



HAL
open science

Optimal Policies with Heterogeneous Agents: Truncation and Transitions

Xavier Ragot, François Legrand

► **To cite this version:**

Xavier Ragot, François Legrand. Optimal Policies with Heterogeneous Agents: Truncation and Transitions. *Journal of Economic Dynamics and Control*, 2023, 156, 10.1016/j.jedc.2023.104737 . hal-04384031

HAL Id: hal-04384031

<https://hal.science/hal-04384031>

Submitted on 11 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Policies with Heterogeneous Agents: Truncation and Transitions*

François Le Grand[†] Xavier Ragot[‡]

August 31, 2023

Abstract

We compare two approaches in their ability to solve for optimal Ramsey policies in heterogeneous-agent models, considering the optimal provision of a public good. First, the “transition” approach makes the problem tractable by assuming a constant path for the planner’s instruments. Second, the “truncation” approach uses a Lagrangian technique, solving the Ramsey problem of a finite state space model that approximates the full model. The truncation approach is shown to compute quantitatively accurate estimates of the actual values of the planner’s instruments, whereas a time-inconsistency issue is found to affect the transition approach.

Keywords: Heterogeneous agents, optimal Ramsey program, transition approach, truncation approach, aggregate shock.

JEL codes: D31, D52, E21.

*This paper has benefited from the comments of many participants in seminars where the truncation approach has been presented. We are very grateful to two anonymous referees, as well as especially the associate editor, for their insightful comments on earlier versions of this paper. We thank Diego Rodrigues for outstanding research assistance, and Martin Bernstein for helpful remarks.

This work was supported by the French National Research Agency (grant number ANR-20-CE26-0018 IRMAC).

Codes and details about the algorithm can be found at https://github.com/RagotXavier/Truncation_Method_Het.

[†]François LeGrand: Rennes School of Business; 2 Rue Robert d’Arbrissel, 35065 Rennes; francois.le-grand@rennes-sb.com.

[‡]Xavier Ragot: Sciences Po, CNRS and OFCE; 27 Rue Saint-Guillaume, 75007 Paris; xavier.ragot@sciencespo.fr.

1 Introduction

A frontier in the heterogeneous-agent literature is the computation of optimal policies in general equilibrium. In heterogeneous-agent models (or, more precisely, in models with incomplete insurance markets for idiosyncratic risk), redistribution generates new trade-offs. For instance, any policy affecting agents' income modifies the incentives to save for the purpose of self-insurance, and consequently affects future real wages and interest rates. This heterogeneously affects agents' welfare. Solving this so-called Ramsey problem is a very difficult task, which to the best of our knowledge cannot be handled in full generality. To tackle it, the literature has relied on various reformulations of the Ramsey problem that lend themselves to tractable approaches. In this paper, we compare two of these reformulations (see the literature review below for the description of other possibilities): the transition approach and the truncation approach. On the one hand, the *transition approach* reduces the dimensionality of the problem by assuming a constant path for the planner's instruments. This approach consists in finding the constant value of policy instruments that maximize aggregate welfare, computed considering transitions. It has been used in a number of papers (see Conesa et al., 2009; Chang et al., 2018; or Ferriere and Navarro, 2023, among many others). On the other hand, the *truncation approach* proposed in LeGrand and Ragot (2022a) reformulates the general Ramsey problem as the Ramsey problem of a finite state-space model approximating the full model. The solution is then derived using the FOCs of the reformulated Ramsey program. We compare the two approaches in their ability to solve for optimal policy in a heterogeneous-agent model. We consider the provision of a public good, financed by lump-sum taxes, in a model where agents face uninsurable income shocks in a production economy à la Huggett (1993) or Aiyagari (1994). In this problem, the planner's only instrument is the tax level in each period.

The transition approach does not optimize over all possible paths of the policy instrument, but considers only constant paths. However, it improves on the solution of Aiyagari and McGrattan (1998), which maximizes steady-state welfare without accounting for transitions. The approach involves the following steps. Given an initial distribution, one can simulate the transition to the long-run equilibrium distribution for any constant value of the planner's instrument, and then compute the aggregate welfare of the economy along the transition. Then, by iterating over the value of that instrument, one can find the value that maximizes aggregate welfare.

The truncation approach, introduced by LeGrand and Ragot (2022a), consists in solving the Ramsey problem in a limited heterogeneity model that approximates the full heterogeneous-agent model. It obtains a finite state-space representation of the full model by pooling agents according to their idiosyncratic histories, which are truncated at a given length. Agents with the same idiosyncratic history over a given number of consecutive past periods are endowed with the same

Abbreviations: EGM—Endogeneous Grid Method; FOC—first-order condition; IRF—impulse response function; PSID—Panel Study of Income Dynamics; SCF—Survey of Consumer Finances; TFP—total factor productivity.

levels of consumption and savings. Each group of agents can thus be represented by a unique agent, which we call a *truncated history*. The model is then expressed in terms of these truncated histories, and some parameters (labeled ξ s) are introduced to account for the heterogeneity within each truncated history. Indeed, at date t , agents pooled in the same truncated history of length N have the same history from dates $t - N + 1$ to t , but differ in their history before period $t - N$. The ξ s aim to account for this within-history heterogeneity. This paper improves on previous implementations of the truncation approach in that it more effectively takes into account nonlinearities in the model equations. The resulting model, called the *truncated model*, allows one to derive the FOCs of the planner in this finite state-space representation, and to simulate the optimal time-varying path of the planner's instrument.

Our comparison of the two approaches has two main objectives: (i) to identify the mechanisms at stake in the two reformulations, and (ii) to use the transition approach to assess the accuracy of the truncation approach in computing the actual optimal instrument path. First, regarding mechanisms, the two approaches involve a Ramsey planner who maximizes aggregate welfare (i.e., taking into account transitions from an initial distribution to a long-run distribution) and who commits either to a constant policy instrument in the full model (transition approach) or to an entire instrument path in a truncated model (truncation approach). Therefore, the optimal constant transition tax rate should be compared to the average tax rate of the optimal time-varying Ramsey path.¹ The comparison with the average truncation tax rate will thus inform us about the accuracy of the truncation approach and about the approximation implied by the truncated model.

The two methods also yield different long-run values of the optimal instrument. In general, the planner wants the instrument to be time-varying due to a time-inconsistency issue, which is discussed below – this is reflected in the non-constant truncation path. This path, especially in early periods, is affected by the initial distribution of agents, while its long-run value is independent of the initial distribution – Açıkgöz (2018); Açıkgöz et al. (2022) provide a formal proof, and we verify that this holds in our computational exercises. Because the solution provided by the transition approach can be seen as the average of the truncation path, this solution also depends on the initial distribution and at the same time differs from the truncation long-run (or steady-state) tax rate. Thus, the transition approach is useful when the constant instrument is a constraint of the environment (from an economic or a political point of view), but it does not approximate the Ramsey steady state, which is the long-run value of the Ramsey problem with time-varying instruments.

Second, regarding the assessment of the accuracy of the truncation approach, we derive both indirect and direct evidence. For indirect evidence, we proceed as follows. We begin with

¹For the sake of clarity, we will use *transition tax rate* to refer to the constant value of the tax path computed by the transition approach. We will use *truncation (tax) path* to refer to the time-varying optimal tax path computed by the truncation approach. We will also use *long-run truncation tax rate* or *steady-state truncation tax rate* to mean the long-run value of the time-varying optimal tax path computed by the truncation approach.

neutralizing the dependence of the transition approach on the initial distribution. To do so, we jointly find a constant instrument value and an initial distribution, such that aggregate welfare is maximized and the final distribution is identical to the initial distribution. Our algorithm works as follows. We start with a guess for the initial distribution. We then compute the transition value that maximizes aggregate welfare, as well as the corresponding long-run distribution. We use the latter to update the initial distribution, and apply the transition approach again. We keep iterating on the initial distribution, until it coincides with the steady-state distribution. We thus obtain the so-called *fixed-point transition tax rate*. However, and interestingly, the fixed-point transition tax rate still differs from the steady-state truncation tax rate. This comes from the time-inconsistency of Ramsey problems: at time 0, although the distribution is the same as in the steady state, the planner does not have the same commitments with respect to agents' past savings as in the steady state. The planner therefore wants to deviate from the steady-state tax rate and to reoptimize the instrument path.² This implies that, despite identical initial and steady-state distributions, the optimal Ramsey path is time-varying: it deviates from the steady-state tax rate at impact before converging to its long-run value thereafter. In the transition approach, since the path is constrained to remain constant, the constant value partly reflects the planner's willingness at time 0 to deviate from the instrument's steady-state value. Using the truncation approach to compute the optimal path, we find that the transition tax rate is indeed almost equal to the "average" discounted value of the instrument along the truncation path. This shows that the truncation approach yields an accurate solution that does not differ too markedly from the actual Ramsey solution. This is an indirect check of the validity of the truncation approach.

We then proceed to a direct assessment of the accuracy of the truncation approach. We use a generalization of the transition approach to verify that the truncation path actually maximizes the aggregate welfare. To do this, we set the truncation path as an exogenous input and compute the aggregate welfare as in the transition approach (with a time-varying deterministic instrument path instead of a constant one). We also neutralize the effect of the initial distribution, following the same method used to compute the fixed-point transition rate above. We proceed similarly for a number of perturbations of the truncation path, and then verify that changing the truncation path reduces aggregate welfare. Since these aggregate welfare computations only involve the full model (and not the truncated one), this shows that truncating the model to derive the Ramsey solution has a modest quantitative impact. We view this exercise as a direct validity check of the accuracy of the truncation approach.³

Our paper belongs to the literature on optimal policies in heterogeneous-agent models. This

²This time-inconsistency is specific to heterogeneous-agent models. The same problem (financing a public good through lump-sum taxation) is not subject to time-inconsistency in the representative-agent economy, because the financial market clearing condition sets the agent's savings. We have chosen this economic setting to highlight this difference between RA and HA models.

³We also run another comparative exercise involving aggregate shocks in Appendix E.

literature can be divided into three strands. A first strand focuses on simplified environments to identify mechanisms where the equilibrium distribution is simple enough to yield a tractable setup.⁴ Challe (2020), Bilbiie (2021), and Bilbiie and Ragot (2021) solve for optimal monetary policy in an environment where a partial insurance structure implies that the equilibrium features only two consumption levels. Acharya et al. (2023) consider a CARA-normal structure without binding credit constraints to easily aggregate consumption.

A second strand uses numerical techniques to compute the optimal dynamics of the planner’s instruments. Yang (2022) solves for the optimal monetary policy by optimizing on the coefficients of a Taylor rule. McKay and Wolf (2023) focus on optimal rules when the optimal steady-state values of the instruments are known, building on Auclert et al. (2021) with a quadratic approximation of the objective. These strategies compute the optimal dynamics of the instruments around a known steady state. How to find the optimal steady-state value of the instrument in Ramsey problems is still an open question. A solution to this question is to numerically optimize over all possible paths to find the optimal long-run value of the instruments, which is a difficult task due to the need to discount the long-run welfare effect of the instrument (see Dyrda and Pedroni, 2022 for an example and discussion).

The third strand circumvents the difficulties related to numerical techniques by exploiting the planner’s FOCs to solve the Ramsey program when the solution is interior (see Aiyagari, 1995 for an early contribution). A first contribution of this approach is to link the normative analysis to the public finance literature, which makes extensive use of marginal valuations (see Heathcote and Tsujiyama, 2021 for a discussion of optimal policies in two-period models and the discussion in Section 2.6 below). Unfortunately, this leads to additional difficulties in intertemporal models, since heterogeneous-agent models typically involve a continuum of Euler equations. Some papers use continuous-time techniques to solve for optimal policies in heterogeneous-agent models (see Dávila and Schaab, 2022, Nuño and Thomas, 2022, and Smirnov, 2022 for recent examples). Other papers develop new techniques to address the issue of aggregate risk. Bhandari et al. (2021) provide a numerical procedure that assumes that credit constraints are not occasionally binding. They rely on the so-called primal approach, which implies substituting the ratio of marginal utilities for the interest rate. Açıkgöz et al. (2022) (who study fiscal policy) use tools from the Lagrangian approach and follow the entire distribution of Lagrange multipliers on Euler equations, in economies where the credit constraints are occasionally binding—which is often the quantitatively relevant case in quantitative models.

LeGrand and Ragot (2022a) (studying optimal unemployment insurance) use the Lagrangian approach but in the context of a finite state-space model, which is derived with the truncation approach. LeGrand and Ragot (2022b) propose a refinement of the truncation approach that allows for consideration of truncated histories of different lengths, but which focuses solely on

⁴We do not review here the large literature on optimal policy with a complete insurance-market but ex ante heterogeneous agents; see, among others, Bassetto (2014).

simulating the model in the presence of aggregate shocks, not on computing optimal policies. LeGrand et al. (2022) study the joint optimal monetary-fiscal policy, and use the truncation approach only as a tool to solve the model. Our current contribution is to verify the accuracy of the truncation approach, using insights from the transition approach. In addition, to improve the accuracy of the truncation approach, we enhance our aggregation procedure by introducing additional history-dependent parameters. This allows for a precise aggregation of the model's non-linear equations (see Section 4.1 below).

The remainder of this paper is organized as follows. Section 2 presents the environment and the Ramsey problem. Section 3 describes how the transition approach can be used to compute an optimal constant value for policy instruments. It also proposes an improvement over current solutions to neutralize the impact of the choice of initial distributions. Section 4 details the computation of the Ramsey solution using the truncation approach. Section 5 contains a numerical exercise quantifying the differences along several dimensions between the two approaches considered in the paper. Section 6 concludes.

2 The environment

We consider an environment similar to that of Den Haan (2010), which is a heterogeneous-agent economy with aggregate productivity risk and exogenous labor supply. The main twist is the introduction of a public good, whose provision enters into private utility. This public good is financed by a benevolent government through a lump-sum tax raised on all agents. The Ramsey problem we study concerns the optimal provision of this public good. We consider a discrete-time economy populated by a continuum of agents of size 1. Agents are distributed according to a non-atomic measure ℓ on a set I : $\ell(I) = 1$. We follow Green (1994) and assume that the law of large numbers holds. We focus the paper on the no-aggregate shock case and introduce aggregate shocks as an extension in Appendix E.

2.1 Risk

Agents are affected by an idiosyncratic labor productivity shock y taking values in a finite set \mathcal{Y} . We assume that the individual productivity process follows a first-order Markov chain with constant transition probabilities $(\Pi_{yy'})_{y,y' \in \mathcal{Y}}$. The size of the agents' population with productivity y is constant and denoted S_y . It is defined through the recursions, $S_y := \sum_{\tilde{y} \in \mathcal{Y}} \Pi_{\tilde{y}y} S_{\tilde{y}}$, holding for all $y \in \mathcal{Y}$, with $\sum_{y \in \mathcal{Y}} S_y = 1$ because the size of the population is one. Finally, an individual history of productivity shocks up to date t is denoted $y^t = \{y_0, \dots, y_t\} \in \mathcal{Y}^{t+1}$, and θ_t is the measure of such histories.

2.2 Preferences

In each period, there are two goods in the economy: a private consumption good and a public consumption good. Households are expected-utility maximizers, who rank streams of private consumption $(c_t)_{t \geq 0}$ and of public consumption $(G_t)_{t \geq 0}$ according to a time-separable intertemporal utility function equal to $\sum_{t=0}^{\infty} \beta^t (u(c_t) + v(G_t))$, where $\beta \in (0, 1)$ is a constant discount factor, and $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $v : \mathbb{R}_+ \rightarrow \mathbb{R}$ are instantaneous utility functions reflecting separable preferences over private and public consumption, respectively. As is standard, we assume that u and v are twice continuously differentiable, increasing, and concave, with $u'(0) = v'(0) = \infty$.

