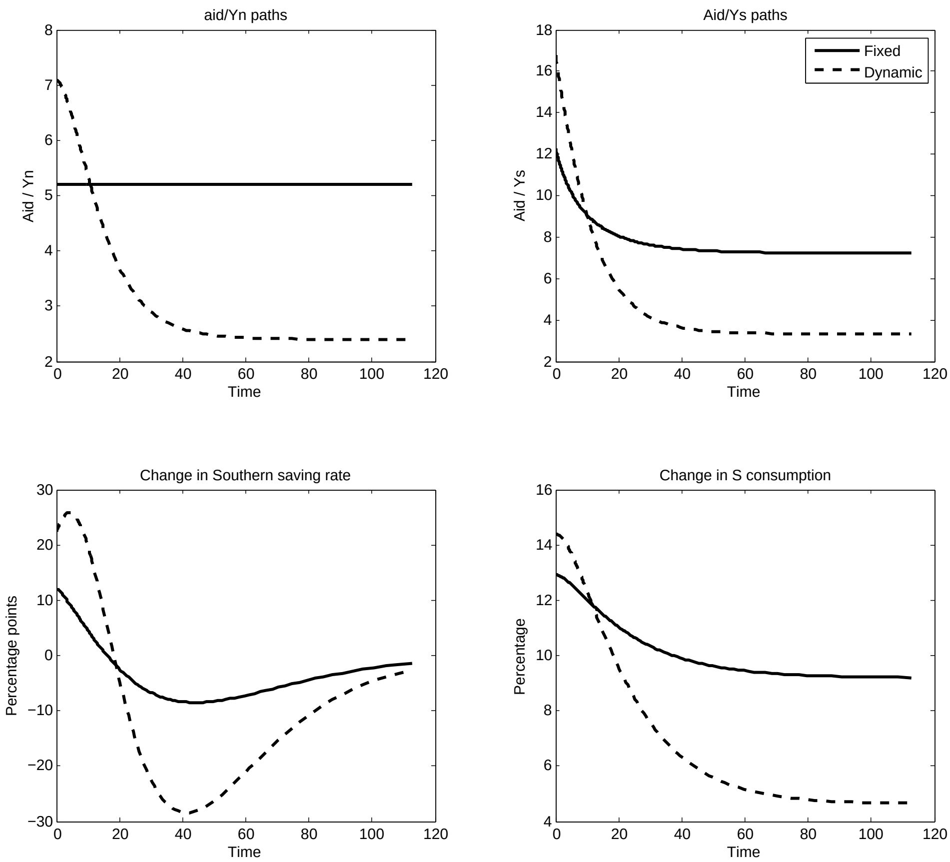

Optimal aid in our baseline parameterization, comparing the unconstrained dynamic solution (dashed line) to the case in which the donor gives a fixed share of its GDP (solid line). The upper two panels show the optimal paths of aid, relative to Northern and Southern GDP. The lower-left panel shows the 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.

