Aggregate implications of the progressivity of fiscal transfers: a heterogeneous-agents approach

Tesis para optar al grado de
Magister en Economía

Profesores Guías: Eduardo Engel y Ernesto Pastén
Alumno: Carlos Eggers Prieto

Santiago, Octubre de 2022

Aggregate implications of the progressivity of fiscal transfers: a heterogeneous-agents approach*

Carlos Eggers Prieto†

October, 2022

Check here for the most recent version and replication materials.

Abstract

This paper introduces the progressivity of fiscal transfers into an otherwise standard heterogeneous-agents (HA) model and shows how this feature has a significant effect on the aggregate dynamics following the windfall of fiscal transfers. A larger Marginal Propensity to Consume (MPC) in the lower part of the income and wealth distribution implies that more progressivity generates a stronger response of consumption and a weaker response of investment, which leads to less capital accumulation after the windfall. This paper has two main contributions. First, it studies the progressivity effects using a model in which the distribution of MPCs is endogenous, as opposed to the usual assumptions found in the previous literature that explores transitional dynamics. The second contribution is a relatively simple method to analyze distributional aspects of fiscal policy in continuous-time HA models, that has rich economic interpretation besides progressivity.

Keywords: Fiscal policy, Progressivity, Heterogeneous-agents, Distribution, Marginal Propensities to Consume, Continuous-time.

JEL Codes: E21, E62, H3

*I am truly grateful for the insightful discussions I had with Eduardo Engel, Ernesto Pastén, and Luis Felipe Céspedes, whose advice was not only a huge contribution to this paper but also to my future research and projects. Thanks to my master's classmate and friend Lucas Rosso, who gave me useful comments during the whole process. To my classmates from the thesis workshop and many other people who listened and gave ideas, thanks as well.

†Department of Economics, University of Chile: ceggers@fen.uchile.cl

1 Introduction

Many macroeconomic models have taken into account, or have been specially developed to study fiscal policy, giving it a relevant role in aggregate dynamics. Traditional representative agent (RA) models, usually treat fiscal policy as a combination of taxes, government purchases, and transfers in a context where Ricardian equivalence holds. More recently, there has been a surge of research studying the macroeconomic effects of fiscal policy, when there are non-Ricardian agents, that have larger Marginal Propensities to Consume (MPCs).[1] Non-Ricardian agents, for example, are present in contexts with financial frictions and heterogeneity in wealth across households, for instance, Heathcote (2005). These settings are important to analyze fiscal policy, as they provide a more realistic aggregate dynamic by addressing the fact that different households have different consumption responses to a same amount transfer. For example, Japelli & Pistaferri (2014), Brinca, Holter, Krusell & Malafry (2016) and Broer, Krusell & Öberg (2021) show how fiscal multipliers are heavily influenced by the presence of credit-constrained households.

In the data, fiscal transfers and subsidies represent a large fraction of governments expense, which is often more than half of it. Moreover, the majority of the large economies have increased that proportion in the last 15 years.[2] Fiscal transfers are not flat across the income or wealth distribution in reality, however, despite the relevance of fiscal transfers in government expense, papers including fiscal policy in non-RA settings often assume they are uniformly distributed across agents, with exceptions that have only limited departures from that assumption.[3] This assumption affects the transitional dynamics that follow a fiscal stimulus, because of heterogeneous MPCs. In this paper, we use a very flexible distribution of fiscal transfers in a heterogeneous-agents (HA) setting where MPCs are endogenous, comparing for many different cases how aggregates react differently after a fiscal shock when we depart from the uniform-transfers case. Typically, market incompleteness makes MPCs larger for those households with both low stock of assets and low labor income, and having this in mind, this paper shows specifically how the progressivity degree affects the aggregate dynamic following a fiscal transfers shock. The results suggest a much lower accumulation of capital following a fiscal transfer, when the transfer distribution is more progressive, because of higher MPCs in the lower part of the distribution that induce higher consumption at the beginning of the windfall.[4] Although these results hold in all cases, the initial size and persistence of the transfer affect their magnitude. The magnitude of the results also depends on the part of the

[1] Marginal Propensity to Consume is a measure of a household's consumption sensitivity to changes in their spendable assets. It can be defined as the fraction of resources windfall, that a household would consume in a given period. However, the definition allows for some variations. For example, the model we use in this paper is in continuous-time, in which *a given period* is actually a time interval that can have different duration and in which consumption is not deterministic.

[2] Figure (17) in Appendix (A) shows that the fraction of government expense represented by transfers and subsidies has increased in 8 out of the 10 largest economies for which we have data.

[3] See for example Bassetto (2014), Bhandari, Evans, Golosov & Sargent (2017) and Bhandari, Evans, Golosov & Sargent (2021).

[4] As it is shown in section (5), consumption is only larger at the beginning when the windfall is more progressive, but then faces a reversal that leads to the welfare trade-off discussed in subsection (5.4).

distribution in which the fiscal transfer is more progressive. Concretely, progressivity has a larger effect if it comes from the lower part of the distribution, that is to say when fewer agents at the bottom receive the majority of transfers.

This paper is located in the macroeconomic literature that studies the distributional feature of fiscal transfers. Papers within this literature have limitations that this paper aims to tackle. First, some papers such as Cantore & Freund (2021) and Spector (2020) explore the transitional dynamics and their dependence on the distribution, but having an MPCs heterogeneity that only arises from *ad hoc* assumptions of types of households, lacking thereby the endogenous MPCs heterogeneity that can be achieved with incomplete markets HA models. Besides the lack of endogenous MPCs, assuming types has the problem that in many cases (like the one we explore in subsection (6.1)), it is harder to analyze complex enough distribution schemes as one can see both in real economies and in HA models because the distribution of agents is usually discrete. Second, some other papers, such as Bakış, Kaymak & Poschke (2015) or Heathcote, Storesletten & Violante (2017), take into account the distribution of fiscal policies, but not for short-term fluctuations. Instead, Bakış, Kaymak & Poschke (2015) looks for the welfare effects of a permanent tax policy, and Heathcote, Storesletten & Violante (2017) compares steady-state welfare for different levels of progressivity, in a model with endogenous skill investment, flexible labor supply, and government purchases. Finally, a paper that is closely related to this paper, is Oh & Reis (2012), in which targeted transfers after the 2008 recession are compared with non-targeted transfers. The rule with which government allocates the transfer, however, does not allow for flexible comparisons that take progressivity degrees explicitly into account.