2.3 Production

A standard profit-maximizing representative firm produces a generic output good, which can then be used either as a private consumption good, or as a public consumption good. At any date t , the firm production function combines labor L_t and capital K_{t-1} , which needs to be installed one period in advance, to produce Y_t units of the consumption good. Individual labor supply is fixed and normalized to 1, and so aggregate labor supply is constant and equal to $\bar{L} := \int_i y_{i,t} \ell(di)$ efficient units. The production function is assumed to be of the Cobb-Douglas type, featuring constant returns-to-scale with parameter $\alpha \in (0, 1)$, and capital depreciation at rate $\delta \in (0, 1)$. TFP is normalized to 1 and the production function is:

$$Y_t = F(K_{t-1}, \bar{L}) = K_{t-1}^\alpha \bar{L}^{1-\alpha} - \delta K_{t-1}. \quad (1)$$

The firm rents labor and capital at respective factor prices w_t and r_t . The profit maximization conditions of the firm imply the following expression for factor prices:

$$w_t = F_L(K_{t-1}, \bar{L}) \text{ and } r_t = F_K(K_{t-1}, \bar{L}). \quad (2)$$

2.4 Government

In each period t , the government finances an endogenous public good expenditure G_t through a lump-sum transfer T_t . In the absence of public debt, the government budget must be balanced in each period:

$$T_t = G_t. \quad (3)$$

We choose to abstract from more complex financing schemes in order to compare the two approaches in a straightforward environment.⁵ As a robustness check, we have also simulated the model with a proportional tax on labor, instead of a lump-sum tax. The comparison of the two approaches generates the same conclusions as with the lump-sum tax (see Section 5.7 below).

⁵See Den Haan (2010) for an early similar strategy for models with aggregate shocks.

2.5 Agents' program, resource constraints, and equilibrium definition

Agents can save in capital shares paying the real interest rate r_t between dates $t - 1$ and t . They face credit constraints, and their savings must remain greater than an exogenous threshold normalized to 0.⁶ In the initial period, each agent i is endowed with initial wealth a_{-1}^i and initial productivity status y_0^i , jointly drawn from an initial distribution Λ_0 , defined over $[-\bar{a}; \infty) \times \mathcal{Y}$. Formally, given this initial endowment and given the stream of public spending $(G_t)_{t \geq 0}$, the agent's program can be expressed as:

$$\max_{(c_t^i, a_t^i)_{t \geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(u(c_t^i) + v(G_t) \right), \quad (4)$$

$$c_t^i + a_t^i = (1 + r_t)a_{t-1}^i + w_t y_t^i - T_t, \quad (5)$$

$$a_t^i \geq 0, c_t^i > 0, a_{-1}^i \text{ given}, \quad (6)$$

where \mathbb{E}_0 is an expectation operator over idiosyncratic shocks. In the initial period, the agent chooses their consumption path $(c_t^i)_{t \geq 0}$ and their saving plan $(a_t^i)_{t \geq 0}$ to maximize their intertemporal utility (4), subject to the budget constraint (5) and the borrowing limit (6).

The solution of the previous program is a set of policy rules $c_t : \mathcal{Y}^t \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and $a_t : \mathcal{Y}^t \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that determine consumption and saving decisions as functions of the idiosyncratic history y_t^i of agent i , and their initial wealth a_{-1}^i . However, to simplify the notation, we will simply write c_t^i and a_t^i (instead of $c_t(y_t^i, a_{-1}^i)$ and $a_t(y_t^i, a_{-1}^i)$). As stated in the next remark, we use the same notation for all variables.

Remark 1 (Simplifying notation) *An agent has an idiosyncratic history y_t^i and initial wealth a_{-1}^i ; X_t^i denotes the realization in state (y_t^i, a_{-1}^i) of any random variable $X_t : \mathcal{Y}^t \times \mathbb{R}^+ \rightarrow \mathbb{R}$.*

A consequence of Remark 1 is that the aggregation of variable X_t in period t over the whole agent population is written as $\int_i X_t^i \ell(di)$, instead of the more involved explicit notation $\int_{a_{-1} \in \mathbb{R}^+} \sum_{y^t \in \mathcal{Y}^t} \theta_t(y^t) X(y^t, a_{-1}) d\Lambda_0(a_{-1}, y_0)$.

Taking advantage of this notation, we denote by $\beta^t \nu_t^i$ the Lagrange multiplier on the agent- i credit constraint. The agent's Euler equation can then be written as:

$$u'(c_t^i) = \beta \mathbb{E}_t \left[(1 + r_{t+1}) u'(c_{t+1}^i) \right] + \nu_t^i. \quad (7)$$

Financial market clearing conditions and the economy-wide resource constraint are:

$$\int_i a_t^i \ell(di) = K_t, \text{ and } \int_i c_t^i \ell(di) + G_t + K_t = Y_t + K_{t-1}. \quad (8)$$

⁶The credit constraint can be arbitrarily set to 0 without loss of generality up to some monotonic transformation of variables (including the wage process). See Açıkgöz (2018, Section 4.1). For a given wage process, changing the credit constraint however affects the economy (through accumulated capital and interest rate).

We can now state our market equilibrium definition.

Definition 1 (Sequential equilibrium) *A sequential competitive equilibrium is a collection of individual plans $(c_t^i, a_t^i, \nu_t^i)_{t \geq 0, i \in \mathcal{I}}$, of aggregate quantities $(K_t, Y_t)_{t \geq 0}$, of price processes $(w_t, r_t)_{t \geq 0}$, and of fiscal policy $(G_t, T_t)_{t \geq 0}$, such that, for an initial wealth and productivity distribution Λ_0 , and for an initial value of capital stock verifying $K_{-1} = \int_{a_{-1} \in \mathbb{R}^+} \sum_{y_0 \in \mathcal{Y}} d\Lambda_0(a_{-1}, y_0)$, we have:*

1. *Given prices and fiscal policy, the functions $(c_t^i, l_t^i, \nu_t^i)_{t \geq 0, i \in \mathcal{I}}$ solve the agent's optimization program in equations (4)–(6);*
2. *Financial and goods markets clear at all dates—for any $t \geq 0$, equations (8) hold;*
3. *The government budget is balanced at all dates—equation (3) holds for all $t \geq 0$;*
4. *Factor prices $(w_t, r_t)_{t \geq 0}$ are consistent with condition (2).*

2.6 The Ramsey problem

The Ramsey problem consists in selecting a fiscal policy that corresponds to a competitive equilibrium with the highest aggregate welfare. Regarding the latter, we opt for the standard ex-ante additive criterion, also known as the utilitarian social welfare function, which attributes an identical weight to all agents. Formally, the aggregate welfare criterion can be expressed as:⁷

$$W_0 := \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \int_i (u(c_t^i) + v(G_t)) \ell(di) \right], \quad (9)$$

which depends on the public spending path $(G_t)_{t \geq 0}$ and on the consumption paths of all agents, $(c_t)_{t \geq 0, i \in \mathcal{I}}$. Other social welfare functions could be considered—and the truncation approach could be used to solve them—but we restrict our attention to this useful benchmark, which has been used in a number of heterogeneous-agent papers since the seminal study of Aiyagari (1995).

Ramsey allocation, reoptimization shock, and time-inconsistency. We now formalize our definition of the Ramsey allocation.

Definition 2 (Ramsey allocation) *Given an initial distribution Λ_0 over initial wealth and productivity levels:*

1. *A Ramsey allocation is a competitive equilibrium in the sense of Definition 1 that maximizes the aggregate welfare W_0 of equation (9) over the set of competitive equilibria;*

⁷In the sequential representation, the explicit expression is $W_0 := \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \int_{a_{-1} \in \mathbb{R}^+} \sum_{y^t \in \mathcal{Y}^t} \theta_t(y^t) (u(c_t(y^t, a_{-1}, Z^t)) + v(G_t)) d\Lambda_0(a_{-1}, y_0)$.

2. The steady-state Ramsey allocation is characterized by the long-run values, when they exist, of the distribution over savings and productivity levels, denoted Λ_∞^{opt} , and of the instrument, denoted T_∞^{opt} .

Definition 2 defines the steady-state Ramsey allocation as the limit of the Ramsey allocation. Indeed, the Ramsey allocation characterizes a path for the distributions (over savings and productivity levels) and for the fiscal instrument, which both depend on the initial distribution Λ_0 . To highlight this relationship, the paths of the distribution and of the fiscal instrument will be denoted $(\Lambda_t^{opt}(\Lambda_0))_t$ and $(T_t^{opt}(\Lambda_0))_t$, respectively. Conversely, the steady-state Ramsey allocation does not depend on initial conditions, as we will verify numerically, and as discussed by Açıkgöz (2018).

Characterizing the path of the instruments and time-inconsistency in heterogeneous-agent models. Definition 2 can be formalized as the outcome of an optimization program. Using the governmental budget constraint (3) to substitute T_t for G_t , the Ramsey program can be written as follows:

$$\max_{(w_t, r_t, T_t, K_t, (a_t^i, c_t^i, \nu_t^i)_{i \in I})_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \int_i (u(c_t^i) + v(T_t)) \ell(di) \right], \quad (10)$$

$$\forall i \in I, a_t^i + c_t^i = (1 + r_t) a_{t-1}^i + w_t y_t^i - T_t, \quad (11)$$

$$u'(c_t^i) = \beta \mathbb{E}_t \left[(1 + r_{t+1}) u'(c_{t+1}^i) \right] + \nu_t^i, \quad (12)$$

$$a_t^i \geq 0, \nu_t^i a_t^i = 0, \nu_t^i \geq 0, c_t^i \geq 0, \quad (13)$$

$$K_t = \int_i a_t^i \ell(di), \quad (14)$$

$$r_t = F_K(K_{t-1}, \bar{L}), w_t = F_L(K_{t-1}, \bar{L}), \quad (15)$$

where $(a_{-1}^i, y_0^i)_i$ is given by the initial distribution Λ_0 . Equation (10) is the planner's objective (9). Equations (11)–(15) are the planner's constraints, which guarantee that the chosen allocation is picked from among the competitive equilibria of Definition 1. Equations (11)–(13) are individual constraints: the budget constraint, the Euler equation, and the positivity and credit constraints, respectively. In the problem under consideration, the consumption positivity constraint should not be neglected because the lump-sum tax is the sole source of financing. This means that the consumption of poorer agents may become negative for large taxes. Equations (14) and (15) are economy-wide constraints, regarding financial market clearing and factor price definitions.

The trade-off faced by the planner in the Ramsey program (10)–(15) is rather straightforward. The planner can increase the provision of the public good at the cost of a higher tax that reduces the consumption of private goods. This higher tax has a heterogeneous effect on agents, because they have different wealth levels and different incomes. Thus, the higher tax heterogeneously

affects agents' saving decisions, which in turn modify the dynamics of the capital stock, the real wage, and the real interest rate. These general equilibrium effects, combined with the redistribution motives, make the Ramsey problem difficult to solve.

To understand the time-inconsistency of the Ramsey problem, we now use the FOCs of the Ramsey program (10)–(15). The quantity $\beta^t \lambda_t^i$ denotes the Lagrange multiplier on the agent- i Euler equation.⁸ To simplify the analysis, we introduce the concept of *social valuation of liquidity for agent i* , denoted ψ_t^i and formally defined as:

$$\psi_t^i := \underbrace{u'(c_t^i)}_{\text{direct effect}} - \underbrace{u''(c_t^i) \left(\lambda_t^i - (1 + r_t) \lambda_{t-1}^i \right)}_{\text{effect on savings incentives}}, \quad (16)$$

which can be seen as the equivalent for the planner of the marginal utility of consumption. Indeed, it measures, from the planner's perspective, the value of one extra unit of consumption for agent i at date t . ψ_t^i includes three terms. The first one, $u'(c_t^i)$, is the private value, for agent i , of consuming one extra unit at date t . The two additional terms reflect the planner's valuation of the savings distortions induced by this extra consumption. The extra consumption at date t affects agents' savings incentives between dates $t - 1$ and t , as well as between dates t and $t + 1$. Since the planner internalizes savings constraints via individual Euler equations, the impact related to savings between $t - 1$ and t is proportional to the shadow cost of the Euler equation between these two dates, which is λ_{t-1}^i . Similarly, the impact related to savings between t and $t + 1$ is proportional to λ_t^i , which is the shadow cost of the Euler equation between t and $t + 1$. Finally, the signs in front of λ_{t-1}^i and λ_t^i are opposite as in the first case the extra consumption means extra future consumption, while in the other it means extra current consumption.⁹

We derive the FOCs in Appendix A. We here discuss the results. The FOC regarding the lump-sum tax T_t can be expressed as follows:

$$v'(T_t) = \int_i \psi_t^i \ell(di). \quad (17)$$

The interpretation of this equation is straightforward. The marginal benefit of increasing the tax (and hence public spending) is $v'(T_t)$ and is common to all agents. This marginal benefit is set equal to the marginal cost, which amounts to taxing one unit of private consumption (valued as ψ_t^i for agent i by the planner) for all agents in the population (hence the integral over i).

Condition (17) can also be rewritten as:

$$v'(T_t) = \int_i u'(c_t^i) \ell(di) - \int_i u''(c_t^i) \left(\lambda_t^i - (1 + r_t) \lambda_{t-1}^i \right) \ell(di), \quad (18)$$

⁸For brevity, we refer to LeGrand and Ragot (2022a), who prove that the Lagrange multipliers can be used with occasionally-binding credit constraints, by taking limits of penalty functions.

⁹In general we cannot predict the sign of λ_t^i , which depends on whether the planner perceives that agent i is over or under saving (see LeGrand and Ragot, 2022a for an illustration).

which shows a time-inconsistency problem due to the presence of current and past values of the Lagrange multiplier on the right-hand side. In an incomplete-market economy, the Ramsey planner generally cannot restore the first-best allocation and must thus use their instruments to close the gap with the first-best allocation. This involves accounting for private saving incentives, which are distorted by the lump-sum tax in a heterogeneous way (because of pervasive heterogeneity among agents). In particular, yesterday's savings choices were affected by the tax amount at that time – hence the presence of λ_{t-1}^i in (18). Should the planner be given the opportunity to re-optimize at some date t_0 (i.e., to set $\lambda_{t_0-1}^i = 0$ for all i , thereby corresponding to the absence of commitment to past promises), they would modify the tax path from date t_0 onwards and be time-inconsistent.¹⁰

Because of this time-inconsistency, there is no reason for the optimal Ramsey tax path to be constant, even when the initial distribution is set to its steady-state value. Indeed, at date 0, there are no past commitments and all $(\lambda_{-1}^i)_i$ are set to zero. However, in the long run ($t \rightarrow \infty$), past promises still matter and past Lagrange multipliers in condition (18) differ from zero. The planner thus faces a different trade-off at date 0 than in the long-run, even though wealth distributions are assumed to be identical. The planner is thus likely to reoptimize at date 0 such that the tax level at impact deviates from the steady-state tax.

In the complete-market environment, there is only one Euler equation for the representative agent, for which the Lagrange multiplier is $\lambda_t^{CM} = 0$, as we formally show in Appendix B. This is because in the absence of heterogeneity, the savings of the representative agent are equal to the aggregate capital stock and hence not influenced by the lump-sum tax. The FOC of the Ramsey planner is:

$$v'(T_t^{CM}) = u'(c_t^{CM}),$$

In this case, the Ramsey problem is time-consistent, because no past value of Lagrange multiplier affects the planner's decision. By the same token, it can be expected that in a complete market economy, should the initial distribution be set to the steady-state distribution, the optimal tax path will remain constant in the absence of aggregate shocks.

3 The transition approach

The transition approach consists in reformulating the Ramsey problem by restricting the planner's instruments to constant paths and then by finding the constant instrument values that maximize the aggregate welfare while accounting for transitions (see Conesa et al. (2009) or Chang et al. (2018), among many others). The model simulation takes as inputs a constant tax rate T and an initial distribution Λ_0 (over wealth and productivity) and generates a long-run distribution through simulation, denoted $\Lambda_\infty(T, \Lambda_0)$. This long-run distribution depends on the initial

¹⁰Setting the past value of Lagrange multipliers to 0 is sometimes called a reoptimization shock in the literature (Debortoli et al., 2014).

distribution and the tax path (hence the notation). We now formalize the definition of the transition approach.