Figure 5: Baseline case: Subsistence version

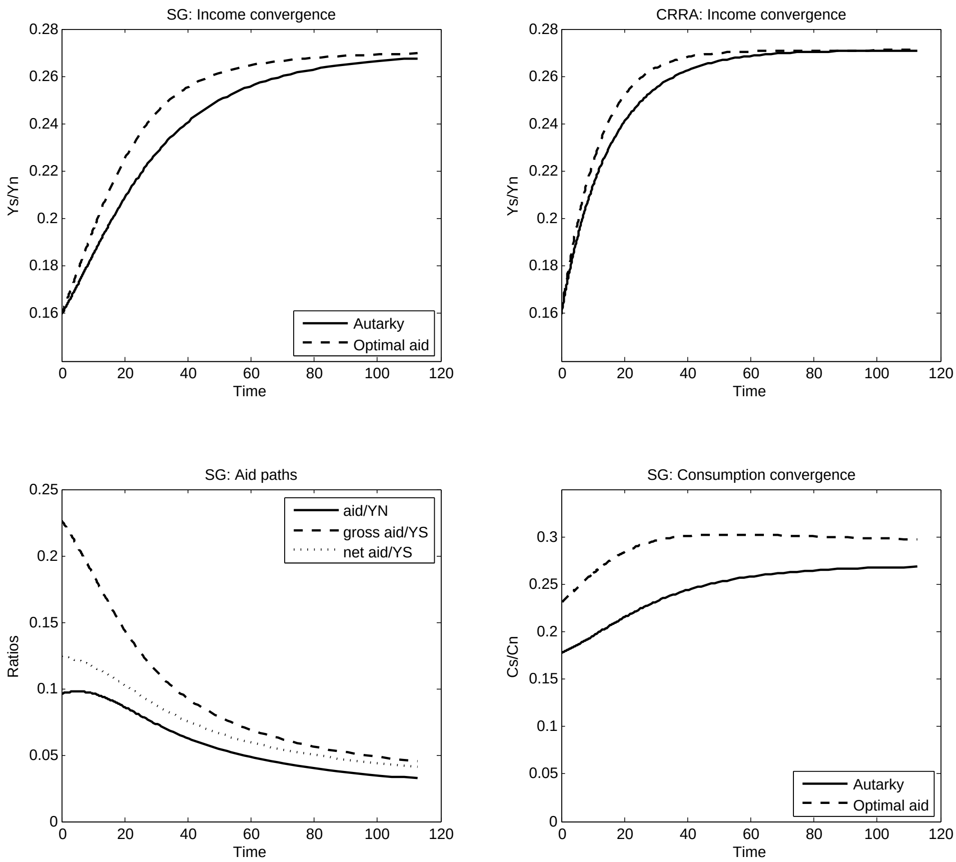

Optimal aid under Stone-Geary preferences. For ease of comparison, the income convergence panel from Figure 1, under CRRA preferences, is shown again here.