In this paper, I use a rather standard HA model with incomplete markets à la Aiyagari-Huggett-Bewley, to analyze the effect of a fiscal transfer on the evolution of aggregates. The first contribution is to show to what extent aggregates in these kinds of models are sensible not only to the size of a fiscal transfer but also to the distribution of the transfer among different households. Concretely, it is found that in the most progressive cases, capital accumulation after a fiscal shock reaches a peak that is one-third lower than in the uniform-transfers case. The MPCs heterogeneity comes endogenously from the parameters of households' idiosyncratic risk, and that is a feature of the model that makes the heterogeneous agents approach convenient when compared with *types* of agents alternatives, in which income distribution is not determined by parameters that affect MPCs endogenously. Besides not relying on *ad hoc* assumptions to determine MPCs heterogeneity, HA models have the property of having stronger or weaker MPCs heterogeneity depending on the size of the windfalls households receive, which introduces a rich discussion over the size and persistence of fiscal transfers.

As we already mentioned, the use of a HA setting allows to explore fiscal rules that are highly flexible from the distributional point of view, where progressivity can be well defined. For this reason, the second contribution of this paper is a method to include a very flexible fiscal rule into a Heterogeneous Agents Continuous Time (HACT) model, that is otherwise very similar to the one solved in Achdou, Han, Lasry, Lions & Moll (2022). The fiscal transfers distribution has the advantage of allowing us to compare not only levels of progressivity, but also the *source* of the progressivity, that is to say, in which part of the distribution is the progressivity stronger. To achieve this detailed analysis of the fiscal transfer progressivity, I propose using a Beta distribution function to shape the Lorenz Curve of the transfer; as this function has two parameters, a very rich distributional analysis can be done with it. Results show that, as MPCs decrease more in the lower part of the distribution, the trade-off between progressivity and accumulation of capital becomes attenuated when the progressivity comes from the upper part of the distribution.

This paper is organized as follows: section (2) presents the model in which the experiment is analyzed; section (3) explains why MPCs behave differently across the distribution, and why more progressive rules are expected to produce a larger consumption stimulus; section (4) shows the transfer distribution rule, how its parameters should be interpreted, and how to make the computation of the function across income and wealth percentiles, which are endogenous and affected for the rule itself; section (5) show the results of fiscal shocks that have the same aggregate magnitudes, but different distributions; section (6) argues why taking into account fiscal progressiveness in a HA environment gives better insights than a simpler model with types of agents; section (7) concludes.

2 The model

2.1 Steady State

2.1.1 Households

Households are heterogeneous on the state (a, z), where a is the household stock of assets and z is the logarithm of labor productivity. A household is described by some combination within the (a, z) space, that provides inelastically one unit of labor and therefore receives labor income $w \cdot e^z$, and capital income ra. The total income of each household is then $we^z + ra$.[5] Households face non-insurable idiosyncratic risk on their labor productivity, which is determined by an Ornstein-

[5] The percentile of this income is then considered for the simulations in section (5) as the input for the fiscal transfer function we describe in section (4).

Uhlenbeck process:[6][7]

$$\mathrm{d}z_t = -\eta\left(z_t + \frac{\sigma^2}{4\eta}\right)\mathrm{d}t + \sigma^2\mathrm{d}W_t. \tag{1}$$

In each instant t, households decide what fraction of their income will go for consumption c_t, solving the following problem:

$$v(a_0, z_0) = \max_{\{c_t\}_{t\in\mathbb{R}_+}} \left(\mathbb{E}_0 \int_0^\infty \mathrm{e}^{-\rho t}\frac{c_t^{1-\gamma}}{1-\gamma}\mathrm{d}t\right). \tag{2}$$

Subject to equations (1) and:

$$\dot{a}_t = w_t\mathrm{e}^{z_t} + r_t a_t - c_t;\ a_t \geq \bar{a}. \tag{3}$$

Notice that a key element of equation (3) is the borrowing constraint $a_t \geq \bar{a}$, which implies that households might not be able to borrow an otherwise optimal amount of assets. As it will be discussed in section (3), this restriction, when placed above the natural borrowing limit, increases the level of precautionary savings that shape MPCs heterogeneity across the (a, z) space.

Households consumption in each state (a, z), for every given time t can be obtained solving the **Hamilton-Jacobi-Bellman** (HJB) equation:

$$\begin{aligned}\rho v(a,z,t) = \max_{\{c\}}\{&\frac{c^{1-\gamma}}{1-\gamma} \\ &+ (w_t\mathrm{e}^z + r_t a - c)\frac{\partial v(a,z,t)}{\partial a} \\ &- \eta\left(z+\frac{\sigma^2}{4\eta}\right)\frac{\partial v(a,z,t)}{\partial z} + \frac{1}{2}\sigma^2\frac{\partial^2 v(a,z,t)}{\partial z^2} + \frac{\partial v(a,z,t)}{\partial t}\}.\end{aligned} \tag{4}$$

The HJB equation gives a policy function $c(a,z,t) = \left[\frac{1}{v_a(a,z,t)}\right]^{\frac{1}{\gamma}}$ that determines the drift of a in each point of (a, z) for each point in time t. As we know the drift of each point of (a, z), we can tell how the density g of (a, z) space will evolve over time. This can be done with the **Kolmogorov-Fokker-Planck** (KFP) equation, which is:

[6]The Ornstein-Uhlenbeck process is the continuous-time equivalent to the discrete-time auto-regressive process AR(1). The key feature of this stochastic process is its stationarity, which allows for the existence of a stationary distribution. As it will be discussed later, this has important implications for the MPC distribution, as the mean reverting feature of the Ornstein-Uhlenbeck process makes high-income agents expect a drop in their income, back to the economy's average. Some examples of these processes in HA incomplete market models are Guerrieri & Lorenzoni (2017), where the log-income process is an AR(1), and Laibson, Maxted & Moll (2021), where they use a Poisson process that works as a discretized version of the Ornstein-Uhlenbeck process (it has discrete values for income, but the mean, variance and autocorrelation are the same as an Ornstein-Uhlenbeck).