Definition 3 *The optimal transition tax rate is the constant fiscal policy T^{cst} , such that, given an initial distribution $\Lambda_0 : [0, \infty) \times \mathcal{Y} \rightarrow \mathbb{R}^+$ over wealth and productivity:*

1. *When the tax path is constant and set to T^{cst} , the Bewley model with the initial distribution Λ_0 converges to a long-run distribution denoted $\Lambda_\infty(T^{cst}, \Lambda_0)$;*
2. *The lump-sum tax T^{cst} maximizes the computed period-0 aggregate welfare, accounting for the transitions when the agents' distribution evolves from Λ_0 at $t = 0$ to $\Lambda_\infty(T^{cst}, \Lambda_0)$ in the long run ($t \rightarrow \infty$).*

Formally, the transition approach consists of finding the constant tax path with level T^{cst} that maximizes aggregate welfare for a given Λ_0 . The algorithm for this solution can be summarized as follows.

Algorithm 1 (The transition approach) *We consider as given an initial distribution Λ_0 . The transition approach involves the following steps:*

1. *Set an initial guess for the constant lump-sum tax path T ;*
2. *Solve for the steady-state distribution $\Lambda_\infty(T, \Lambda_0)$;*
3. *Compute the transition and the welfare during the transition of the economy from the initial distribution Λ_0 toward $\Lambda_\infty(T, \Lambda_0)$;*
4. *Update T and start again at Step 2 until the computed welfare is maximal.*

The resulting tax rate, denoted by T^{cst} , is the optimal transition tax rate.

How does the transition approach compare to the Ramsey program (10)–(15)? The transition and the truncation approaches select a competitive equilibrium (in the sense of Definition 1) that maximizes aggregate welfare (9). The Ramsey program (10)–(15) of the truncation approach allows for a time-varying tax path, while the transition approach selects a constant tax path.¹¹ Formally, the transition approach solves a problem that could be written as a Ramsey program with one additional condition, stating that the tax path is constant: $T_t = T$ for all t . In particular, the transition approach also features time-inconsistency. Indeed, should the truncation planner be given the opportunity to choose a new (constant) tax path at $t = 1$, then they would choose a different tax path than at $t = 0$ – for the same reason as in the Ramsey program.

¹¹Note that in both approaches, the initial distribution can be freely calibrated.

However, imposing a constant path is a non-trivial restriction. Loosely speaking – and as we will see later in our numerical applications – this means that the transition approach has to trade-off short-term aspects with long-term ones. In other words, the initial distribution will affect the outcome of the transition approach. For instance, starting from an economy with a very low initial capital stock (close to 0) implies a very low T , because the tax base in the initial periods is close to 0. Alternatively, if the initial capital stock is very large, then the planner can set a higher value for its instrument and speed the convergence toward a lower optimal capital stock. Consequently, we denote this optimal tax as $T^{cst}(\Lambda_0)$ to highlight the dependence on the initial distribution.

Addressing the dependence in the initial condition. To neutralize the effect of the initial distribution, we modify the previous computation by iterating on the initial distribution, until initial and long-run distributions coincide for the chosen tax level. Formally, we iterate on the initial distribution Λ_0 and compute a new optimal instrument value $T^{cst}(\Lambda_0)$ until the long-run distribution equals the initial distribution: $\Lambda_\infty(T^{cst}(\Lambda_0), \Lambda_0) = \Lambda_0$. When convergence has been reached, we obtain a tax and a long-run distribution, denoted (Λ_∞^c, T^c) , which verify: $T^c = T^{cst}(\Lambda_\infty^c)$ and $\Lambda_\infty^c = \Lambda_\infty(T^{cst}(\Lambda_\infty^c), \Lambda_\infty^c)$. Formally, we are looking for the fixed point of the mapping $\Lambda \mapsto \Lambda_\infty(T^{cst}(\Lambda), \Lambda)$ defined for distributions over initial wealth and initial productivity status. Hence, this procedure will be called the *fixed-point transition approach* and the resulting tax rate, the *fixed-point transition tax rate*. To our knowledge, this study is the first to propose this procedure. We formalize its definition below.¹²

Definition 4 *The fixed-point transition tax rate is the constant fiscal policy T^c such that:*

1. *When the fiscal policy is constant and set to T^c , the Bewley model with the initial distribution $\Lambda_0 = \Lambda_\infty^c$ converges to the same long-run distribution $\Lambda_\infty(T^c, \Lambda_\infty^c) = \Lambda_\infty^c$;*
2. *The tax rate T^c maximizes the period-0 aggregate welfare computed when accounting for the transitions when the agents' distribution evolves from $\Lambda_0 = \Lambda_\infty^c$ at $t = 0$ to Λ_∞^c in the long run.*

The fixed-point transition tax rate T^c is such that: (i) the initial and the long-run distributions coincide when the fiscal policy is constant and set to T^c ; and (ii) the aggregate welfare with transitions is maximal when the fiscal policy is set to T^c .

We formalize the computation of T^c in the algorithm below.

Algorithm 2 (The fixed-point transition approach) *The fixed-point transition approach involves the following steps:*

¹²To avoid confusion, the “transition tax rate” refers to Definition 3, where the initial distribution influences the outcome. The “fixed-point transition tax rate” refers to Definition 4, where the initial and long-run distributions are identical. As will be defined below, the “optimal truncation tax rate” refers to the long-run optimal tax rate computed using the truncation approach of Algorithm 3.

1. Choose an initial guess for the tax rate T ;
2. Choose an initial guess for the initial distribution Λ_0 :
 - (a) Compute the long-run distribution $\Lambda_\infty(T, \Lambda_0)$;
 - (b) If the initial and long-run distributions coincide (i.e., $\Lambda_\infty(T, \Lambda_0) = \Lambda_0$), then proceed to Step 3. Otherwise, update the initial distribution, $\Lambda_0 \leftarrow \Lambda_\infty(T, \Lambda_0)$, and start at 2(a).
3. Compute the aggregate welfare during the transitions of the economy from the initial distribution Λ_0 toward $\Lambda_\infty(T, \Lambda_0) = \Lambda_0$;
4. Update T and start again at Step 2 until the computed welfare at Step 3 is maximal.

The resulting tax rate, denoted by T^c , is the optimal fixed-point transition tax rate.

Definition 4 and Algorithm 2 neutralize the influence of the initial distribution on the optimal transition tax rate. However, this still does not mean that the fixed-point transition instrument T^c is equal to the Ramsey steady-state value. Indeed, and as explained in Section 2.6 above, the Ramsey planner is time-inconsistent. Even though the initial and long-term distributions coincide with each other, planner will take advantage of the absence of initial commitment in the first period to deviate at impact from the long-term tax rate and choose a non-constant tax path. By contrast, in the case of the fixed-point transition approach, the planner must implement a constant tax path by construction. Because of the absence of short-run deviations, the transition tax must balance short-run and long-run aspects. The fixed-point transition tax rate will again be close to the average Ramsey tax path and thus differ from the Ramsey steady-state value. The gap between the two tax rates will reflect the severity of the planner's time-inconsistency. It will be large when the time-0 deviation due to time-inconsistency is large.

Finally, it should be observed that Definition 4 and Algorithm 2 provide a solution that differs from the optimal tax maximizing the steady state welfare (as in Aiyagari and McGrattan, 1998 for instance). Indeed, the objective solved by the planner is not the same in the two cases: in one case it is the intertemporal aggregate welfare, while in the other it is the steady-state aggregate welfare. We illustrate this point in our quantitative exercise.

4 The truncation approach

We now present the truncation approach of LeGrand and Ragot (2022a), which consists in solving the Ramsey program of a finite state-space model approximating the full model. The finite state-space model is called the *truncated model*.

Because the present study improves on this approach, we present it in several steps:

1. We first provide the construction of the truncated model in theory and in practice in Section 4.1;
2. We explain how to solve for steady-state Ramsey policies in the truncated model in Section 4.2;
3. Finally, we show how to compute the optimal Ramsey solution in the truncated model in Section 4.3.

Before turning to the details, we start with an intuitive overall summary of the truncation approach.

An intuitive description of the truncation approach. The truncation approach consists in solving the Ramsey problem of a finite state space model approximating the full model. This finite state-space model is called the truncated model and results from an aggregation of the full model, which groups agents together according to their recent idiosyncratic history and expresses the model in terms of these groups of agents. These groups are called *truncated histories*. As an example, we assume that there are 2 productivity levels: $y_l < y_h$ and we group agents according to their histories over the 3 last periods. Agents with the productivity history (y_l, y_l, y_l) over the three last periods will thus be assigned to the same truncated history; the same goes for history (y_l, y_h, y_l) , and so on. Agents move from one truncated history to another when their productivity history is updated in the beginning of the next period. For example, an agent with productivity history (y_l, y_l, y_l) who draws the productivity y_h will have history (y_l, y_l, y_h) – and hence will move from truncated history (y_l, y_l, y_l) to (y_l, y_l, y_h) .

All agents of a given truncated history are assigned the same allocation, equal to the average allocation of all agents belonging to the group. *Each* truncated history is thus akin to a “representative agent”. A key aspect of the aggregation is that the truncated allocation has to be consistent with the fact that the representative agents of a truncated history are expected-utility maximizers. This does not raise any difficulty for aggregating linear equations, such as budget constraints. Indeed, the average budget constraint is equal to the budget constraint with average quantity (at least as long as there are no non-linearities, such as progressive taxation). However, this is not the case in general, for aggregating non-linear equations. For example, the average marginal utility of individual consumption is not equal to the marginal utility of average consumption. The discrepancy is due to the heterogeneity within each truncated history (and the non-linearity of the marginal utility function). Hence, we introduce parameters, called ξ s, that are necessary to correctly aggregate nonlinear equations. In our case, the non-linearities come from the utility function and its derivatives: u , u' and u'' . As such, we introduce three sets of correction coefficients. They are denoted by ξ^0 , ξ^1 , and ξ^2 and correspond to u , u' and u'' respectively.¹³ This extends LeGrand and Ragot (2022a), who for the sake of simplicity only

¹³Other ξ s could be introduced for other non-linearities – such as progressive taxation for instance.

considered a single set ξ .

The truncation approach for solving the Ramsey problem then proceeds as follows. We guess a value for the optimal instrument and compute the corresponding Bewley allocation (using standard methods such as EGM, for instance). We then compute the coefficients ξ^j ($j = 0, 1, 2$) to properly aggregate non-linear equations and obtain the truncated model. Finally, we derive the Ramsey FOCs of the truncated model and verify whether they hold given the guess of the instrument value and the implied truncated allocation. If they do, then this instrument value is optimal. Otherwise, we iterate on the instrument value and redo the whole process (which includes re-solving the Bewley model).

Two remarks are important. First, the truncation approach converges to an instrument value that corresponds to an existing Bewley equilibrium (by construction). Second, the truncation length of the productivity history (equal to 3 in the example above) used to build the set of histories is an exogenous model parameter. LeGrand and Ragot (2022a) have shown that when this length increases, the truncated economy converges to the Bewley economy, but that in practice small truncation lengths offer very good accuracy.

4.1 Constructing the truncated model

The truncated model results from an aggregation of the full model. The procedure can be applied to any heterogeneous-agent model. We explain the process formally in Section 4.1.1 and document how to use it in practice in Section 4.1.2.

4.1.1 A formal description

The starting point is the sequential solution of the full-fledged incomplete-market model, which can be written as a set of policy rules, mapping histories into choices:¹⁴ $a_t(y^t, a_{-1}, Z^t)$ and $c_t(y^t, a_{-1}, Z^t)$, for $y^t \in \mathcal{Y}^{t+1}$, $Z^t \in \mathcal{Z}^{t+1}$ using previous notation. The main idea of the truncated model is to group together agents who have the same productivity history for a given number of consecutive past periods, and then to state the model in terms of this finite number of groups of agents. We use the term *truncation length* for the exogenous parameter setting the length of each shared productivity history, denoted $N > 0$. The construction of the truncated model consists in truncating idiosyncratic histories and following a finite number of representative agents with adjusted weights. A key step in the truncated model is the computation of these adjusted weights.

Consider an agent with complete idiosyncratic history $y^\infty = (\dots, y_{t-N-1}, y_{t-N}, y_{t-N+1}, y_{t-N+2}, \dots, y_{t-1}, y_t)$ at date t (y_t being the current productivity status). If their history over the last N periods is such that $(y_{t-N+1}, \dots, y_{t-1}, y_t) = (y_{-N+1}^N, \dots, y_{-1}^N, y_0^N)$, this agent will be assigned to

¹⁴We consider the sequential representation to ease exposition, whereas the actual implementation uses the recursive representation, which is the standard input of computational methods as discussed in Section 4.1.2.

truncated history $y^N := (y_{-N+1}^N, \dots, y_{-1}^N, y_0^N)$ at date t , independent of earlier productivity levels (i.e., of the sequence $(\dots, y_{t-N-1}, y_{t-N})$). Because the number of productivity levels is finite and equal to $n_y := \text{Card}(\mathcal{Y})$, the number of truncated histories of length N is also finite and equal to $N_{tot} := n_y^N$. Because every agent draws a new idiosyncratic status in every period, a given agent is in general assigned to a different truncated history in each period. For instance, if the previous agent with history y^∞ at t is endowed with productivity y_{t+1} at $t+1$, their $t+1$ -history will be: $\tilde{y}^\infty := (\dots, y_{t-N-1}, y_{t-N}, y_{t-N+1}, \dots, y_{t-1}, y_t, y_{t+1})$ and they will be assigned at date $t+1$ to truncated history $\tilde{y}^N = (y_{-N+2}^N, \dots, y_1^N, y_0^N, \tilde{y}_0^N)$, where $\tilde{y}_0^N := y_{t+1}$.¹⁵ The probability, denoted $\Pi_{y^N \tilde{y}^N}$, that an agent transitions from history y^N to history \tilde{y}^N is the probability that the agent transitions from productivity level y_0 to \tilde{y}_0 , or formally:

$$\Pi_{y^N \tilde{y}^N} = 1_{\tilde{y}^N \succeq y^N} \Pi_{y_0^N \tilde{y}_0^N}, \quad (19)$$

where $1_{\tilde{y}^N \succeq y^N} = 1$ if \tilde{y}^N is a possible continuation of y^N , or formally: $\tilde{y}_{-1}^N = y_0^N$, $\tilde{y}_{-2}^N = y_{-1}^N, \dots$, $\tilde{y}_{-N+1}^N = y_{-1}^N$; $1_{\tilde{y}^N \succeq y^N} = 0$ otherwise. The population of agents with truncated history y^N can be defined recursively from the previous probabilities as:

$$S_{y^N} = \sum_{\hat{y}^N \in \mathcal{Y}^N} S_{\hat{y}^N} \Pi_{\hat{y}^N y^N}. \quad (20)$$

Because the truncated model aims to express the economy using truncated histories, we must derive each truncated history's consumption level and end-of-period savings, denoted $c_{t,y^N}(Z^t)$ and $a_{t,y^N}(Z^t)$ respectively (or simply c_{t,y^N} and a_{t,y^N} when there is no ambiguity). These quantities are defined as the corresponding average values among agents sharing the same truncated history y^N . For instance, for savings:

$$a_{t,y^N} := \frac{1}{S_{y^N}} \int_{a_{-1}} \sum_{\hat{y}^t \in \mathcal{Y}^{t+1} | (\hat{y}_{t-N+1}, \dots, \hat{y}_t) = y^N} \theta_t(\hat{y}^t) a_t(\hat{y}^t, a_{-1}, Z^t) d\Lambda_0(a_{-1}, y_0).$$

To compute beginning-of-period savings, denoted \tilde{a}_{t,y^N} , we must account for the possibility that agents with current truncated history y^N had different truncated histories \hat{y}^N in the previous period. Formally:

$$\tilde{a}_{t,y^N} = \frac{1}{S_{y^N}} \sum_{\hat{y}^N \in \mathcal{Y}^N} S_{\hat{y}^N} \Pi_{\hat{y}^N y^N} a_{t-1, \hat{y}^N}. \quad (21)$$