Figure 6: Fixed aid: CRRA versus SG

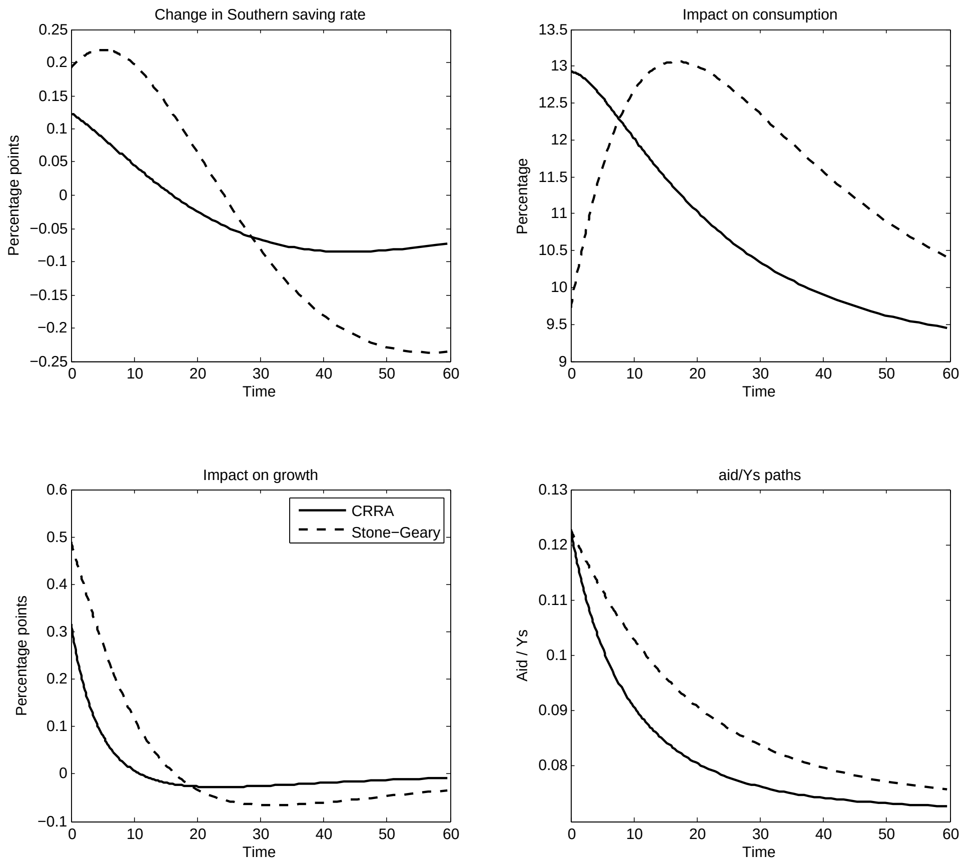

This figure compares outcomes under Stone-Geary (dashed line) and CRRA (solid line) preferences, for the same fixed flow of aid, equal to 5 per cent of Northern GDP. The lower-left panel shows the difference made to the growth rate by aid compared to autarky. Descriptions of other panels as before.

Figure 7: Fixed versus dynamic: Subsistence case

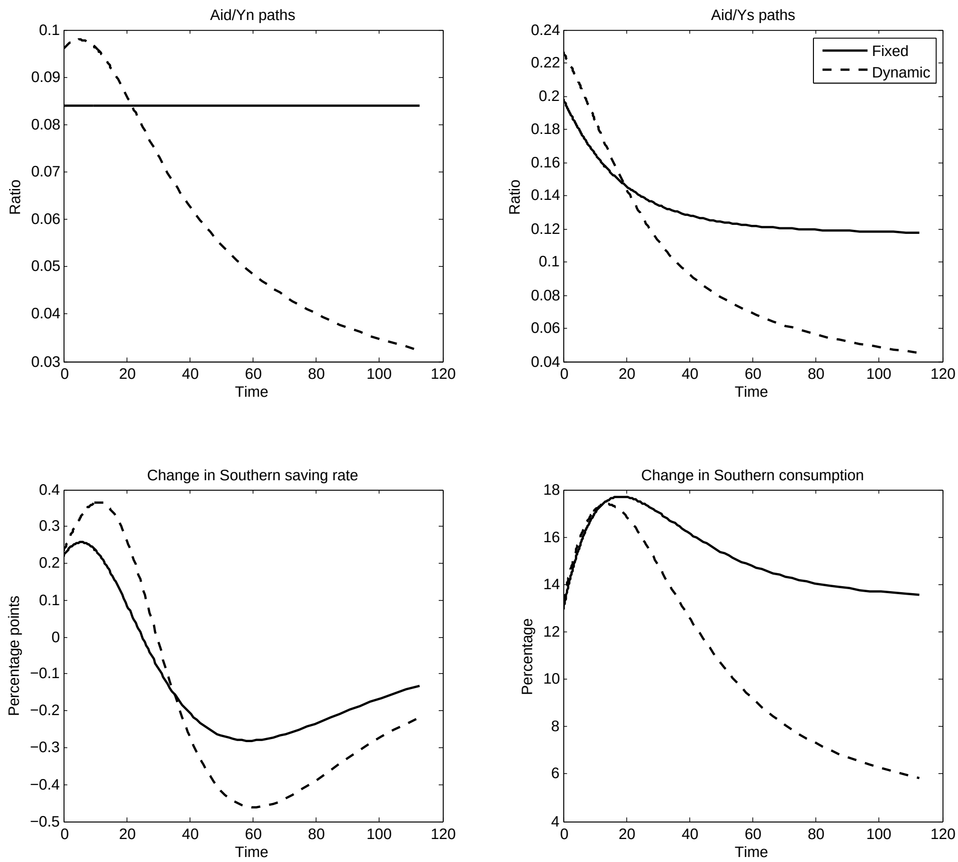

Optimal aid under Stone-Geary prefences, comparing the unconstrained dynamic solution (dashed line) against the case where the donor gives a fixed share of its GDP (solid line). The upper two panels show the paths of aid, relative to Northern and Southern GDP. The lower-left panel shows the change in 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.