[7]Notice that $\mathbb{E}(z_t) = -\frac{\sigma^2}{4\eta}$ and $\mathbb{V}(z_t) = \frac{\sigma^2}{2\eta}$, so that $\mathbb{E}(\mathrm{e}^{z_t}) = 1$. This simplifies the calculations, as firms will assume that each unit of labor they hire, has a level of productivity equal to 1.

$$
\begin{aligned}
\frac{\partial g(a,z,t)}{\partial t} = & \\
& + \frac{\partial}{\partial a}\left(\left[w_t \mathrm{e}^{z_t} + r_t a - c(a,z,t)\right] g(a,z,t)\right) \\
& - \frac{\partial}{\partial z}\left[\eta\left(z+\frac{\sigma^2}{4\eta}\right) g(a,z,t)\right] + \frac{\partial^2}{\partial z^2}\left[\frac{1}{2}\sigma^2 g(a,z,t)\right].
\end{aligned} \tag{5}
$$

Given a steady state value for wages $w_\infty = w$ and interest rate $r_\infty = r$, imposing $\frac{\partial v(a,z,t)}{\partial t} = 0$ and removing time subscripts from equation (4) gives a solution to the HJB, that maps the consumption $c(a,z)$ of each household in steady state. The stationary distribution is found by solving the case $\frac{\partial g(a,z,t)}{\partial t} = 0$ of the KFP.[8] Now that we have a stationary distribution g^* that maps each point of the (w, r, η, σ^2) space to a density function in the (a, z) space, we can know what will be the **households aggregate stock of assets** in steady state, given (w, r, η, σ^2):

$$
\int_{\bar{a}}^{\infty}\int_{-\infty}^{\infty} a g^*(a, z|w, r)\mathrm{d}z\mathrm{d}a. \tag{6}
$$

The equilibrium in this economy will be a condition that the aggregate stock must satisfy (it will have to match the capital demand of firms).

2.1.2 Firms

Firms solve a static problem, and they demand a level of capital K and employment L, maximizing their profits subject to (w, r, τ, δ), where A is the aggregate productivity and δ is the depreciation rate of capital:

$$
\max_{\{K,L\}} \left[K^\alpha L^{1-\alpha} - wL - (r+\delta)K\right]. \tag{7}
$$

The first-order conditions are:

$$
L = \left[\frac{(1-\alpha)}{w}\right]^{\frac{1}{\alpha}} K; K = \left(\frac{\alpha}{r+\delta}\right)^{\frac{1}{1-\alpha}} L. \tag{8}
$$

Combining the first-order conditions gives the following relationship between r and w:

$$
w(r) = (1-\alpha)\left(\frac{\alpha}{r+\delta}\right)^{\frac{\alpha}{1-\alpha}}. \tag{9}
$$

The first order conditions by themselves only determine $\frac{K}{L}$, but not the *size* of the firm. The size of the firm, in the long run, is determined by imposing $L = 1$, which comes from the inelastically

[8] Solving the KFP is trivial once the HJB is solved. This is because the differential operator $\mathcal{A}$ that defines the HJB is the adjoint differential operator of the KFPs operator. This fact implies that the transition matrix A we use to discretize and solve the HJB is the transpose of KFP's discretized operator.

provided unit of labor of households, that has an average productivity of 1 (see equation (1)). Imposing $L = 1$ would imply:

$$K(r) = \left(\frac{\alpha}{r+\delta}\right)^{\frac{1}{1-\alpha}}. \tag{10}$$

2.1.3 Equilibrium Condition

Recalling that, for each point in (w, r) space we had a stationary distribution that gave an aggregate level of households assets, the goal is to find an interest rate $\bar{r}$ such that:

$$K(\bar{r}) = \int_{\bar{a}}^{\infty} \int_{-\infty}^{\infty} a g^* \left[a, z | w(\bar{r}), \bar{r}\right] \mathrm{d}z \mathrm{d}a. \tag{11}$$

Notice that equation (11) implies that aggregate output and aggregate demand are equal. This is because, the competitive equilibrium of the firm is $Y = \bar{r}K(\bar{r}) + \delta K(\bar{r}) + w(\bar{r})$, where the right hand side is equal to households income plus the depreciation that the firm reinvests in each period.

The steps to get a numerical solution of the equilibrium are:

(1.) Fix a range for r, so $r \in \left[r^{\min}, r^{\max}\right]$, guess some level $r^0 \in \left[r^{\min}, r^{\max}\right]$.

(2.) If $K(r^0) > \int_{\bar{a}}^{\infty} \int_{-\infty}^{\infty} a g^* \left[a, z | w(r^0), r^0\right] \mathrm{d}z \mathrm{d}a$, raise r to $r^1 = \frac{1}{2}(r^0 + r^{\max})$ and reduce the range to $r \in \left[r^0, r^{\max}\right]$. If $K(r^0) < \int_{\bar{a}}^{\infty} \int_{-\infty}^{\infty} a g^* \left[a, z | w(r^0), r^0\right] \mathrm{d}z \mathrm{d}a$, lower r to $r^1 = \frac{1}{2}(r^0 + r^{\min})$ and reduce the range to $r \in \left[r^{\min}, r^0\right]$.