We can aggregate the individual budget constraint (5) along a common truncated history and obtain the following truncated-history budget constraint:

$$c_{t,y^N} + a_{t,y^N} = (1 + r_t) \tilde{a}_{t,y^N} + w_t y_0^N - T_t, \quad (22)$$

¹⁵For consistency, we denote future truncated histories with a tilde, past ones with a hat, and current ones without decoration.

where y_0^N is the current productivity level for y^N . This aggregation is straightforward because budget constraints are linear. Aggregating utility or its derivatives is less straightforward, because utility levels and marginal utilities are not linear in consumption. To proceed with the aggregation of non-linear equations (and in particular of the utility function and its derivatives), we define the following history-dependent parameters:

$$\xi_{t,y^N}^0 := \frac{1}{u(c_{t,y^N})} \int_{a_{-1}} \sum_{\hat{y}^t \in \mathcal{Y}^{t+1} | (\hat{y}_{t-N+1}, \dots, \hat{y}_t) = y^N} \theta_t(\hat{y}^t) u(c_t(\hat{y}^t, a_{-1}, Z^t)) d\Lambda_0(a_{-1}, y_0), \quad (23)$$

$$\xi_{t,y^N}^1 := \frac{1}{u'(c_{t,y^N})} \int_{a_{-1}} \sum_{\hat{y}^t \in \mathcal{Y}^{t+1} | (\hat{y}_{t-N+1}, \dots, \hat{y}_t) = y^N} \theta_t(\hat{y}^t) u'(c_t(\hat{y}^t, a_{-1}, Z^t)) d\Lambda_0(a_{-1}, y_0), \quad (24)$$

$$\xi_{t,y^N}^2 := \frac{1}{u''(c_{t,y^N})} \int_{a_{-1}} \sum_{\hat{y}^t \in \mathcal{Y}^{t+1} | (\hat{y}_{t-N+1}, \dots, \hat{y}_t) = y^N} \theta_t(\hat{y}^t) u''(c_t(\hat{y}^t, a_{-1}, Z^t)) d\Lambda_0(a_{-1}, y_0). \quad (25)$$

These parameters enable us to reconcile the aggregation of the transformation of a given quantity with the transformation of the aggregation of the quantity. For instance, the aggregate utility in period t can be expressed as the sum over all histories and all initial asset holdings of individual utility levels:

$$\int_{a_{-1}} \sum_{y^t \in \mathcal{Y}^{t+1}} \theta_t(y^t) u(c_t(y^t, a_{-1}, Z^t)) d\Lambda_0(a_{-1}, y_0).$$

Generally, because the utility function is not linear, it differs from the utility derived from truncated-history consumption levels. The role of ξ^0 is precisely to reconcile both, as the previous aggregate utility is also: $\sum_{y^N \in \mathcal{Y}^N} S_{y^N} \xi_{t,y^N}^0 u(c_{t,y^N})$, which is the sum of truncated-history utilities, weighted by ξ^0 .

These ξ^j ($j = 0, 1, 2$) parameters capture the residual heterogeneity within each truncated history attributable to the fact that agents experienced different idiosyncratic histories N periods ago and before. Indeed, each truncated history groups together agents sharing the same history over the last N periods, while ignoring the distant past. In the absence of within-truncated-history heterogeneity, the ξ^j would be all equal to 1.

On the theoretical side, the ξ^j ($j = 0, 1, 2$) parameters can be shown to converge toward 1 as the length of the truncation N increases. However, because N remains small in practice, this asymptotic result has little practical implication. Fortunately, as we check in our quantitative exercise in Section 5, even for short truncation lengths, the ξ^j allow us to obtain accurate results. We explain in Section 4.1.2 how to easily compute the ξ^j for $j = 0, 1, 2$.

Finally, $\mathcal{C}_{t,N}$ denotes the set of credit-constrained truncated histories at date t . With this

notation, the Euler equations can be written as follows:

$$\forall y^N \in \mathcal{Y}^N, \xi_{t,y^N}^1 u'(c_{t,y^N}) = \beta(1+r_{t+1}) \sum_{\tilde{y}^N \in \mathcal{Y}^N} \Pi_{y^N \tilde{y}^N} \xi_{t+1,\tilde{y}^N}^1 u'(c_{t+1,\tilde{y}^N}) + \nu_{t,y^N}, \quad (26)$$

$$\forall y^N \in \mathcal{C}_{t,N}, a_{t,y^N} = 0, \quad (27)$$

where the expectation over the future idiosyncratic state has been addressed explicitly. Equation (26) is the Euler equation at the truncated history level for non-credit-constrained histories, and equation (27) corresponds to credit-constrained truncated histories holding zero assets.

We have thus characterized the truncated model, whose main advantage is to feature limited heterogeneity, and is thereby characterized by a finite number of equations and unknowns.

4.1.2 Computing the truncated model in practice

We now explain how the truncated allocation and (ξ^0, ξ^1, ξ^2) can be computed in practice, relying on the standard recursive approach instead of the sequential one of Section 4.1.1.

As proved in the literature (for example by Açıkgöz, 2018, for recent results), the solution of the Bewley model is characterized by a steady-state wealth distribution, $\Lambda_\infty : (a, y) \in [0, +\infty) \times \mathcal{Y} \rightarrow \mathbb{R}_+$, and a set of policy rules for savings, denoted $g_a : (a, y) \in [0, +\infty) \times \mathcal{Y} \rightarrow \mathbb{R}_+$. These quantities can efficiently be computed using the EGM procedure for instance.

To compute the truncated allocation, we must obtain the wealth distribution with respect to truncated histories of length N , which will be denoted $\tilde{\Lambda}_N : (a, y^N) \in [0; +\infty) \times \mathcal{Y}^N \rightarrow \mathbb{R}_+$, such that $\tilde{\Lambda}_N(da, y^N)$ is the measure of agents with wealth in $[a, a + da)$ and truncated history $y^N = (y_{-N+1}, \dots, y_{-1}, y_0)$. The measure $\tilde{\Lambda}_N$ can be computed by starting from the wealth distribution of agents in state y_{-N+1} , which is $\Lambda_\infty(\cdot, y_{-N+1})$, and then successively applying the sequence of savings policy functions corresponding to $y^N = (y_{-N+1}, \dots, y_{-1}, y_0)$, which is $g_a(\cdot, y_{-N+1}), g_a(\cdot, y_{-N+2}), \dots, g_a(\cdot, y_0)$. From a practical perspective, computing $\tilde{\Lambda}_N$ is both straightforward and very fast, because the process consists of multiplying a vector—modelling the initial distribution—with N different transition matrices—modelling the policy functions.

Using the measure $\tilde{\Lambda}_N$, we deduce that the end-of-period savings, a_{y^N} , for a truncated history y^N can be computed as:

$$a_{y^N} = \int_{a \in [0, \infty)} a \tilde{\Lambda}_N(da, y^N), \quad (28)$$

where in our case the savings are actually bounded from above (see Açıkgöz, 2018). Again, this quantity is very simple to compute in practice, as it only involves basic linear algebra manipulations. We can then deduce from (21) the beginning-of-period savings $(\tilde{a}_{y^N})_{y^N}$, as consumption $(c_{y^N})_{y^N}$ from the truncated-history budget constraint (22), as well as the Lagrange multipliers $(\nu_{y^N})_{y^N}$ on the credit constraints.

The set \mathcal{C} of credit constrained histories is determined based on these Lagrange multipliers

$(\nu_{y^N})_{y^N}$. We take the n_C histories, which have the largest multiplier values (ν_{y^N}) and such that the total size of the n_C histories is as close as possible to the share of credit-constrained agents in the Bewley model.

Finally, the ξ^j parameters need to be computed for properly aggregating non-linear equations.

We deduce from equations (20)–(27) that the steady-state economy is then characterized by the following set of equations:

$$\tilde{a}_{y^N} = \frac{1}{S_{y^N}} \sum_{\hat{y}^N \in \mathcal{Y}^N} S_{\hat{y}^N} \Pi_{\hat{y}^N, y^N} a_{\hat{y}^N}, \quad (29)$$

$$c_{y^N} + a_{y^N} = (1+r)\tilde{a}_k + w y_0^N - T_\infty, \quad (30)$$

$$\xi_{y^N}^1 u'(c_{y^N}) = \beta(1+r) \sum_{\tilde{y}^N=1}^{N_{tot}} \Pi_{y^N, \tilde{y}^N} \xi_{\tilde{y}^N}^1 u'(c_{\tilde{y}^N}) + \nu_{y^N}, \quad (31)$$

$$y^N \in \mathcal{C}, \quad a_{y^N} = 0, \quad (32)$$

where y_0^N is the current productivity level of history y^N .

4.2 The truncation approach at the steady-state

We now present the truncation approach at the steady state. It relies on the previous truncated model to solve for optimal policies in the steady state. The truncation approach takes advantage of the finite state-space property of the truncated model and proceeds in three steps:

1. For a given steady-state tax level T_∞ , we compute the truncated model as explained in Section 4.1.2;
2. We derive the FOCs of the Ramsey program for the truncated model, and then compute all associated Lagrange multipliers;
3. The Lagrange multipliers allow us to check whether the planner's FOC characterizing the optimal value of T_∞ holds. If the constraint holds, then T_∞ is the optimal steady-state tax. If not, the procedure must be repeated for an updated value of T_0 .

We provide the derivation of all three steps and then present the algorithm.

4.2.1 FOCs of the Ramsey program in the truncated model

Details of the computation are provided in Appendix C. Before stating the conditions, we introduce the quantity $\tilde{\lambda}_{t, y^N}$, defined as follows:

$$\tilde{\lambda}_{t, y^N} = \frac{1}{S_{t, y^N}} \sum_{\tilde{y}^N \in \mathcal{Y}^N} S_{t-1, \tilde{y}^N} \Pi_{t, \tilde{y}^N, y^N} \lambda_{t-1, \tilde{y}^N}. \quad (33)$$

This quantity corresponds to the previous period Lagrange multiplier on the Euler equation for agents with truncated history y^N at date t . These agents may have different truncated histories in the previous period, which explains equation (33), similar to the beginning-of-period wealth, \tilde{a}_{t,y^N} , in equation (21).

We use equation (33) to express the quantity ψ_{t,y^N} , which is the social valuation of liquidity for truncated history y^N , and is thus the parallel of the individual quantity ψ_t^i in equation (16). The formal definition features the within-heterogeneity parameters ξ^1 and ξ^2 and is:

$$\psi_{t,y^N} = \xi_{t,y^N}^1 u'(c_{t,y^N}) - (\lambda_{t,y^N} - \tilde{\lambda}_{t,y^N}(1+r_t)) \xi_{t,y^N}^2 \xi_{t,y^N} u''(c_{t,y^N}). \quad (34)$$

With this notation, the FOCs for the Ramsey allocation in the truncated economy can be written as follows:

$$v'(T_t) = \sum_{y^N \in \mathcal{Y}^N} S_{y^N} \psi_{t,y^N}. \quad (35)$$

These FOCs are very similar to condition (17). Equations characterizing the dynamics of the Ramsey model are provided in Appendix C.

4.2.2 Computing the Ramsey allocation at the steady state using matrix notation

We now explain how to compute the steady-state Ramsey allocation, including the value of the planner's instrument, for the truncated model. The truncated model can be represented very efficiently by relying on matrix notation, which enables us to derive Lagrange multipliers using simple linear algebra. To implement the numbering of truncated histories in practice, a convenient solution is to use the enumeration in base n_y . A truncated history $y^N = (y_{-N+1}, \dots, y_{-1}, y_0)$ will be assigned the index $k := 1 + \sum_{j=0}^{N-1} n_y^k (n_{y_{-j}} - 1)$, where $n_{y_{-j}} \in \{1, \dots, n_y\}$ is the position of productivity level y_{-j} in the set \mathcal{Y} , from 1 for the smallest productivity level to n_y for the largest. The index belongs by construction to the set $\{1, \dots, N_{tot}\}$.

We then introduce the following matrix notation:

- $\mathbf{S} = (S_k)_{k=1, \dots, N_{tot}}$, the N_{tot} -vector of sizes;
- $\mathbf{c} = (c_k)_{k=1, \dots, N_{tot}}$, the N_{tot} -vector of consumption levels;
- $\mathbf{a} = (a_k)_{k=1, \dots, N_{tot}}$ and $\tilde{\mathbf{a}} = (\tilde{a}_k)_{k=1, \dots, N_{tot}}$, the N_{tot} -vector of end-of-period and beginning-of-period asset holdings, respectively;
- $u'(\mathbf{c}) = (u'(c_k))_{k=1, \dots, N_{tot}}$, the N_{tot} -vector of marginal utilities;
- $\boldsymbol{\xi}^j = (\xi_k^j)_{k=1, \dots, N_{tot}}$, the vector of residual-heterogeneity parameters ($j = 0, 1, 2$);

- Finally, “ \circ ” denotes the term-by-term product of two vectors of the same size, which is another vector of the same size: $\mathbf{x} \circ \mathbf{z} = (x_{y^N}) \circ (z_{y^N}) = (x_{y^N} z_{y^N})$.¹⁶

We can now state our result for the computation of Lagrange multipliers.

Proposition 1 (Steady-state Lagrange multipliers) *Consider a steady-state tax value T_∞ (not necessarily optimal), for which the truncated model can be computed.¹⁷ Then, there exist two matrices, \mathbf{M}_1 and \mathbf{M}_2 , depending only on the equilibrium allocation, such that:*

$$\boldsymbol{\lambda} = \mathbf{M}_1(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})) \text{ and } \boldsymbol{\psi} = \mathbf{M}_2(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})). \quad (36)$$

Proposition 1 states that the steady-state values of Lagrange multipliers and the social value of liquidity can be deduced from the truncated allocation, using basic linear algebra. This computation is possible for any value of the instrument for which the steady-state equilibrium can be computed. The proof can be found in Appendix D, which provides a step-by-step computation of the expressions of matrices \mathbf{M}_1 and \mathbf{M}_2 – the coefficients $\boldsymbol{\xi}^2$ s appear in the expression of the matrix \mathbf{M}_2 . After the social value of liquidity vector $\boldsymbol{\psi}$ has been computed, it is straightforward to check the optimality of the steady-state tax level T_∞ using condition (35), which can be written in a matrix form as:

$$v'(T_\infty) = \mathbf{S}^\top \boldsymbol{\psi}, \quad (37)$$

where $\mathbf{S}^\top \boldsymbol{\psi}$ is a scalar.

This result provides the basis for the following algorithm summarizing the successive steps to compute the steady-state Ramsey allocation using the truncation approach.

Algorithm 3 (The truncation approach: Steady state) *Set a precision criterion $\varepsilon > 0$ and a truncation length N .*

1. *Set an initial value for the steady-state lump-sum tax T_∞ .*
2. *Solve the full-fledged Bewley model for the value of the instrument T_∞ .*
3. *Construct the truncated model for truncation length N using the Section 4.1.2. We obtain the steady-state truncated allocation $(\mathbf{a}, \mathbf{c}, \boldsymbol{\xi}^0, \boldsymbol{\xi}^1, \boldsymbol{\xi}^2)$ and the set of credit-constrained histories \mathcal{C} , associated with T_∞ .*
4. *Compute matrices \mathbf{M}_1 and \mathbf{M}_2 using equations (70) and (71) of Appendix D.*
5. *Compute the vectors $\boldsymbol{\lambda}$ and $\boldsymbol{\psi}$ using equation (36).*

¹⁶This operation is also known as the Hadamard product.

¹⁷We characterize the full path $(T_t)_{t \geq 0}$ below.

6. If $|v'(T_\infty) - \mathbf{S}^\top \boldsymbol{\psi}| < \varepsilon$, then the algorithm stops, and the optimal steady-state Ramsey truncated allocation is given by $(T_\infty, \mathbf{a}, \mathbf{c}, \boldsymbol{\xi}^0, \boldsymbol{\xi}^1, \boldsymbol{\xi}^2)$ and T_∞ is the optimal steady-state truncation tax rate. Otherwise, update T_∞ and start again at Step 2.