Figure 8: Two recipients

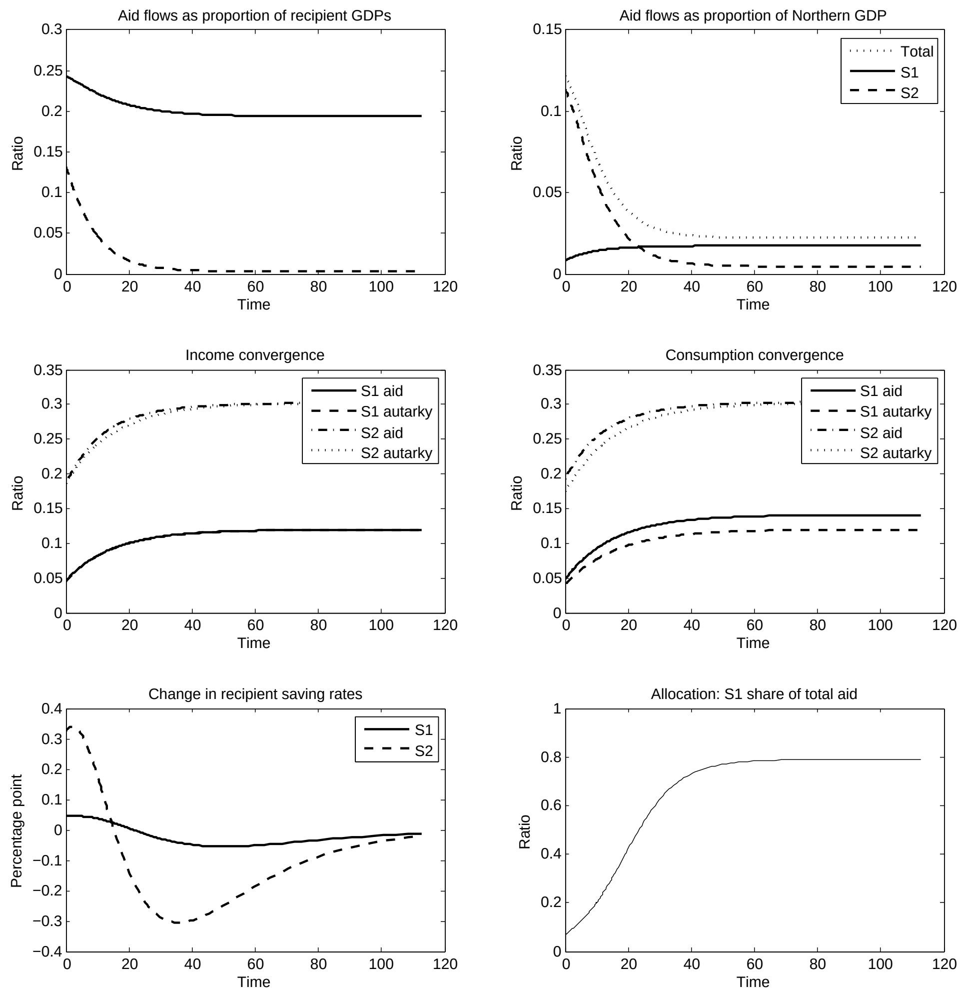

A two-recipient case: $S1$ small and poor, $S2$ middle-income and populous. The middle panels show the ratios of income and consumption in the two recipients to Northern levels, under optimal aid and in autarky; for income in the case of S2, the effect is too small to be seen.