(3.) Repeat step (2.) to get r^2, r^3, etc. until convergence.

2.2 Dynamics

Modelling dynamics requires solving the same PDEs system, but allowing for (r_t, w_t) to be variable over time, and both $\frac{\partial v(a,z,t)}{\partial t}$ and $\frac{\partial g(a,z,t)}{\partial t}$ be different from zero. In this case, we are interested in solving the case of an *MIT* shock, whose details are described in section (5). In such a scenario, the shock is transitory, so the economy starts at the steady state, and converges back to it. The PDEs system is then solved by imposing two border conditions on the equations: for the HJB, in the last period (T)[9], $v(a, z, T)$ must be the same as in the steady state, and then we solve it backwards in time; for the KFP the first-period distribution is set to be equal to the stationary distribution and then we solve it forward in time.[10]

The equilibrium condition, just as above, is given by the balance between the aggregate stock of households' assets and the capital demand of the firms. In the dynamic context, in which the economy deviates from the steady state and then comes back to it, a trajectory $\bar{r}_t$ for interest rates is an equilibrium trajectory if there exists some distribution function $g^*(a, z, t)$ such that $K(\bar{r}_t) =$

[9] In theory, there is no *last* period, as the steady state is only a limit of $t \longrightarrow \infty$. However, by making the simulation large enough in the t space, this is not a problem.

[10] Hence the PDEs system name: *Backward-Forward Mean Field System*

$\int_{\bar{a}}^{\infty}\int_{-\infty}^{\infty} ag^*(a,z,t)\,\mathrm{d}z\mathrm{d}a \;\forall\, t\in\mathbb{R}_+$. If such condition holds, then $Y_t = \bar{r}_t K(\bar{r}) + \delta K(\bar{r}_t) + w(\bar{r}_t)$ and goods market clears, while agents maximize their objective functions.

3 MPCs across the distribution

This paper explores differences in aggregate responses to fiscal transfers due to differences in MPCs across the distribution. Therefore, understanding the different responses of aggregates requires a detailed understanding of how MPCs behave across the distribution and why the model produces that difference. This is a non-trivial question because, in a deterministic non-constrained setting, MPCs are affected neither by a nor z.

In discrete time, MPCs have a straightforward interpretation, which is the fraction of a windfall that individuals consume in the period in which the windfall is received. In a continuous time environment, however, MPCs have a slightly less intuitive interpretation, because there is uncertainty immediately after the windfall is received. MPCs in continuous time, therefore, are interpreted not as the consumption induced by a windfall in a given period, but as the consumption that is **expected** to be induced in a given **interval**. Concretely, the precise formula for MPCs over a given time interval $[0,\tau]$ is:

$$\mathrm{MPC}_\tau(a,z) = \partial_a \mathbb{E}_0\left[\int_0^\tau c(a,z)\mathrm{d}t\right]. \tag{12}$$

Using the Feynman-Kac formula, we know that $\mathbb{E}_0\left[\int_0^\tau c(a,z)\mathrm{d}t\right] = \Gamma(a,z,0)$, where $\Gamma(a,z,t)$ satisfies the following PDE:

$$0 = c(a,z) + s(a,z)\partial_a\Gamma(a,z,t) - \eta\left(z + \frac{\sigma^2}{4\eta}\right)\partial_z\Gamma(a,z,t) + \frac{\sigma^2}{2}\partial_{z^2}\Gamma(a,z,t) + \dot{\Gamma}(a,z,t). \tag{13}$$

On $[\bar{a},\infty)\times\mathbb{R}\times[0,\tau]$, with terminal condition $\Gamma(a,z,\tau) = 0 \;\forall\,(a,z)\in[\bar{a},\infty)\times\mathbb{R}$. To solve equation (13) we use the terminal condition, and compute the rest of the function backwards in time, using sufficiently small steps.

3.1 Analyzing MPCs

Computed MPCs are shown in Figure (1), and have some key features that are worth noting:

(1.) For households with low z, MPCs become larger as they get closer to the borrowing constraint, while for households with high z, MPCs are almost flat across the a distribution.

(2.) MPCs decrease in z, but they decrease more, the closer households get to the borrowing constraint.

(3.) MPCs are larger when there is more idiosyncratic risk.