Algorithm 3 shows how the truncation approach computes the steady-state optimal lump-sum tax as a fixed point of an iteration procedure. The algorithm starts with a guess for the steady-state value of the lump-sum tax T_∞ . It then computes the allocation of the Bewley model corresponding to the lump sum tax T_∞ . We can then deduce the steady-state allocation of the truncated model. The matrix notation (further details are provided in Appendix D) enables us to compute the social valuation of liquidity $\boldsymbol{\psi}$. Finally, the optimality of the steady-state lump-sum tax T_∞ is checked using equation (37).

Three remarks are in order. First, Step 2 of Algorithm 3 implies that we compute the Bewley allocation for each value of the steady-state lump-sum tax. Consequently, the algorithm converges by construction to a Bewley equilibrium that does exist. Second, the computational implementation of Algorithm 3 is fast. At every step, the computationally intensive task is to simulate the Bewley model for the steady-state lump-sum tax T_∞ . The other steps (in particular, 3 to 6) only involve linear algebra and are very fast to perform (taking less than a second). Third, the computation of the steady-state tax only involves steady-state FOCs and is therefore by construction immune to the choice of initial distribution. Compared to Algorithm 2 for the transition approach, the truncation approach only requires a fixed point over the instrument and not a joint fixed point over the instrument and the initial distribution.

We check in Section 5 that Algorithm 3 yields an accurate solution, and we compare it to that of the transition approach.

4.3 Computing the truncation tax path

One feature of the truncation approach is that it separates the computation of the steady-state allocation (including the planner's instruments) from the computation of the full instrument path. As explained above, this makes the steady-state computation immune to the choice of initial distribution. This also greatly simplifies the computation of the full path of the planner's instruments. Indeed, the steady state being known, standard perturbation methods can be used to compute the path.

The algorithm below summarizes the various steps to compute the instrument path, after the steady-state allocation has been computed using Algorithm 3.

Algorithm 4 (The truncation approach: Simulating the full instrument path) *We consider as given a truncation length, $N > 0$, and a precision criterion, $\varepsilon > 0$.*

1. Steady state. *We use Algorithm 3 to compute the optimal steady-state Ramsey truncated allocation with precision ε and truncation length N .*

2. Truncated model. *The dynamics of the truncated model with truncation length N are characterized by equations (47)–(55) of Appendix C. The within-heterogeneity coefficients and the set of credit-constrained histories remain equal to their steady-state values, (ξ^0, ξ^1, ξ^2) and \mathcal{C} , respectively.*
3. Initial conditions. *The initial wealth distribution can be set arbitrarily and the initial distribution of Lagrange multipliers $(\lambda_{-1, y^N})_{y^N}$ is set to 0.*
4. Simulation. *Based on the steady-state allocation of Step 1, the model equations of Step 2, and the initial conditions of Step 3, we simulate the path of variables (including planner’s instruments) in the truncated model using a perturbation method.*

Algorithm 4 describes a straightforward solution to simulate the full path of model variables for the Ramsey truncated solution. The core of the algorithm is the perturbation method of Step 4, which can be performed using existing and well-tested software, such as Dynare (Adjemian et al., 2011) (which is already widely adopted for solving DSGE models).

Algorithm 4 involves two main assumptions. The first is that the ξ^j coefficients remain equal to their steady-state values along the path. This implies that the heterogeneity within truncated histories is constant through time and is the same as in the steady state. Importantly, this assumption does not preclude the existence of heterogeneity within each truncated history. The second assumption is that credit-constrained truncated histories are determined at the steady state and must remain unchanged along the transition. This can be checked, especially if the initial distribution markedly differs from the steady-state one.¹⁸

We use Algorithm 4 to compute a reoptimization shock in Section 5.4, which allows us to assess the severity of time-inconsistency.

5 Quantitative exercise

We now turn to the quantitative exercise. This section is structured as follows:

1. We specify our calibration in Section 5.1.
2. We compute the optimal steady-state truncation tax rate in Section 5.2.
3. We compute the optimal fixed-point transition tax rate in Section 5.3.
4. In Section 5.4, we compute the optimal truncation tax path and document the severity of time-inconsistency issue. We also explain the gap between the steady-state truncation tax rate and the fixed-point transition tax rate.

¹⁸Assuming that the set of credit constrained histories is time-varying would require iterating over the path of this set. We leave that for future research.

5. In Section 5.5, we check that the time-0 truncation tax path actually generates the highest aggregate welfare when the economy is simulated – showing that the solution of the truncation approach is close to the Ramsey solution of the full model.
6. We perform several robustness checks in Sections 5.6 and 5.7.

5.1 The calibration

The period is a quarter. The discount factor is set to $\beta = 0.99$. The period utility function for the private good is $\log(c)$. The utility function for the public good is $v(G) = G^\vartheta$. We set the parameter ϑ to 24%, to target a value of steady-state public consumption over GDP of 8.0%; this roughly corresponds to US government consumption of specific final goods minus public investment. In the production function of equation (1), the capital share is set to $\alpha = 36\%$ and the depreciation rate is set to $\delta = 2.5\%$, as in Krueger et al. (2018), among others.

Idiosyncratic productivity is modeled as an AR(1) productivity process: $\log y_t = \rho_y \log y_{t-1} + \varepsilon_t^y$, with $\varepsilon_t^y \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_y^2)$. We calibrate the parameters ρ_y and σ_y to obtain a realistic income process, following the estimates of Krueger et al. (2018). We use a quarterly persistence of $\rho_y = 0.996$ and a quarterly standard deviation of $\sigma_y = 4.39\%$; for the log of earnings, these generate an annual persistence of 0.9849 and an annual standard deviation of 8.71%. The Rouwenhorst (1995) procedure is then used to discretize the productivity process into five idiosyncratic states with a constant transition matrix.

Table 1 summarizes the model parameters.

Parameter	Description	Value
β	Discount factor	0.98
α	Capital share	0.36
δ	Depreciation rate	0.025
$\tau = T/Y$	Tax rate	8%
ϑ	Curvature of util. public good	24%
ρ_y	Autocorrelation idio. income	0.996
σ_y	Standard dev. idio. income	4.39%

Table 1: Parameter values in the baseline calibration. See text for further descriptions and targets.

We can now compute the steady-state equilibrium of the model, using the standard EGM method.¹⁹ The implied capital-output ratio is $K/Y = 2.67$; the consumption-output ratio is $C/Y = 0.65$. Table 2 provides descriptive statistics for the wealth distribution in the data and

¹⁹We use the EGM method with 100 points on an exponential grid point for wealth, following Carrol (2006) and Boppart et al. (2018), among others.

model. We use data from the Panel Study of Income Dynamics (PSID) in 2006 and from the Survey of Consumer Finances (SCF) in 2007, to abstract from the effects of the 2008 global financial crisis.

Wealth statistics	Data		Model
	PSID, 06	SCF, 07	
Q1 (minimum)	-0.9	-0.2	0.0
Q2	0.8	1.2	0.3
Q3 (median)	4.4	4.6	5.6
Q4	13.0	11.9	21.4
Q5 (maximum)	82.7	82.5	72.7
Gini	0.77	0.78	0.71

Table 2: Steady-state wealth distribution.

The model reproduces the wealth distribution relatively well. It is known that other mechanisms must be introduced to match the very top of the wealth distribution (such as entrepreneurship or stochastic β s).

5.2 Results with the truncation approach

We first compute the steady-state truncation tax rate using Algorithm 3. We consider a truncation length $N = 5$, which implies $5^5 = 3,125$ histories. The model by construction generates the steady-state truncation tax-to-GDP ratio $\tau_\infty^p := T_\infty^p/Y = 8\%$, which was the targeted value. (The superscript p refers to optimal tax computed with the truncation approach.) We perform a sensitivity test on the choice of truncation length N in Section 5.6, and the value of N (beyond $N \geq 2$) appears to have a very modest quantitative impact on the optimal provision of the public good.

We now provide further explanation of the different contributing forces to the truncation tax-to-GDP ratio of 8%. To evaluate the contribution of indirect effects (through savings) to the optimal tax, we can compute the value of the tax $T_t^{partial}$ that would correspond to a planner valuing only direct effects—thus discarding the role of Lagrange multipliers on Euler equation (18). Formally, this corresponds to:

$$v'(T_t^{partial}) = \int_i u'(c_t^i) \ell(di).$$

The numerical quantification implies that $T_t^{partial}$ is 4.3% lower than T_∞^p , and that the tax-to-GDP ratio decreases from $\tau_\infty^p = 8.00\%$ to $\tau^{partial} = 7.76\%$. A higher tax raises precautionary savings and hence boosts capital and aggregate consumption. When the planner internalizes the

savings distortions (through the term in $u''(c_t^i)$ in equation (18)), the tax is higher because the benefits of this higher tax are factored in by the planner.

5.3 Results with the transition approach

We first solve for the *optimal transition tax rate* as presented in Definition 3. To do this, we implement the transition approach of Algorithm 1, which requires an initial distribution Λ_0 as input. We solve for the transition tax rate with two different initial distributions. We first consider the steady-state distribution presented in Table 2 of Section 5.1. We then multiply the wealth level of all agents by 0.9 (implying a 10% decrease in their initial wealth). We call this distribution the low distribution, which corresponds to the transition tax-to-GDP ratio denoted τ^{low} . The second initial distribution, called *high*, is a 10% increase in the wealth of all agents starting from the distribution of Table 2. The corresponding transition tax-to-GDP ratio is denoted τ^{high} .

The computation of the two optimal transition tax rates yields $\tau^{low} = 6.4\%$ and $\tau^{high} = 8.45\%$. When the initial wealth is low, the tax—that is imposed to remain constant throughout the path—affects private consumption at the beginning of the transition when agents have few resources and hence a high marginal utility for private consumption. This contributes to setting a low tax. In contrast, when the initial wealth is high, the tax can also be higher because in the first periods agents have a relatively low marginal utility for private consumption. This illustrates that in the context of the transition approach, the initial distribution has a sizable impact on the optimal outcome.

To neutralize the effect of the initial distribution, we compute the optimal fixed-point transition tax rate T_c , as presented in Definition 4. The tax rate T_c is such that the initial distribution Λ_c equals the long-run distribution: $\Lambda^c = \Lambda_\infty(T_c, \Lambda^c)$. The computation yields a tax-to-GDP ratio $\tau^c = 7.8\%$, which lies between the low and high values, τ^{low} and τ^{high} , that we computed above.

To illustrate the optimality of τ^c , in Figure 1 we plot aggregate welfare as a function of the tax (plain blue line), with an initial distribution equal to the long-run one. More precisely, for each tax value T , we iterate on the initial distribution Λ_0 to compute the fixed point verifying $\Lambda_0 = \Lambda_\infty(T, \Lambda_0)$, and we then compute the aggregate welfare with transitions when the distribution evolves from Λ_0 to $\Lambda_\infty(T, \Lambda_0) = \Lambda_0$. The tax rate, on the x -axis, is reported as tax-to-GDP ratio $\tau = T/Y$, and welfare, on the y -axis, is reported as the percentage loss in consumption compared to the optimal welfare. In addition to τ^c , Figure 1 also shows three other tax rates (as orange dashed vertical bars): τ^{low} corresponding to the low-wealth initial distribution, τ^{high} corresponding to the high-wealth initial distribution, and τ_∞^p corresponding to the optimal truncation tax rate.

Figure 1 confirms that the transition tax rate is highly sensitive to the initial distribution,




Figure 1: Welfare with transitions and fixed-point initial distribution computed as a function of the tax rate. See the text for the definitions of other tax rates.

and we can find transition tax rates that are greater or smaller than the fixed-point transition tax rate. Figure 1 also illustrates that the fixed-point transition tax rate τ^c is close to, but lower than, the truncation tax rate τ_∞^p . Although we neutralize the effect of the initial distribution in the computation of the transition tax rate, both tax rates differ. We quantitatively confirm below that this comes from the fact that the Ramsey planner prefers to choose a non-constant tax path.

5.4 Indirect evidence of the accuracy of the truncation approach

We now further investigate the gap between the truncation and transition tax rates τ_∞^p and τ^c . We compute the truncation tax path following Algorithm 4. We neutralize the effect of the initial wealth distribution by setting it equal to its steady-state value. This can be seen as a pure reoptimization shock in period 0, in the sense that the economy starts from the steady state but we turn off the planner's commitment. Figure 2 plots the optimal truncation tax path. We denote it by $(T_t^p)_{t \geq 0}$, and denote the corresponding optimal tax-to-GDP path $(\tau_t^p)_{t \geq 0}$. By construction, the tax path converges to the steady-state tax T_∞^p in the long-run. Despite starting from the steady-state distribution, the truncation tax path is not constant.

The tax rate drops at impact to a level that is smaller than the fixed-point transition tax rate τ^c . The intuition is as follows. Because $\lambda_{-1}^i = 0$, the planner is not committed to maintain




Figure 2: The truncation tax path.

the tax rate equal to τ_∞^p and chooses to decrease the tax. This allows agents, and especially those with a high marginal propensity to consume, to increase private consumption in the short run. The tax path then increases and converges back to the long-run steady-state value τ_∞^p . This increases agents' incentives to build up savings, which increases the capital stock and, in turn, the consumption of private goods. The consumption of the public good also increases with the tax. The reoptimization that occurs in the truncation tax path stems from the time-inconsistency of the Ramsey program in heterogeneous-agent models, discussed in Section 2.6.

Figure 2 allows us to understand the gap between the steady-state truncation tax τ_∞^p and the fixed-point transition tax τ^c . As explained above, the transition tax can be seen as a solution to a Ramsey program, to which we add the constraint that the tax path be constant. Similarly, the transition approach can be seen as selecting the “average” tax rate over the whole Ramsey path, so as to balance the benefits and costs along the transition path. We check this statement by computing the average discounted tax rate over the transition path $(\tau_t^p)_{t \geq 0}$, where we discount future tax values using the discount factor β to obtain an approximation of the discounted value of the optimal path of taxes. We denote this weighted discounted tax rate τ^{weight} , which is formally defined as:

$$\tau^{weight} = \frac{\sum_{t=1}^{400} \beta^{t-1} \tau_t^p}{\sum_{t=1}^{400} \beta^{t-1}}.$$

The computation yields $\tau^{weight} = 7.84\%$, which is very close to the tax rate τ^c computed with the transition approach. This computation shows that the gap between the fixed-point transition

and the truncation tax rates is well explained by the severity of the time-inconsistency in the Ramsey program. Since the tax path is computed using the truncation approach, this is also indirect evidence that the truncation approach does a good job in computing the actual Ramsey solution. Computing the Ramsey solution of the truncated model is a very good approximation of the actual Ramsey solution.

5.5 Direct evidence of the accuracy of the truncation approach

We now verify the previous indirect evidence and directly check that the truncation approach delivers an accurate solution to the full Ramsey program. We do this by checking that any perturbation of the truncation tax path in Figure 2 implies a decrease in aggregate welfare. Importantly, we compute aggregate welfare using an extended transition approach, where the given tax path is not constant but time-varying and deterministic. This means that the truncated model (or any of the other simplifying assumptions of the truncation approach) is not involved in the computation of aggregate welfare.

We first formalize the perturbation of the truncation tax path of Figure 2. The perturbed path, denoted by $(T_t^\kappa)_t$ for any real value κ , is defined as:

$$T_t^\kappa = T_\infty^p + (1 + \kappa)(T_t^p - T_\infty^p), t \geq 0, \quad (38)$$

where T_∞^p is the steady-state value of the truncation tax path. Independent of κ , any path T_t^κ converges at the steady state toward T_∞^p , which means that there is no steady-state deviation. The parameter κ modifies the initial drop in the tax rate and the speed of convergence to the steady-state value. When $\kappa = 0$, we implement the truncation path, and when $\kappa > 0$ ($\kappa < 0$), we implement a higher (lower) path.