Figure 9: Two recipients: S2 population equal to S1

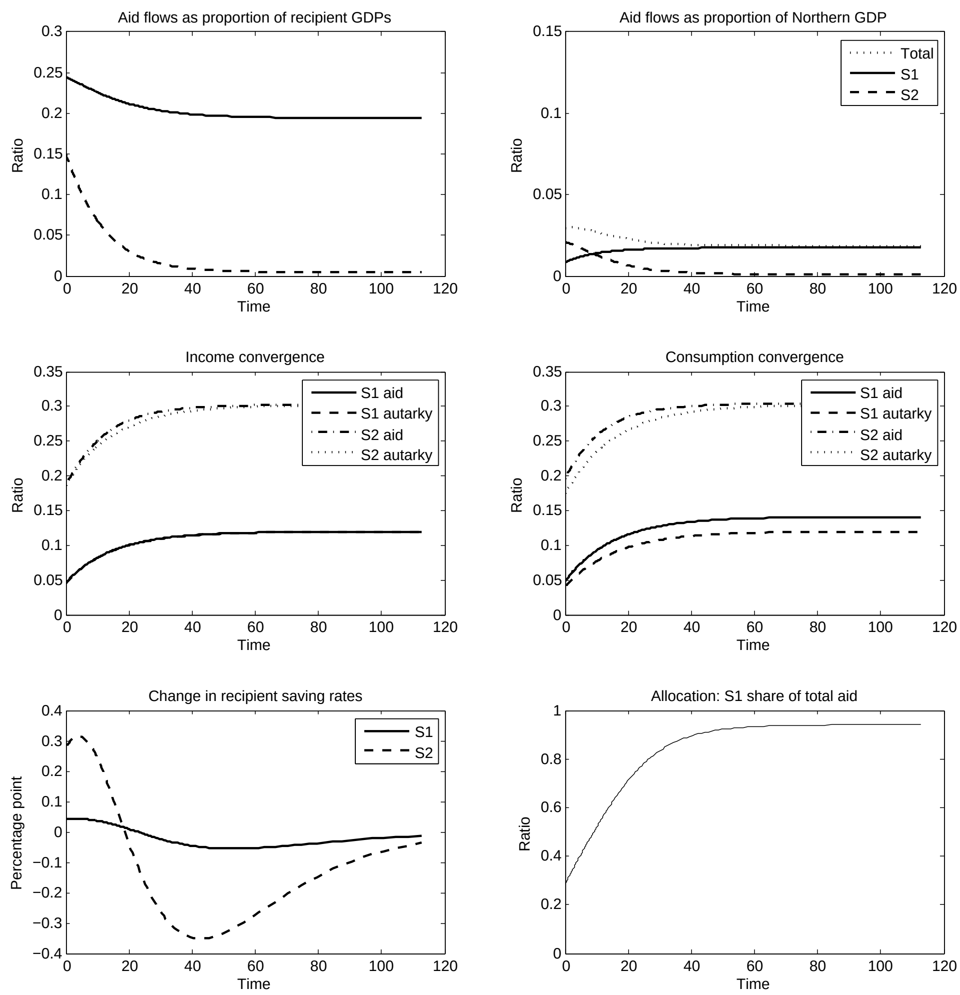

A two-recipient experiment, with the population of $S2$ now reduced to match $S1$.