Figure 1: $\text{MPC}_1(a, z)$ over the (a, z) distribution

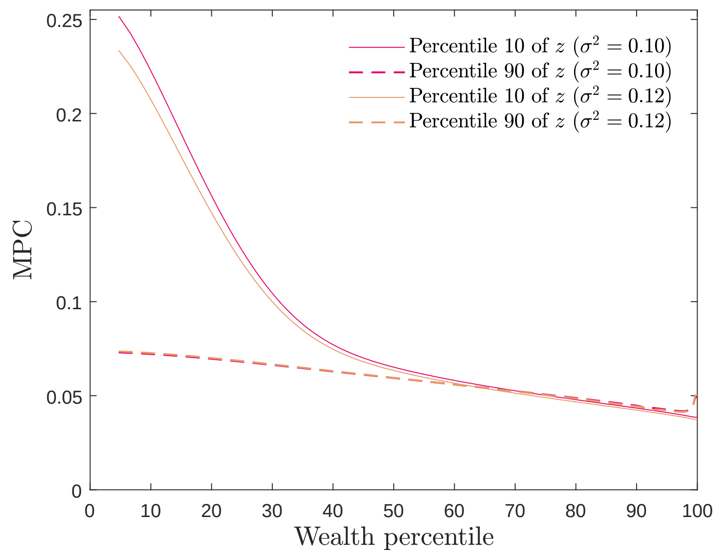

Notice that these facts are not the case in a deterministic non-constrained setting. In these last kinds of settings, households simply consume a fraction $r - \frac{r-\rho}{\gamma}$ of their assets plus income present value. Their consumption is $c(a, z) = \left(r - \frac{r-\rho}{\gamma}\right)\left[a + \int_0^\infty e^{-rt} w e^{-\frac{\sigma^2}{4\eta} + e^{-\eta t}\left(z + \frac{\sigma^2}{4\eta}\right)}\right]$, and when equation (12) is solved for this consumption function, we see that $\text{MPC}_1(a, z) = \frac{\rho + (\gamma-1)r}{\gamma}\ \forall\ (a, z) \in \mathbb{R}^2$[11]. These three facts can be explained because of the precautionary behavior of households across the distribution, that arises when uncertainty and financial restrictions are introduced. These components of the model make households internalize a positive probability of hitting the borrowing constraint in a finite amount of time. The probability of hitting the borrowing constraint increases the closer the household gets to the borrowing constraint, and also decreases as z gets higher. The key to understanding precautionary saving is that even if the household is not constrained at a given instant, and is allowed to borrow, having a positive probability of hitting the constraint at some point in time in the future, makes it rational for the household to consume less than it would in the case where there was no constraint at all. In other words, households behave *as if* they were constrained even if they are not, to avoid being constrained in the near future. This difference between the consumption households get, and the consumption they would get without risk or financial restrictions is the precautionary saving component of consumption. This component varies across the distribution, because as was mentioned, the probability of hitting the constraint on a finite interval of time, changes with both a and z.

[11]See Appendix (B.1) for detailed proof.

As households in a deterministic non-constrained setting consume a fraction $\frac{\rho+(\gamma-1)r}{\gamma}$ of the present value of their income, any windfall is consumed in that exact proportion. However, if a household receives resources when there is a precautionary saving component, it will not only increase the value of its assets but will also decrease the probability of hitting the constraint, thus reducing precautionary saving. This reduction in precautionary saving is what adds a further reaction of consumption to the windfall, which explains why MPCs are not only larger than $\frac{\rho+(\gamma-1)r}{\gamma}$ when there is uncertainty and financial restrictions, but also heterogeneous. As precautionary saving arises with the probability of hitting the constraint, households with a lower probability of hitting the constraint (i.e. high a or z), will have low precautionary saving and therefore their consumption will not behave much differently from the case with no uncertainty and complete markets[12], and their MPCs will therefore be close to $\frac{\rho+(\gamma-1)r}{\gamma}$. MPCs become larger when windfalls imply a reduction in precautionary saving, which happens only when the probabilities of hitting the constraint are high enough. This is what explains facts (1.) and (2.) about MPCs across the distribution. With regards to the fact (3.), what happens is that larger levels of risk imply that the probability of hitting the constraint decreases less as a or z increase, making precautionary saving flatter across the distribution.

Points discussed above, can be seen in Figure (2) and Figure (3), where precautionary saving is computed simply like $c(a,z) - \left[r - \frac{r-\rho}{\gamma}\right]\left[a + \int_0^\infty \mathrm{e}^{-rt} w \mathrm{e}^{-\frac{\sigma^2}{4\eta}+\mathrm{e}^{-\eta t}\left(z+\frac{\sigma^2}{4\eta}\right)}\right]$[13], where $c(a,z)$ is the optimal consumption in the uninsurable idiosyncratic risk setting.

Figure 2: Precautionary saving in a distribution (different zs)

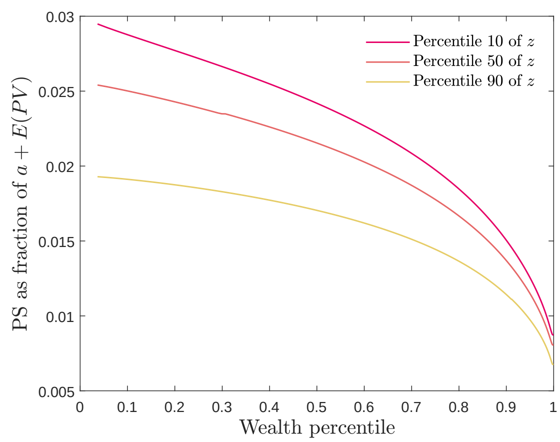

Figure 3: Precautionary saving in a distribution (different σ^2s)

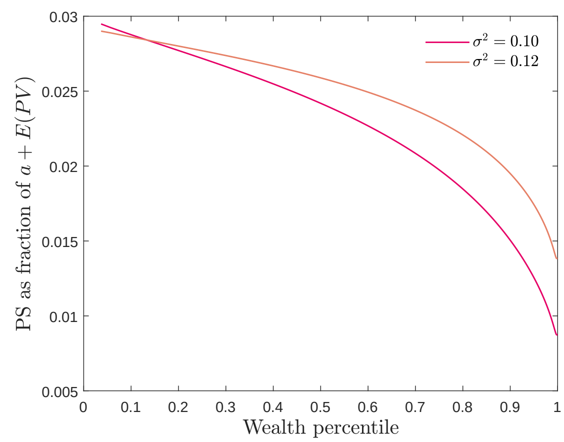

[12] Although the level of precautionary savings does not fall to zero, the ratio between actual consumption and consumption with no uncertainty and complete markets, converges to one, as the in the Carroll (1997) traditional buffer-stock precautionary savings model.

[13] $\mathrm{e}^{-\frac{\sigma^2}{4\eta}+\mathrm{e}^{-\eta t}\left(z+\frac{\sigma^2}{4\eta}\right)}$ is computed numerically. Code is available in the replication materials.