We incorporate equation (38) for a given κ as an exogenous rule into the Bewley model. We set the initial distribution equal to the steady-state distribution of the truncated model. As the tax path is exogenous, the model simulation does not require any optimization and follows the same lines as the transition approach, but with a time-varying, deterministic tax path. For each value of κ , we compute the aggregate welfare in the economy, and we report the results in Figure 3. Welfare is expressed as the consumption-equivalent drop in welfare from its optimum. Figure 3 shows that the welfare is maximal for a path that is extremely close to the optimal path $(T_t^p)_t$ ($\kappa = 0$). The welfare computation does not involve the truncation approach and is thus an external validity check. This exercise provides direct evidence that the truncation approach yields a solution that is very close to the Ramsey solution in the full model. This approach is thus well-suited to compute optimal policies in heterogeneous-agent models.

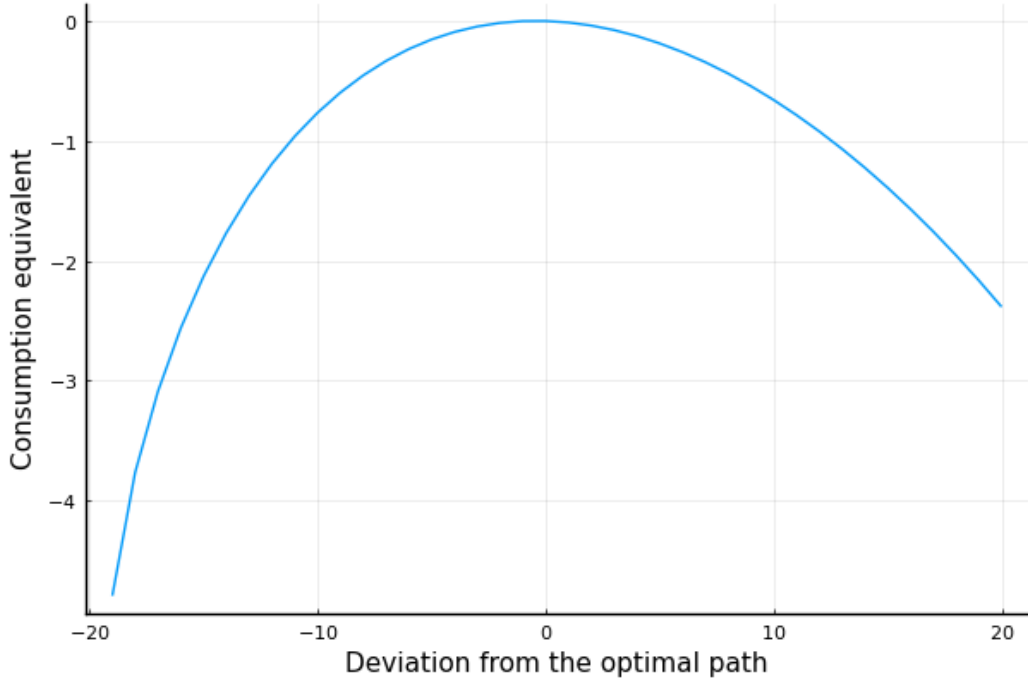


Figure 3: Welfare impact of a deviation from the optimal path. The x -axis plots the value of κ in equation (38) and the y -axis the welfare variation (in percent of consumption equivalent).

5.6 Additional robustness checks

We consider two robustness checks, one on the truncation length N , and another on the concavity parameter of the public good utility function, ϑ . The resulting tax rates are reported in Table 3. We report the following optimal tax rates: the long-run truncation tax rate τ_{∞}^p ; the fixed-point transition tax rate τ^c ; the transition tax rate with low-wealth initial distribution, τ^{low} ; and the transition tax with high-wealth initial distribution, τ^{high} . All these quantities are defined in Section 5.3.

We find that the long-run truncation tax rate is basically unchanged when we move from $N = 5$ to $N = 7$. This shows that the ξ^j efficiently capture the overall heterogeneity, even when the truncation length is not very long. The second column of Table 3 considers a change in the concavity parameter ϑ , and an increase from 24% to 65%. This raises optimal public good provision and thus yields a higher tax rate $\tau_{\infty}^p(\vartheta = 65\%) = 15\%$. The optimal transition tax rate, τ^c , is also higher and remains close to $\tau_{\infty}^p(\vartheta = 24\%)$ of the benchmark case. Finally, the tax rates τ^{low} and τ^{high} are also higher and the ranking $\tau^{low} < \tau^c < \tau^p < \tau^{high}$ is preserved. The corresponding figures are provided in Appendix F.

	Trunc. ($N = 7$)	Trunc. ($\vartheta = 65\%$)
τ^p	8.0%	15%
τ^c	7.8%	14.45%
τ^{low}	6.4%	13.92%
τ^{high}	8.45%	15.40%

Table 3: Robustness checks.

5.7 Proportional labor income tax

We now report the results in an economy where public spending is not funded by a lump-sum tax, but by a proportional labor income tax. We denote this tax by τ_t^L . More formally, the agent's budget constraint (5) becomes:

$$c_t^i + a_t^i = (1 + r_t)a_{t-1}^i + (1 - \tau_t^L)w_t y_t^i,$$

instead of $c_t^i + a_t^i = (1 + r_t)a_{t-1}^i + w_t y_t^i - T_t$. Since aggregate labor is \bar{L} , the government budget constraint becomes:

$$G_t = \tau_t^L w_t \bar{L}.$$

We use the same calibration as in Section 5, and in particular set the steady-state value τ^L such that public spending to GDP is equal to 8%. This corresponds to $\tau^p = 12.5\%$. For the transition approach, we repeat the exercise in Section 5.3. We compute the fixed-point transition tax rate, which is found to be equal to 22.2%. The transition tax rates for the low and high distributions (a decrease and increase of 10% of all initial wealth levels relative to the steady-state distribution) are equal to 20.4% and 23.6%, respectively. Finally in Figure 4 below, we plot aggregate welfare for different values of the tax rate τ^L (and for each tax value, we compute a fixed point on the initial distribution so that it equals the long-run distribution).

The gap between the fixed-point transition tax rate and the truncation tax rate is quite large. This is due to the fact that the labor tax affects more high-productivity agents than low-productivity ones – unlike the lump-sum tax. Thus, the planner is tempted to raise the tax in the short run in order to extract resources from high-productivity agents and finance public expenditures that are valued equally by all agents. The time inconsistency thus appears to be more severe in this environment than in the lump-sum tax environment. Lump-sum taxation therefore appears to generate less distortion.




Figure 4: Welfare with transition approach and fixed-point initial distribution computed as a function of the labor tax rate.

6 Conclusion

We focus on the resolution of a Ramsey program with commitment in a heterogeneous-agent model, where the planner finances a public good by lump-sum taxes. Since the Ramsey solution of the full economy cannot be computed exactly, we focus on two reformulations of the Ramsey program. The transition approach reduces the problem’s dimensionality by assuming that the planner’s instrument path is constant. We also propose a modification of the current technique to avoid the influence of initial distributions on the optimal value of the instrument. The truncation approach solves the Ramsey program of a so-called truncated model, which approximates the full model and features a finite state space. Because the latter method is recently developed, we explain in detail the algorithm and the implementation strategy, which relies on the FOCs of the truncated Ramsey program. We provide direct and indirect evidence that the solution of the truncation approach is accurate and close to the Ramsey solution of the full model. Solving the Ramsey program in the truncated model is thus a valid approximation, even for small truncation lengths. However, despite the success of the truncation approach in the setting of this paper, it is possible that in some environments, the length of the truncation required for correct accuracy would imply a high-dimensional problem.

We summarize the characteristics of the two approaches in Table 4.

Declarations of interest: None.

Characteristic	Transitions	Truncation
Instrument path	constant	time-varying
Long-run instrument value	dependent on initial distribution	indep. of initial distribution
Initial distribution	flexible	flexible but difficult when initial and long-run distributions differ a lot
Aggregate shocks	cannot be accommodated	can be accommodated (via perturbation method)

Table 4: Summary of the two approaches.

Appendix

A Computing the FOCs of the Ramsey program

The Ramsey program (10)–(15) can be written using two instruments only: savings $(a_t^i)_i$ and the lump-sum tax T_t . We also incorporate the Euler equation into the planner's objective. Recalling that the Euler equation Lagrange multiplier is $\beta^t \lambda_t^i$, the Ramsey program (10)–(15) can equivalently be expressed as follows:

$$\begin{aligned} \max_{(T_t, (a_t^i)_{i \in I})_{t \geq 0}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\int_i u(c_t^i) \ell(di) + v(T_t) \right) \\ - \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \int_i (\lambda_t^i - \lambda_{t-1}^i (1 + F_K(\int_i a_{t-1}^i \ell(di), \bar{L}))) u'(c_t^i) \ell(di), \end{aligned} \quad (39)$$

where: $c_t^i = (1 + F_K(\int_i a_{t-1}^i \ell(di), \bar{L})) a_{t-1}^i - a_t^i + F_L(\int_i a_{t-1}^i \ell(di), \bar{L}) y_t^i - T_t$.

FOC with respect to a_t^i . Computing the derivative of (39) with respect to a_t^i yields:

$$\begin{aligned} 0 = & \left(u'(c_t^i) - (\lambda_t^i - \lambda_{t-1}^i (1 + r_t)) u''(c_t^i) \right) \frac{\partial c_t^i}{\partial a_t^i} \\ & + \beta \mathbb{E}_t \left[\int_j \left(u'(c_{t+1}^j) - (\lambda_{t+1}^j - \lambda_t^j (1 + r_{t+1})) u''(c_{t+1}^j) \right) \frac{\partial c_{t+1}^j}{\partial a_t^i} \ell(dj) \right] \\ & + \beta \mathbb{E}_t \left[\int_j \lambda_t^j F_{KK,t} u'(c_{t+1}^j) \ell(dj) \right], \end{aligned}$$

where $F_{KK,t} = F_{KK}(K_t, \bar{L})$, and similarly for $F_{KL,t}$. Observe that we have:

$$\frac{\partial c_t^i}{\partial a_t^i} \frac{\partial c_{t+1}^j}{\partial a_t^i} = -1, \text{ and } = (1 + r_{t+1}) 1_{i=j} + a_t^j F_{KK,t} + F_{KL,t} y_{t+1}^j.$$

With definition (16) of ψ_t^j , this yields:

$$\psi_t^i = \beta \mathbb{E}_t \left[(1 + r_{t+1}) \psi_{t+1}^i \right] + \beta \mathbb{E}_t \left[\underbrace{\int_j \psi_{t+1}^j \left(F_{KK,t} a_t^j + F_{KL,t} y_{t+1}^j \right) \ell(dj)}_{\text{Indirect effect of the change in saving}} \right] \quad (40)$$

$$+ \beta \mathbb{E}_t \left[\underbrace{\int_j \lambda_t^j F_{KK,t} u'(c_{t+1}^j) \ell(dj)}_{\text{Effect of int. rate on saving}} \right], \text{ if } i \text{ is not credit-constrained;}$$

$$\lambda_t^i = 0, \text{ if } i \text{ is credit-constrained.} \quad (41)$$

FOC with respect to T_t . Computing the derivative of (39) with respect to T_t yields:

$$0 = v'(T_t) + \int_i \left(u'(c_t^i) - (\lambda_t^i - \lambda_{t-1}^i (1 + r_t)) u''(c_t^i) \right) \frac{\partial c_t^i}{\partial T_t} \ell(di).$$

Using definition (16) of ψ_t^i and $\frac{\partial c_t^i}{\partial T_t} = -1$, we obtain:

$$v'(T_t) = \int_i \psi_t^i \ell(di).$$

B Complete market economy

The complete-market economy is used as a benchmark to identify the effects of heterogeneity. We start with the first-best allocation, which maximizes aggregate welfare subject to a given level of initial capital K_{-1} and to the economy-wide resource constraint. Let C_t denote the consumption of the representative agent; the economy-wide resource constraint can be written as:

$$C_t + G_t + K_t = F(Z_t, K_{t-1}, \bar{L}) + K_{t-1},$$

which is obviously similar to equation (8) in the general case. The first-best allocation is determined as the solution of the following program:

$$\max_{(K_t, C_t, G_t)_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (u(C_t) + v(G_t)) \right], \quad (42)$$

$$K_t + C_t + G_t = F(Z_t, K_{t-1}, \bar{L}) + K_{t-1}, \quad (43)$$

where we use the governmental budget constraint $G_t = T_t$ and K_{-1} is given. The two FOCs of the first-best program can be written as:

$$u'(C_t) = \beta \mathbb{E}_t \left[(1 + F_K(Z_{t+1}, K_t, \bar{L})) u'(C_{t+1}) \right], \quad (44)$$

$$v'(G_t) = u'(C_t). \quad (45)$$

Equations (44) and (45) together with budget constraint (43) determine a dynamic system in $(C_t, K_t, G_t)_{t \geq 0}$ for a given initial capital K_{-1} and characterize the first-best allocation. The first-best allocation can easily be decentralized by setting the following prices: $r_t = F_K(Z_t, K_t, \bar{L})$ and $w_t = F_L(Z_t, K_t, \bar{L})$.

In that case, the individual budget constraint can be written as:

$$C_t + K_t = K_{t-1} + F_K(Z_t, K_t, \bar{L})K_{t-1} + F_L(Z_t, K_t, \bar{L})\bar{L} - T_t, \quad (46)$$

where we use the financial market and labor market clearing conditions. Combining (46) with the constant returns to scale property of the production function and the governmental budget constraint implies that the individual budget constraint is identical to resource constraint (43). Therefore, because the representative agent is endowed with the whole amount of capital at the initial date and because the first-best FOC (44) is identical to the Euler equation of the representative agent, a competitive allocation in which the fiscal policy is the same as the first-best will be identical to the first-best allocation.²⁰ The Ramsey planner can thus implement the first-best allocation by choosing a fiscal policy according to FOC (45).

To draw the parallel with the Lagrangian approach we used in the general case, observe that in the absence of heterogeneity, equation (40) simplifies into a linear equation in λ_{t-1}^{CM} and λ_t^{CM} with no other terms, with $\lambda_t^{CM} = 0$ as a unique solution.²¹ Intuitively, the agents' Euler equation is not a constraint for the planner because it corresponds to the first-best intertemporal allocation of capital. Consequently, the Lagrange multiplier on the Euler equation is null and the Lagrangian approach implies $\psi_t^{CM} = u'(C_t) = v'(G_t)$, such that FOCs (17) and (45) are actually identical.

A direct consequence is that in our setup, with an instrument that has only within-period effects, there is no time-inconsistency issue in the complete market setup.

²⁰Formally, the two dynamic systems will have the same dynamic equations and the same initial conditions. They will therefore coincide at each date.

²¹Indeed, the term in $F_{KK,t}a_t^{CM} + F_{KL,t}y_{t+1}^{CM}$ would be zero, and after using the Euler equation to simplify further, equation (40) could be written as $A_t\lambda_t^{CM} + B_{t-1}\lambda_{t-1}^{CM} = 0$ for non-zero coefficients A_t and B_{t-1} . With $\lambda_{-1}^{CM} = 0$, this implies by induction $\lambda_t^{CM} = 0$ at all dates.