Figure 10: Two recipients: S2 started in steady-state

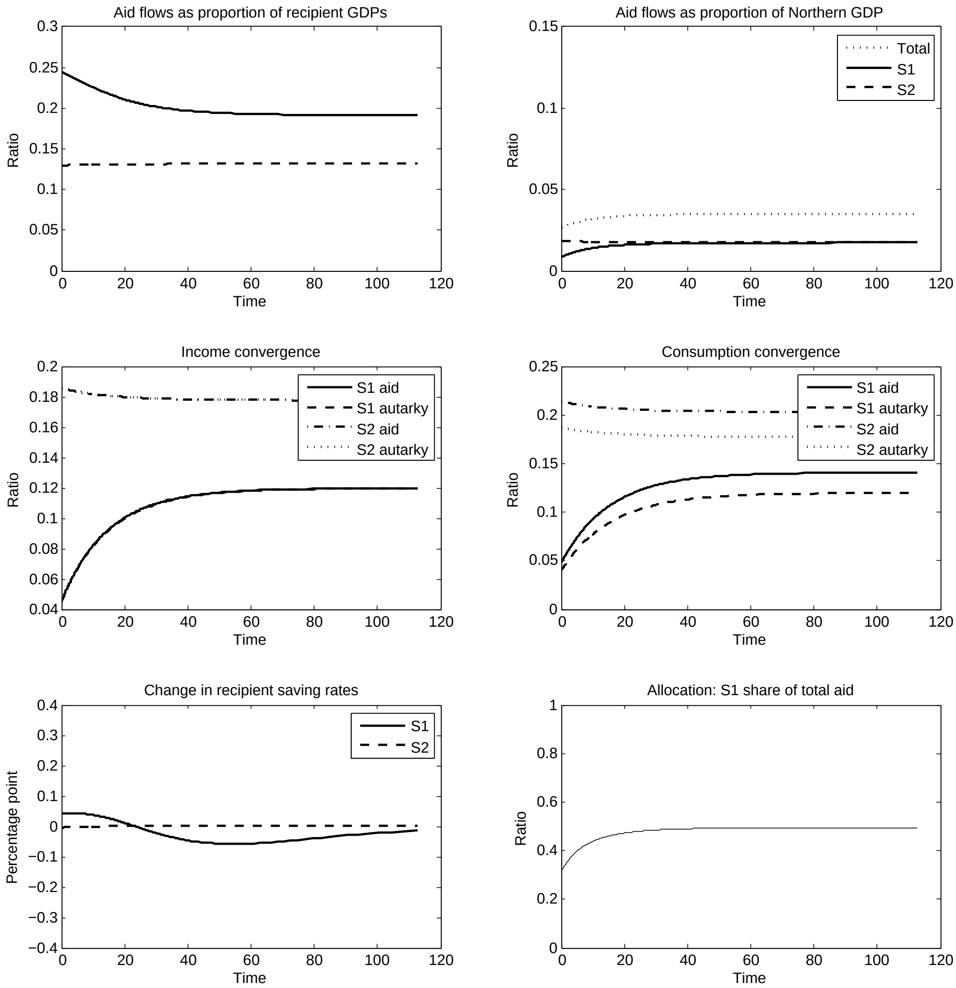

A further two-recipient experiment, starting $S2$ in steady-state. The slight divergence of consumption and income in $S2$ from Northern levels, evident in the middle panels, arises because the North is initially slightly below steady-state.