3.2 MPCs under alternative HA models

As we just have seen, MPCs are larger near the borrowing constraint, which makes poorer households, consume a larger fraction of a windfall. However, advances in HA models, have shown some interesting properties, of MPCs when a second asset is added. Concretely, Kaplan & Violante (2014) have proposed the introduction of *wealthy hand-to-mouth* households into theoretical models, to address the fact observed in micro-data, that many households with large stocks of illiquid assets have a nevertheless small amount of liquid assets. These households, despite being rich, are actually constrained, just as poorer households are because illiquid assets have adjustment costs that overcome the benefit of consuming their value. Therefore, including wealthy hand-to-mouth households has relevant theoretical implications, as a fraction of rich households has a consumption behavior that is not different from poor households' behavior. A sophisticated heterogeneous agents model that takes into account the existence of illiquid assets is further developed in Kaplan, Moll & Violante (2018), where wealthy hand-to-mouth households are shown to have a relevant role in monetary policy transmission. Finally, Kaplan & Violante (2022) shows how the introduction of this second asset, solves the problem of one-asset HA models to replicate some key features of the wealth distribution, and the MPCs magnitudes simultaneously.

For the purpose of this paper, which is to take a measure of how different the predictions of a model that departs from the uniform fiscal transfers case can be, it is not necessary to include a second asset. However, the implications of doing so would be worth mentioning for further research. On the one hand, the main trade-off of progressivity and savings would be attenuated, as part of the rich households would also have large MPCs. On the other hand, non-convexities on the cost of adjusting illiquid assets would make rich households more prone to save in form of illiquid assets, which are the assets that are transformed into physical capital, as opposed to liquid assets, which are the kind of assets that would have a larger participation in poorer households savings, and do not translate into greater production.

4 The fiscal transfer function

The shock we will explore later has a fiscal transfer for households, that will be a function of (a, z). In this model, the government makes transfers such that their *Lorenz Concentration Curve* matches a cumulative Beta distribution.

The Beta distribution ($\mathcal{B}$) is defined for a continuous variable $x \in [0, 1]$. The distribution has two parameters (α, β), and $x \sim \mathcal{B}(\alpha, \beta)$ means its density function $f : [0, 1] \longrightarrow \mathbb{R}_+$ is defined by:[14]

$$f(x|\alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{\mathrm{B}(\alpha, \beta)}; \text{ where } \mathrm{B}(\alpha, \beta) = \int_0^1 y^{\alpha-1}(1-y)^{\beta-1}\mathrm{d}y. \tag{14}$$

[14]Function $\mathrm{B}(\alpha, \beta)$ is known as Beta function, hence the name Beta for the distribution.

The cumulative density function $F : [0,1] \longrightarrow \mathbb{R}_+$ is defined by the *regularized incomplete beta function*:

$$F(x|\alpha,\beta) = I_x(\alpha,\beta) \equiv \frac{\mathrm{B}(x,\alpha,\beta)}{\mathrm{B}(\alpha,\beta)}; \text{ where } \mathrm{B}(x,\alpha,\beta) = \int_0^x y^{\alpha-1}(1-y)^{\beta-1}\mathrm{d}y. \tag{15}$$

The transfers distribution planner will choose, will be a distribution within the set for which the *Lorenz Concentration Curve* (LCC) matches the cumulative density function of some Beta distribution. That is to say, the space of transfer functions T over the $we^z + ra$ percentile $p_g(a,z)$ for (a,z) distribution g is:

$$\mathbb{T} := \{T : [0,1] \longrightarrow \mathbb{R}_+ \ / \ \exists\ (\alpha,\beta) \in \mathbb{R}^2_{++} : \ \frac{\int_0^x T(p)\mathrm{d}p}{\int_0^1 T(p)\mathrm{d}p} = I_x(\alpha,\beta)\}. \tag{16}$$

After a short proof[15], we know that $T \in \mathbb{T} \iff T(p|\alpha,\beta) = C\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}$, with $(\alpha,\beta,C) \in \mathbb{R}^3_{++}$, which means the planner must choose the parameters (α,β) to give shape to the LCC, and a scaling parameter C which determines the aggregate size of the transfers.

To gain economic intuition, we redefine parameters to be $\alpha = 1-\theta_1$ and $\beta = \frac{1}{1-\theta_2}$. The advantage is that $(\theta_1,\theta_2) \in (-\infty,1)^2$, and the area under the LCC, grows with each θ_i. As the area under the curve (within the whole [0,1] domain) can be interpreted as the level of *progressivity*[16] of the transfers distribution, we can make the case that larger θ_i implies more progressive distributions.

A given level of progressivity for the transfer distribution can be achieved with different combinations of (θ_1,θ_2). In fact, there is an injective mapping $\theta_2 : (\infty,1) \longrightarrow (\infty,1)$ that we call *iso-progressive curve* (IPC), which assigns a value for θ_2 for any value of θ_1 to maintain a given area under the LCC. In Figure (4), different iso-progressive curves are displayed for three levels of progressivity. Notice, that higher levels of progressivity require either higher θ_1 or higher θ_2.

[15] See Appendix B.2

[16] As progressivity measure $\mathcal{P}$ we will use $\mathcal{P}(T) = 2\int_0^1 \frac{\int_0^x T(p)\mathrm{d}p}{\int_0^1 T(p)\mathrm{d}p}\mathrm{d}x - 1$. This measure is convenient for interpretation purposes, as the most regressive transfer possible has $\mathcal{P}(T) = -1$, the uniform-transfer case has $\mathcal{P}(T) = 0$, and the most progressive case has $\mathcal{P}(T) = 1$. This is progressivity measure that is very similar to the *Kakwani Index* presented in Kakwani (1977), as they are both linear functions of the area under the LCC.