C Truncated model

The aggregation of the equations characterizing the individual model implies that the full dynamics of the truncated model can be written as follows:

$$v'(T_t) = \sum_{k=1}^{N_{tot}} S_k \psi_{k,t}, \quad (47)$$

$$K_t = \sum_{k=1}^{N_{tot}} S_k a_{k,t}, \quad (48)$$

$$k = 1 \dots N_{tot} : \tilde{a}_{k,t} = \frac{1}{S_k} \sum_{k'=1}^{N_{tot}} S_{k'} \Pi_{k',k} a_{k',t-1}, \quad (49)$$

$$c_{k,t} + a_{k,t} = (1 + r_t) \tilde{a}_{k,t} + w_t y_0, \quad (50)$$

$$\tilde{\lambda}_{k,t} = \frac{1}{S_{k,t}} \sum_{k'=1}^{N_{tot}} S_{k'} \Pi_{k',k} \lambda_{k',t-1}, \quad (51)$$

$$\psi_{k,t} = \xi_k^1 u'(c_{k,t}) - (\lambda_{t,k} - \tilde{\lambda}_{t,k} (1 + r_t)) \xi_k^2 u''(c_{k,t}). \quad (52)$$

$$\xi_k^1 u'(c_{k,t}) = \beta \mathbb{E}_t \left[(1 + r_{t+1}) \sum_{k'=1}^{N_{tot}} \Pi_{k,k'} \xi_{k'}^1 u'(c_{k',t+1}) \right] + \nu_{k,t}, \quad (53)$$

$$k \notin \mathcal{C} : \psi_{k,t} = \beta \mathbb{E}_t \left[(1 + r_{t+1}) \sum_{k'=1}^{N_{tot}} \Pi_{k,k'} \psi_{k',t+1} \bar{y}^N \right] \quad (54)$$

$$\begin{aligned} & + \beta \mathbb{E}_t \left[\sum_{k'=1}^{N_{tot}} S_{k'} \psi_{k',t+1} (\tilde{a}_{k',t+1} F_{KK}(K_t, L) + y_{k'} F_{LK}(K_t, L)) \right], \\ & + \beta F_{KK}(K_t, L) \mathbb{E}_t \left[\sum_{k'=1}^{N_{tot}} S_{k'} \tilde{\lambda}_{k',t+1} \xi_{k'}^1 u'(c_{k',t+1}) \right], \end{aligned}$$

$$k \in \mathcal{C}, \quad a_{k,t} = \lambda_{k,t} = 0. \quad (55)$$

D Matrix representation

A very convenient way to express the truncated model allocation involves using matrix notation. This notation is very powerful for computing optimal policies at the steady state, as shown below. We define the following elements:

- $\mathbf{S} = (S_k)_{k=1 \dots N_{tot}}$ is the N_{tot} -vector of sizes;
- $\mathbf{\Pi} = (\Pi_{kk'})_{k,k'=1 \dots N_{tot}}$ is the transpose of the $N_{tot} \times N_{tot}$ matrix of transition probabilities across histories;
- $\mathbf{c} = (c_k)_{k=1 \dots N_{tot}}$, $\mathbf{a} = (a_k)_{k=1 \dots N_{tot}}$, $\tilde{\mathbf{a}} = (\tilde{a}_k)_{k=1 \dots N_{tot}}$ are the N_{tot} -vectors of allocations (consumption, end-of-period, and beginning-of-period savings, respectively);

- $\mathbf{y}_0 = (y_{k,0})_{k=1\dots N_{tot}}$ is the vector of current productivity levels across histories;
- $\boldsymbol{\xi}^j = (\xi_k^j)_{k=1\dots N_{tot}}$ ($j = 0, 1, 2$) are the vectors of residual-heterogeneity parameters defined in equation (23)–(25) (at the steady state);
- \mathcal{C} is the set of indices of credit-constrained history;
- $\mathbf{P} = \text{diag}((p_k)_{k=1\dots N_{tot}})$, with $p_k = 1$ if $k \notin \mathcal{C}$ and $p_k = 0$ if $k \in \mathcal{C}$, is the diagonal $N_{tot} \times N_{tot}$ -matrix; the matrix \mathbf{P} selects the non-constrained histories;
- $\mathbf{D}_{\mathbf{u}'(\mathbf{c})} = \text{diag}((u'(c_k))_{k=1\dots N_{tot}})$ is the diagonal $N_{tot} \times N_{tot}$ -matrix with $u'(c_k)$ on the diagonal for history k , and 0 elsewhere;
- \mathbf{I} is the identity matrix;
- $\mathbf{1}_{N_{tot}}$ is the N_{tot} -vector of 1.

We also introduce the following operations:

- \circ is the term-by-term product of two vectors of the same size, which is another vector of the same size: $\mathbf{x} \circ \mathbf{z} = (x_{y^N}) \circ (z_{y^N}) = (x_{y^N} z_{y^N})$;²²
- \times is the usual matrix product: e.g., for a matrix \mathbf{M} and a vector \mathbf{x} (of length equal to the number of columns of \mathbf{M}), $\mathbf{M} \times \mathbf{x}$ is the vector $(\sum_{k'} M_{kk'} x_{k'})_k$.

We still denote without a sign the usual scalar multiplication, which is assumed to apply to matrix and vectors (e.g., $\lambda \mathbf{M} = (\lambda M_{kk'})_{k,k'}$) and denote component-wise addition with $+$, which is extended to matrices and vectors of the same size (e.g., $\mathbf{x} + \mathbf{z} = (x_k + z_k)_k$). We also keep the same notation for functions that apply element-wise to vectors: $f(\mathbf{x}) = (f(x_k))_k$.

We can rewrite the equations characterizing the steady-state of the truncated economy using this notation. We start with equation (20):

$$\mathbf{S} = \mathbf{\Pi}^\top \times \mathbf{S}, \quad (56)$$

which makes it clear that the vector of sizes, \mathbf{S} , is the eigenvector of matrix $\mathbf{\Pi}^\top$ associated to the eigenvalue 1, where the sum of the eigenvector coordinates is normalized to 1.²³ The vector \mathbf{S} is thus straightforward to compute.

Second, we consider equation (29) for per-capita beginning-of-period wealth $\tilde{\mathbf{a}}$, which yields:

$$\tilde{\mathbf{a}} = (1/\mathbf{S}) \circ (\mathbf{\Pi}^\top \times (\mathbf{S} \circ \mathbf{a})), \quad (57)$$

²²This operation is also known as the Hadamard product.

²³The existence of a positive eigenvector is guaranteed by the Perron-Frobenius theorem for a positive matrix whose rows sum to 1.

where $1/\mathbf{S} = (1/S_k)_k$ is the vector of size inverses and \mathbf{a} is the given vector of end-of-period wealth. Note that if the size of truncated history is $S_k = 0$, we can set $1/S_k = 0$, which with a null-size history yields a null wealth.

Third, at the steady state (where the tax is T_∞), the budget constraint (30) becomes:

$$\mathbf{c} + \mathbf{a} = (1 + r)\tilde{\mathbf{a}} + w\mathbf{y}_0 - T_\infty \mathbf{1}_{N_{tot}}, \quad (58)$$

which allows us to obtain consumption levels using the given vector \mathbf{a} of end-of-period wealth, and the vector of beginning-of-period wealth from equation (57).

Euler equations (7) can be written as follows with matrix notation:

$$\xi^1 \circ u'(\mathbf{c}) = \beta(1 + r)\mathbf{\Pi} \left(\xi^1 \circ u'(\mathbf{c}) \right) + \boldsymbol{\nu},$$

which defines the vector $\boldsymbol{\nu}$.

We then define the matrix $\mathbf{\Pi}^\lambda$:

$$\Pi_{k,k'}^\lambda = S_{k'} \Pi_{k',k} \frac{1}{S_k}, \quad (59)$$

or equivalently, $\mathbf{S} \circ (\mathbf{\Pi}^\lambda \mathbf{x}) = \mathbf{\Pi}^\top (\mathbf{S} \circ \mathbf{x})$ for any vector $\mathbf{x} \in \mathbb{R}^{N_{tot}}$. $\boldsymbol{\lambda}$, $\tilde{\boldsymbol{\lambda}}$, and $\boldsymbol{\psi}$ denote the vectors corresponding to $(\lambda_k)_k$, $(\tilde{\lambda}_k)_k$, and $(\psi_k)_k$, respectively.

Using (59), definitions (33) and (34) for $\tilde{\lambda}_k$ and ψ_k imply in matrix form:

$$\tilde{\boldsymbol{\lambda}} = \mathbf{\Pi}^\lambda \boldsymbol{\lambda}, \quad (60)$$

$$\boldsymbol{\psi} = \xi^1 \circ u'(\mathbf{c}) - \mathbf{D}_{\xi^2 \circ u''(\mathbf{c})} (\mathbf{I} - (1 + r)\mathbf{\Pi}^\lambda) \boldsymbol{\lambda}. \quad (61)$$

We start by expressing the FOC (40) with respect to saving choices, which only holds for unconstrained truncated histories:

$$\begin{aligned} \mathbf{P}\boldsymbol{\psi} &= \beta(1 + r)\mathbf{P}\mathbf{\Pi}\boldsymbol{\psi} + \beta\mathbf{P}\mathbf{1}_{N_{tot}} (\mathbf{S} \circ (F_{KK}(K, L)\tilde{\mathbf{a}} + F_{LK}(K, L)\mathbf{y}))^\top \boldsymbol{\psi} \\ &\quad + \beta F_{KK}(K, L)\mathbf{P}\mathbf{1}_{N_{tot}} \left(\mathbf{S} \circ \xi^1 \circ u'(\mathbf{c}) \right)^\top \mathbf{\Pi}^\lambda \boldsymbol{\lambda}, \end{aligned} \quad (62)$$

where it should be observed that $\mathbf{1}_{N_{tot}} (\mathbf{S} \circ (F_{KK}(K, L)\tilde{\mathbf{a}} + F_{LK}(K, L)\mathbf{y}))^\top$ and $\mathbf{1}_{N_{tot}} \left(\mathbf{S} \circ \xi^1 \circ u'(\mathbf{c}) \right)^\top$ are $N_{tot} \times N_{tot}$ matrices. The pre-multiplication in (62) by the matrix \mathbf{P} is because FOC (40) holds only for unconstrained histories. We define the following two matrices:

$$\mathbf{L}_0 = \mathbf{I} - \beta(1 + r)\mathbf{\Pi} - \beta\mathbf{1}_{N_{tot}} (\mathbf{S} \circ (F_{KK}(K, L)\tilde{\mathbf{a}} + F_{LK}(K, L)\mathbf{y})), \quad (63)$$

$$\mathbf{L}_1 = \beta F_{KK}(K, L)\mathbf{1}_{N_{tot}} \left(\mathbf{S} \circ \xi^1 \circ u'(\mathbf{c}) \right)^\top \mathbf{\Pi}^\lambda, \quad (64)$$

such that FOC (40) becomes:

$$\mathbf{L}_0\boldsymbol{\psi} = \mathbf{L}_1\boldsymbol{\lambda}. \quad (65)$$

Using definition (61) to express $\boldsymbol{\psi}$ using $\boldsymbol{\lambda}$, we deduce:

$$\mathbf{P}(\mathbf{L}_1 + \mathbf{D}_{\xi^2 \circ u''(\mathbf{c})}(\mathbf{I} - (1+r)\boldsymbol{\Pi}^\lambda))\boldsymbol{\lambda} = \mathbf{P}\mathbf{L}_0(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})). \quad (66)$$

For constrained histories, we simply have $\lambda_{y^N} = 0$, or equivalently using matrix notation:

$$(\mathbf{I} - \mathbf{P})\boldsymbol{\lambda} = 0. \quad (67)$$

Combining equations (66) and (67) yields the important result:

$$\boldsymbol{\lambda} = \mathbf{L}_2^{-1}\mathbf{P}\mathbf{L}_0(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})), \quad (68)$$

$$\text{with: } \mathbf{L}_2 = \mathbf{I} - \mathbf{P} + \mathbf{P}(\mathbf{L}_1 + \mathbf{D}_{\xi^2 \circ u''(\mathbf{c})}(\mathbf{I} - (1+r)\boldsymbol{\Pi}^\lambda)). \quad (69)$$

Equation (68) provides a closed-form expression for the vector $\boldsymbol{\lambda}$ as a function of steady-state allocations, through matrices \mathbf{L}_0 , \mathbf{L}_1 , and \mathbf{L}_2 (which only depend on the allocation) of equations (63), (64), and (69). Matrix \mathbf{L}_2 is invertible when $r > 0$.

Finally, we deduce from (65) and (67):

$$\boldsymbol{\lambda} = \mathbf{M}_1(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})),$$

$$\boldsymbol{\psi} = \mathbf{M}_2(\boldsymbol{\xi}^1 \circ u'(\mathbf{c})),$$

$$\text{where: } \mathbf{M}_1 := \mathbf{L}_2^{-1}\mathbf{P}\mathbf{L}_0, \quad (70)$$

$$\mathbf{M}_2 := \mathbf{I} - \mathbf{D}_{\xi^2 \circ u''(\mathbf{c})}(\mathbf{I} - (1+r)\boldsymbol{\Pi}^\lambda)\mathbf{M}_1. \quad (71)$$

E Introducing aggregate shocks

As a further check of the truncation approach, we compare the dynamics of the same economy with aggregate shocks when it is simulated using the truncation approach and using the Reiter method.

We start by introducing an aggregate risk that affects the TFP, denoted by Z_t . The production function (1) becomes: $Y_t = Z_t K_{t-1}^\alpha \bar{L}^{1-\alpha} - \delta K_{t-1}$.

E.1 Computing the full instrument path with aggregate shocks

Algorithm 4 can be adapted to handle aggregate shocks.

Algorithm 5 (Simulating the full instrument path with aggregate shocks) *We consider as given a truncation length, $N > 0$, and a precision criterion, $\varepsilon > 0$.*

1. Steady state. *We use Algorithm 3 to compute the steady-state Ramsey allocation with precision ε and truncation length N .*
2. Truncated model. *The dynamics of the truncated model with truncation length N are characterized by equations (47)–(55) of Appendix C. The within-heterogeneity coefficients and the set of credit-constrained histories remain equal to their steady-state values, (ξ^0, ξ^1, ξ^2) and \mathcal{C} , respectively.*
3. Aggregate shock. *We specify a functional form for the dynamics of TFP.*
4. Initial conditions. *We have two exclusive cases:*
 - (a) Timeless perspective. *We set the initial wealth distribution $(a_{-1,y^N})_{y^N}$ and the initial distribution of Lagrange multipliers $(\lambda_{-1,y^N})_{y^N}$ to their steady-state values.*
 - (b) Time-0 perspective. *The initial wealth distribution can be set arbitrarily and the initial distribution of Lagrange multipliers $(\lambda_{-1,y^N})_{y^N}$ is set to 0.*
5. Simulation. *Based on the steady-state allocation of Step 1, the model equations of Step 2, the aggregate shock of Step 3, and the initial conditions of Step 4, we simulate the path of model variables (including planner’s instruments) using perturbation techniques.*

Algorithm 5 proposes a versatile approach that covers time-0 and timeless approaches. The timeless approach assumes that the economy has converged to its steady state and is then hit by a shock. The planner optimally reacts to the shock by changing the instrument $(T_t)_t$ and letting the economy adapt to the shock and the new instrument path. As stated by McCallum and Nelson (2000), this timeless perspective is the closest notion to “optimal policy making according to a rule”. In the time-0 perspective, the economy is simply assumed to start from date 0 (with or without a shock). The main difference between the two perspectives is the commitment of the planner. In the timeless perspective, the planner is assumed to be constrained by their past commitments (i.e., those of the steady state), whereas in the time-0 perspective there is no past commitment. The truncation approach makes it easy to switch between the two perspectives because it mostly reduces to adapting the initial values of Lagrange multipliers on Euler equations that enter into the Ramsey FOCs (18). The timeless approach involves setting the initial distribution of Lagrange multipliers, as well as that for wealth, to their steady-state values, thereby mimicking the commitment of the planner and the convergence to the steady state. This is Step 4(a) of Algorithm 5. Conversely, the time-0 approach (Step 4(b)) involves setting past Lagrange multipliers to zero, reflecting the absence of past commitment. The initial wealth distribution can be arbitrary, enabling one to study the effect of initial inequality on the optimal instrument path. However, in the remainder of this paper, we neutralize this effect and set Lagrange multipliers to their steady-state values, similarly as in the transition approach.

Both perspectives can accommodate the presence of an aggregate shock (Step 3).²⁴ A standard specification for the dynamics of TFP assumes that the log of TFP follows a standard AR(1) process – as we do in Section E.2 –, but the algorithm can handle more complex processes. The model equations to be simulated are summarized in Appendix C. As for Algorithm 4, the perturbation method can be performed using existing software.