Figure 11: Two recipients: $v = 1$ for S1

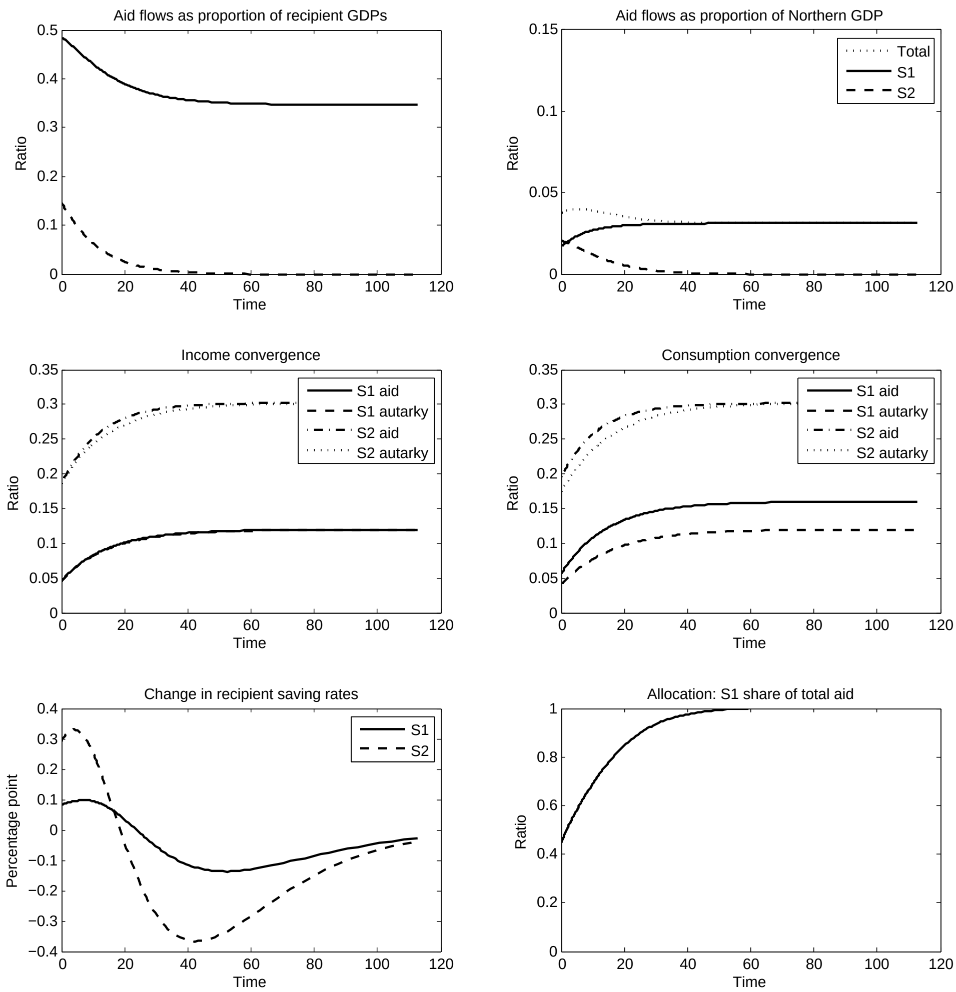

A further two-recipient experiment, relaxing the aid absorption constraint in $S1$.

Figure 12: Southern Planner

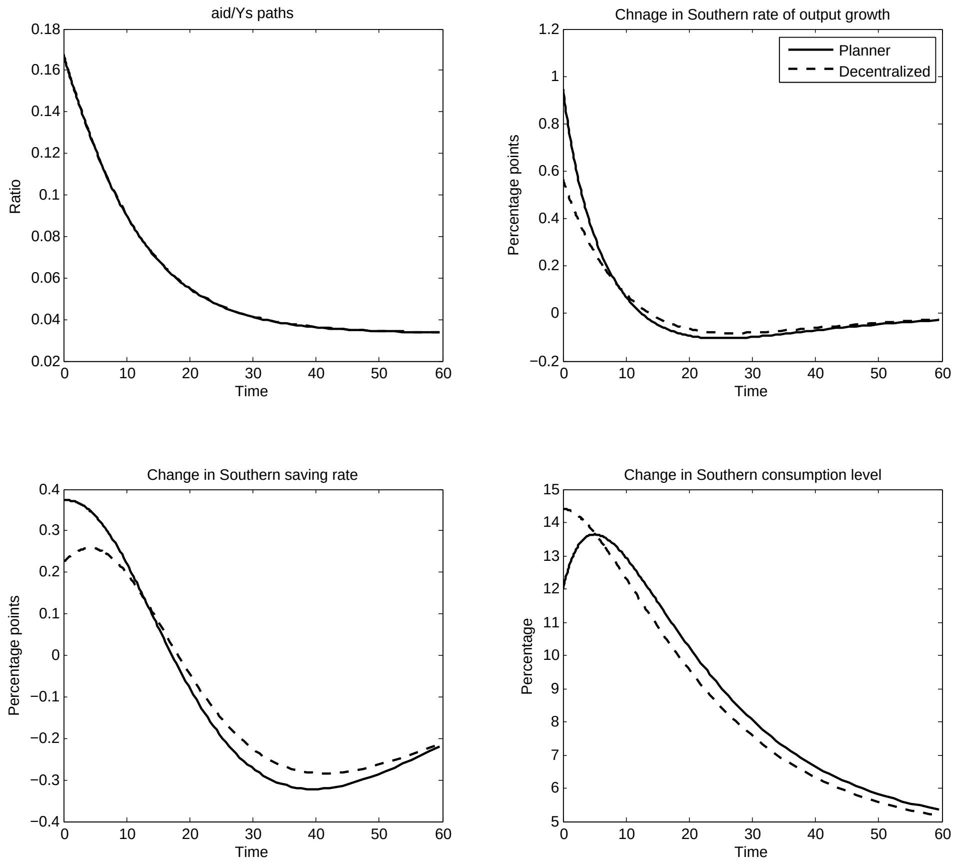

Optimal aid with a Southern planner (solid line) compared to the decentralized case (dashed line). The differences are driven by the aid absorption externality. In the top-left panel, note that the effect on aid generosity is too small to be seen in the figure.