Figure 4: Iso-progressive curves

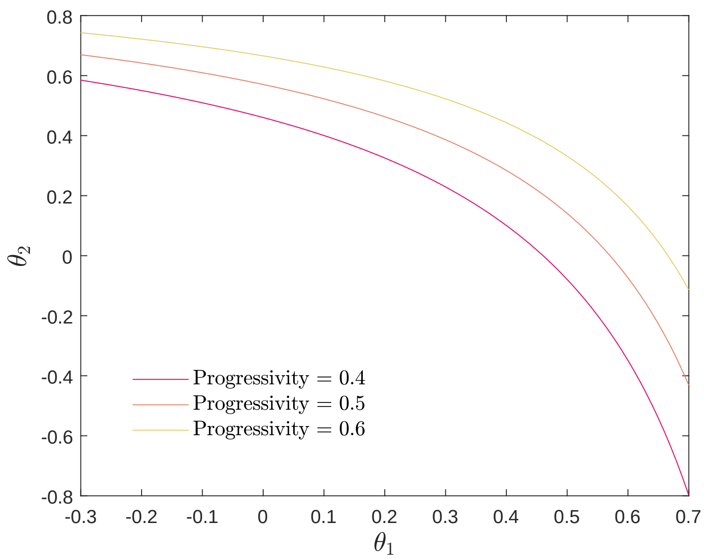

This is economically relevant because a distribution can be progressive for two reasons:

1. The lower tail of the distribution has a high enough concentration of the transfers, which could be interpreted as a *premium* for the lower part of the distribution.

2. The upper tail of the distribution has a low enough concentration of the transfers, which could be interpreted as a *penalty* for the upper part of the distribution.

The relation between the above reasons and the Beta transfers distribution is that, given a level of progressivity, higher levels of θ_1 will concentrate a lot of mass within the lower part of the distribution while keeping the transfer flatter on the rest of the distribution, including the upper part. On the other hand, low levels of θ_1 will keep the transfer flatter on almost the whole distribution, except the upper part, in which mass concentration will be closer to zero.

This intuition can be seen graphically in the following figures. Figure (5) shows how different combinations of (θ_1, θ_2) can have the same area under the LCC, and yet differ strongly on their shapes when we look at Figure (6).

Figure 5: Iso-Progressive transfer functions

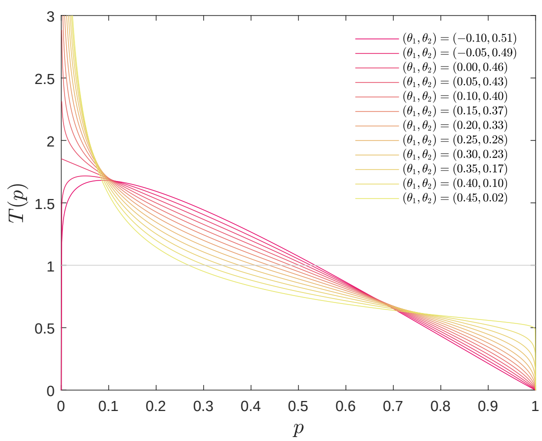

Note: Each transfer function has the same level of progressivity ($\mathcal{P} = 0.3$)

Figure 6: Equal Progressivity LCCs

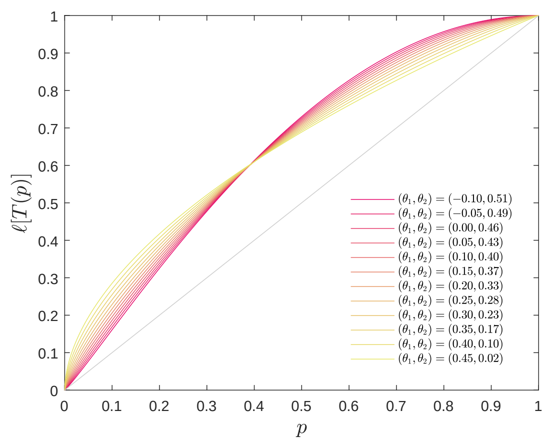

Note: Each Lorenz Concentration Curve has the same level of progressivity ($\mathcal{P} = 0.3$)

Concretely, in Figure (6), the low θ_1 case has a more linear LCC on the bottom of the distribution, which means the progressivity comes primarily from a wealth *penalty*. Meanwhile, the high θ_1 case has a more linear LCC at the top of the distribution, which means the progressivity comes primarily from a poverty *premium*.

This distinction between equally progressive yet differently shaped LCCs is quite important for the progressivity trade-off that motivates this paper. This is because the progressivity trade-off comes from the heterogeneity of MPCs. As MPCs heterogeneity is higher in the lower part of the distribution (see Figure (1)), a natural consequence is that the progressivity trade-off becomes stronger when the fiscal rule is less linear in the lower part of the distribution. In our particular analysis, that would be the case when we have high θ_1. Subsection (5.2) simulates fiscal windfalls that have the same level of progressivity, arising from different combinations of (θ_1, θ_2). Results suggest that the progressivity trade-off becomes in fact stronger when θ_1 becomes larger.

5 Simulation and results

The *experiment* we want to study in this economy, is an *MIT* shock (i.e. an unexpected zero probability event) consisting of a fiscal stimulus transfer financed with an endowment shock, that is not flat across the distribution. The purpose of the experiment is to show how the departure from the uniform distribution assumption of the fiscal stimulus can largely affect the aggregate response following afterwards. The transfers each household receive at each instant t are determined by their percentile of $we^{z_t} + r_t a_t$, following the Beta transfer rule that was detailed in section (4). The transfers received by each household in each instant t are denoted by $F_t(a, z, \theta_1, \theta_2)$. The transfer is transitory, and its aggregate size follows an exponential decay with rate ψ_F. We also simulate the shock in a set-up where the government taxes production and redistributes both in the steady state and during the whole transition back to it. In this alternative set-up, welfare gains from redistribution during the windfall are non-trivial (because redistribution is already happening without the windfall). This gives a welfare trade-off in which an optimal level of progressivity is implied by the fact that larger levels of progressivity cause a lower level of capital and wages for future periods. Simulating windfalls for different levels of θ_1 within this set-up, we calculate its optimal level. Finally, a last variation of the welfare exercise is explored, in which progressivity is costly, so the aggregate size of the transfer decreases with θ_1.