Algorithm 5 involves the same two assumptions as Algorithm 4: (i) constant ξ^j along the transition, and (ii) a constant set of credit-constrained truncated histories. The latter constraint means that aggregate shocks must remain small enough not to affect the set of credit-constrained histories.

E.2 Comparison with Reiter in the timeless perspective

As is standard, we assume that the TFP process follows an AR(1) process, with $Z_t = \exp(z_t)$ and $z_t = \rho_z z_{t-1} + \varepsilon_t^z$, where $\varepsilon_t^z \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_z^2)$. We use the common of values $\rho_z = 0.95$ and $\sigma_z = 0.31\%$ to obtain a deviation of the TFP shock z_t equal to 1% at a quarterly frequency (see Den Haan, 2010, for instance). We then follow Algorithm 4 in a timeless perspective to simulate the economy with aggregate TFP shocks over a simulation of 10,000 periods. This allows us to obtain the paths of the optimal truncation tax $(T_t)_t$ and economic aggregates, including aggregate consumption $(C_t)_t$ and GDP $(Y_t)_t$. We use these data to approximate the optimal tax path using two observable aggregates of the model, capital and GDP. More precisely, we run the following regression:²⁵

$$T_t = T_\infty + a_Y (Y_t - Y) + a_C (C_t - C) + \varepsilon_t^T, \text{ where } \varepsilon_t^T \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_T^2). \quad (72)$$

T_∞ , C , and Y are the steady-state values computed in the Bewley model. We obtain the following values: $a_Y = 0.0442$, and $a_C = 0.0672$ with an R^2 equal to 0.9982. This allows us to capture the optimal dynamics of the tax path using only GDP and capital.

The estimated equation (72) is then incorporated into a full-fledged heterogeneous-agent model with an aggregate shock as an exogenous rule. The model is simulated using the Reiter method. Indeed, the tax path is thus considered as exogenous and the model simulation does not involve any optimization, making simulation via the Reiter method possible.²⁶

We compare the outcomes of the truncation approach and Reiter method, and report the impulse response functions (IRFs) for the main variables in Figure 5. The two solutions are labeled “Truncation” and “Reiter.” We also plot the aggregate welfare in the two cases’ economies, using the aggregate welfare, as the percentage change in consumption equivalent.

It can be observed that the two simulation solutions generate very similar results, along the

²⁴The timeless perspective combined with no aggregate shock simply involves constant paths for the instrument and allocation, and is therefore of little interest.

²⁵We have considered more involved regressions, but these give little actual improvement to the fit.

²⁶We implement the Reiter method using 100 wealth bins and idiosyncratic states, and perform a first-order perturbation of the policy rules as a function of the aggregate shock.




Figure 5: Simulated IRFs after a TFP shock simulated using the truncation approach and Reiter method. See the text for the details of the implementations.

tax path (by construction), aggregate quantities (GDP, consumption, and capital), and prices (interest rate and wages). A small difference between the two simulations arises from the fact that the rule estimated in equation (72) is close, but not exactly equal, to the actual dynamics of T , which are difficult to capture in the very first periods. We complement the findings of Figure 5 by reporting in Table 5 the second-order moments in the two simulations (the Reiter method and truncation approach for $N = 5$). As was the case for the IRFs, the second-order moments are very similar in the two cases.

Methods		Reiter	Trunc.
GDP	Mean	3.793	3.793
	Std/mean (%)	1.288	1.281
C	Mean	2.475	2.475
	Std/mean (%)	0.978	0.979
K	Mean	40.590	40.590
	Std/mean (%)	1.225	1.201
Corr(C, C_{-1})		0.9945	0.9942
Corr(GDP, GDP_{-1})		0.9695	0.9692
Corr(GDP, C)		0.9242	0.9294

Table 5: Moments of the simulated model for different computational techniques.

E.3 Computational results

We provide here additional statistics about the convergence of the truncation approach. The previous section showed that the Reiter model and the truncation approach with $N = 5$ are very close (up to the approximation of the optimal tax function of equation (72)). We here compare the outcomes of the truncation approach with different truncation lengths N . We also report computational times. Results are provided in Table 6.

mN	2	3	4	5	6
Distance d_N	5.5×10^{-4}	4.1×10^{-4}	2.7×10^{-4}	1.3×10^{-4}	-
Comp. time	3s	5s	20s	2.20min	26min

Table 6: Simulation outcomes of the truncation approach for different N . See the text and equation (73) for the definition of the distance d_N .

We apply the truncation approach with truncation lengths from $N = 2$ to $N = 6$. In each case, we simulate the economy for 10,000 periods. We then compare the truncation tax paths for the different truncation lengths. We denote by $(T_{N,t})_{t=1,\dots,10000}$ the simulated truncation tax

path for the truncation length $N = 2, \dots, 6$. We then define the distance d_N as:

$$d_N = \max_{t=1, \dots, 10000} \left(\text{abs} \left(\frac{T_{N,t} - T_{6,t}}{T_{6,t}} \right) \right), \quad (73)$$

which is the maximum relative deviation in tax paths of the truncation approach between arbitrary N and $N = 6$ (which is the maximum truncation length).

First, one can observe that the maximum relative distance d_N is decreasing with N (line “Distance d_N ” of Table 6). An interesting result is that the case $N = 2$ does already a fairly good job and allows one to compute a realistic optimal path for the instrument. Compared to LeGrand and Ragot (2022a), where this result is already discussed, the introduction of multiple ξ^j , instead of a unique set, further improves the accuracy. A proportional deviation of 0.055% (case $N = 2$) implies that the time series for $N = 2$ and $N = 6$ can barely be distinguished.

Second, we report the computing time to simulate the model (Line “Comp. time” of Table 6), once the steady state of the Bewley model (which is the same for all simulations) has been computed.²⁷ For $N = 2$ the simulation of the model for 10,000 periods take less than 3s. The computation times increases very rapidly with N , as the number of truncated histories increases. For $N = 6$, the computation time is 26 minutes.

F Robustness checks

We summarize below the results for an economy with a different value for the curvature of the public good ($\vartheta = 65\%$). In this economy, the optimal value of the tax as a share of GDP is 15%, as computed by the truncation approach. The optimal value of tax-to-GDP computed with the transition approach is 14.45%. Again, if we consider different initial distributions, we will end up with different values for the optimal tax computed with the transition approach. Figure 6 plots welfare as a function of tax-to-GDP.

We also plot in Figure 7 the tax path in the projected economy (after a pure reoptimization shock). As in the baseline calibration, the difference in the two tax rates is due to time-inconsistency.

Finally, we also report a robustness check for the second-order moments in the presence of aggregate risk, as described in Section E.2. Table 7 makes it clear that the truncation length has little impact on second-order moments.

²⁷The computation of the steady state takes less than 2s, using a EGM method with 100 grid points for each productivity level, exponentially spaced. In any case, this computational time is mostly independent on N .

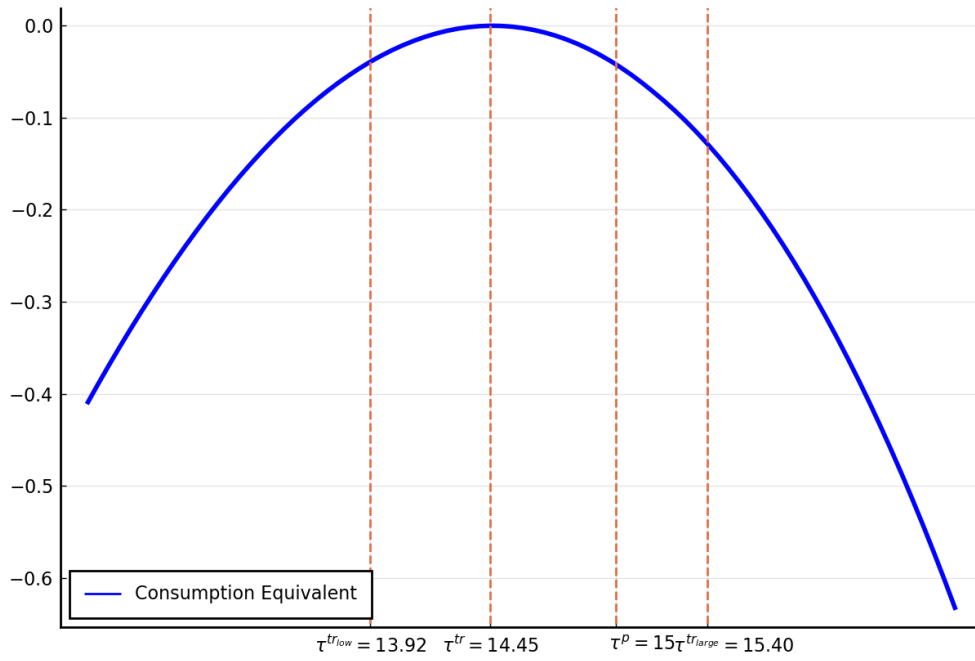


Figure 6: Welfare with transitions and a fixed-point initial distribution, computed as a function of the tax rate.




Figure 7: Dynamic of the tax after a reoptimization shock when the curvature of v is $\vartheta = 65\%$.

Methods		Reiter	Trunc. ($N = 2$)	Trunc. ($N = 3$)	Trunc. ($N = 4$)	Trunc. ($N = 5$)
Economies		(1)	(2)	(3)	(4)	(5)
GDP	Mean	3.793	3.793	3.793	3.793	3.793
	Std/mean (%)	1.288	1.280	1.280	1.281	1.281
C	Mean	2.475	2.475	2.475	2.475	2.475
	Std/mean (%)	0.978	0.980	0.980	0.979	0.979
K	Mean	40.590	40.590	40.590	40.590	40.590
	Std/mean (%)	1.225	1.196	1.198	1.199	1.201
Corr(C, C_{-1})		0.9945	0.9941	0.9941	0.9941	0.9942
Corr(GDP, GDP_{-1})		0.9695	0.9691	0.9691	0.9692	0.9692
Corr(GDP, C)		0.9242	0.9300	0.9298	0.9296	0.9294

Table 7: Moments of the simulated model for different computational techniques.

References

- AÇIKGÖZ, O., M. HAGEDORN, H. HOLTER, AND Y. WANG (2022): “The Optimum Quantity of Capital and Debt,” Working Paper, University of Oslo.
- AÇIKGÖZ, O. T. (2018): “On the Existence and Uniqueness of Stationary Equilibrium in Bewley Economies with Production,” *Journal of Economic Theory*, 173, 18–55.
- ACHARYA, S., E. CHALLE, AND K. DOGRA (2023): “Optimal Monetary Policy According to HANK,” *American Economic Review*, 113, 1741–1782.
- ADJEMIAN, S., H. BASTANI, M. JUILLARD, F. KARAMÉ, J. MAIH, F. MIHOUBI, G. PERENDIA, J. PFEIFER, M. RATTO, AND S. VILLEMOT (2011): “Dynare: Reference Manual, Version 4.6,” Dynare Working Paper 1, CEPREMAP.
- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109, 659–684.
- (1995): “Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting,” *Journal of Political Economy*, 103, 1158–1175.
- AIYAGARI, S. R. AND E. R. MCGRATTAN (1998): “The Optimum Quantity of Debt,” *Journal of Monetary Economics*, 42, 447–469.
- AUCLERT, A., B. BARDÓCZY, M. ROGNLIE, AND L. STRAUB (2021): “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models,” *Econometrica*, 89, 2375–2408.
- BASSETTO, M. (2014): “Optimal Fiscal Policy with Heterogeneous Agents,” *Quantitative Economics*, 5, 675–704.
- BHANDARI, A., D. EVANS, M. GOLOSOV, AND T. J. SARGENT (2021): “Inequality, Business Cycles, and Monetary-Fiscal Policy,” *Econometrica*, 89, 2559–2599.
- BILBIIE, F. (2021): “Monetary Policy and Heterogeneity: An Analytical Framework,” Working Paper, Lausanne University.
- BILBIIE, F. AND X. RAGOT (2021): “Optimal Monetary Policy and Liquidity with Heterogeneous Households,” *Review of Economic Dynamics*, 41, 71–95.

- BOPPART, T., P. KRUSELL, AND K. MITMAN (2018): “Exploiting MIT Shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative,” *Journal of Economic Dynamics and Control*, 89, 68–92.
- CARROL, C. D. (2006): “The Method of Endogenous Grid Points for Solving Dynamic Stochastic Optimization Problems,” *Economic Letters*, 91, 312–320.
- CHALLE, E. (2020): “Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy,” *American Economic Journal: Macroeconomics*, 2, 241–283.
- CHANG, B., Y. CHANG, AND S.-B. KIM (2018): “Pareto Weights in Practice: Income Inequality and Tax Reform,” *Review of Economic Dynamics*, 28, 181–204.
- CONESA, J. C., S. KITAO, AND D. KRUEGER (2009): “Taxing Capital? Not a Bad Idea after All!” *American Economic Review*, 99, 25–48.
- DÁVILA, E. AND A. SCHAAB (2022): “Welfare Assessments with Heterogeneous Individuals,” Working Paper, Yale University.
- DEBORTOLI, D., J. MAIH, AND R. NUNES (2014): “Loose Commitment in medium-scale Macroeconomic Models: Theory and Applications,” *Macroeconomic Dynamics*, 18, 175–198.
- DEN HAAN, W. J. (2010): “Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents,” *Journal of Economic Dynamics and Control*, 1, 79–99.
- DYRDA, S. AND M. PEDRONI (2022): “Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Shocks,” *Review of Economic Studies*, 90, 744–780.
- FERRIERE, A. AND G. NAVARRO (2023): “The Heterogeneous Effects of Government Spending: It’s All About Taxes,” Working Paper, Paris School of Economics.
- GREEN, E. (1994): “Individual-Level Randomness in a Nonatomic Population,” Working Paper, University of Minnesota.
- HEATHCOTE, J. AND H. TSUJIYAMA (2021): “Optimal Income Taxation: Mirrlees Meets Ramsey,” *Journal of Political Economy*, 129, 3141–3184.
- HUGGETT, M. (1993): “The Risk Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17, 953–969.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2018): “On the Distribution of the Welfare Losses of Large Recessions,” in *Advances in Economics and Econometrics: Volume 2, Eleventh World Congress of the Econometric Society*, ed. by B. Honoré, A. Pakes, M. Piazzesi, and L. Samuleson, Cambridge University Press, 143–184.
- LEGRAND, F., A. MARTIN-BAILLON, AND X. RAGOT (2022): “Should Monetary Policy Care About Redistribution? Optimal Fiscal and Monetary Policy with Heterogeneous Agents,” Working Paper, SciencesPo.
- LEGRAND, F. AND X. RAGOT (2022a): “Managing Inequality over the Business Cycle: Optimal Policies with Heterogeneous Agents and Aggregate Shocks,” *International Economic Review*, 63, 511–540.
- (2022b): “Refining the Truncation Method to Solve Heterogeneous-Agent Models,” *Annals of Economics and Statistics*, 146, 65–92.
- MCCALLUM, B. AND E. NELSON (2000): “Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models,” NBER Working Paper 7915, National Bureau of Economic Research.
- MCKAY, A. AND C. WOLF (2023): “Optimal Policy Rules in HANK,” Working Paper, FRB Minneapolis.

- NUÑO, G. AND C. THOMAS (2022): “Optimal Redistributive Inflation,” *Annals of Economics and Statistics*, 146, 3–64.
- ROUWENHORST, G. K. (1995): “Asset Pricing Implications of Equilibrium Business Cycle Models,” in *Structural Models of Wage and Employment Dynamics*, ed. by T. Cooley, Princeton: Princeton University Press, 201–213.
- SMIRNOV, D. (2022): “Optimal Monetary Policy in HANK,” Working paper, Universitat Pompeu Fabra.
- YANG, Y. (2022): “Redistributive Inflation and Optimal Monetary Policy,” Working paper, Princeton University.