The exponential aggregate size of the stimulus is formalized in equation (17), and the (F_0, ψ_F) combination is subject to a fixed present value $\hat{\mathcal{F}}$, so that equation (18) holds. The interest rate assumed for the government restriction is just an arbitrary calibration of 0.02. This government trade-off between initial size and persistence of the shock, as we show later, has effects on the aggregate response despite being neutral in terms of present value. This is because the borrowing constraint makes households non-Ricardian, and that constraint becomes weaker as the fiscal stimulus becomes larger in $t = 0$.

$$\int_{\bar{a}}^{\infty}\int_{-\infty}^{\infty} F_t(a, z, \theta_1, \theta_2) g(a, z, t)\mathrm{d}z\mathrm{d}a = \mathrm{e}^{-\psi_F t}\int_{\bar{a}}^{\infty}\int_{-\infty}^{\infty} F_0(a, z, \theta_2, \theta_2) g(a, z, 0)\mathrm{d}z\mathrm{d}a. \tag{17}$$

$$\int_0^{\infty} \mathrm{e}^{-rt}\mathrm{e}^{-\psi_F t}\mathcal{F}_0\mathrm{d}t = \hat{\mathcal{F}} \iff \mathcal{F}_0 = (r + \psi_F)\hat{\mathcal{F}}. \tag{18}$$

As the purpose of this paper is not the evaluation of actual counterfactuals in some specific economy, we just use a standard calibration in the literature[17], in which the differences induced by the progressivity level can be seen clearly. Parameters used in all simulations are detailed in Table (1):

Table 1: Fundamental parameters calibration

ρ	γ	η	σ^2	δ	α
0.050	2.000	0.115	0.230	0.070	0.350

5.1 Comparing levels of progressivity

For the first experiment, we choose $\hat{\mathcal{F}}$ to be 0.02 times the steady state product and $\psi_F = 0.6$. Also, the experiment is simulated for different levels of progressivity. The left subfigure in Figure (7) shows the evolution of capital after fiscal stimulus with different progressivities, that go from the uniform case $(\theta_1, \theta_2) = (0.0, 0.0)$ to the most progressive case $(\theta_1, \theta_2) = (0.6, 0.0)$. The subfigure on the right, does the same but changing θ_2, so parameters are $(\theta_1, \theta_2) = (0.0, 0.6)$ in the most

[17] We follow a calibration as used in Achdou et al. (2022) online appendix, except for σ^2, which is taken from Kaplan, Moll & Violante (2018) and η which is chosen to match a Gini coefficient of 0.5 in labor income.

progressive case.[18] Results show that, compared to the uniform case, progressive transfers induce a smaller increase in aggregate capital stock. The difference between the most extreme cases, which are $(\theta_1, \theta_2) = (0.0, 0.0)$ and $(\theta_1, \theta_2) = (0.6, 0.0)$, is quite substantial. Concretely, capital deviates approximately 0.250% from its steady-state level in the uniform case, while it only deviates around 0.163% in the $(\theta_1, \theta_2) = (0.6, 0.0)$ case; which is a third less. This smaller deviation means that, on the aggregate, households save a smaller fraction of the stimulus when it becomes more progressive. This is consistent with the MPCs heterogeneity discussion of section (3), and it is shown more clearly in Figure (8), which shows how the aggregate consumption response is significantly bigger when the stimulus gets more progressive. Before analyzing Figure (8) further, notice that the right side of Figure (7) shows how the effect of progressivity, becomes attenuated when progressivity becomes from θ_2 instead of θ_1.

Figure 7: Capital evolution after fiscal transfer with different progressivities

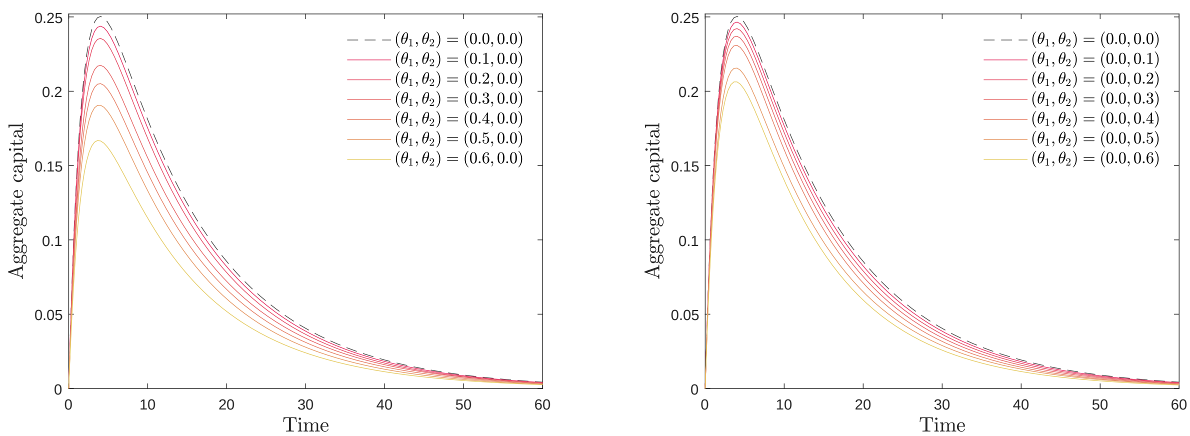

Note: Results are presented as percent deviation from the steady state.

Note: Results are presented as percent deviation from the steady state.

The evolution of consumption shown in Figure (8) makes clear there is a welfare intertemporal trade-off that arises from progressivity. More progressive fiscal transfers, induce a larger response of consumption at the beginning, which is nevertheless reverted at some point in time, because of lower capital accumulation. As more progressive transfers produce a weaker response of capital, wages and labor are therefore lower in the future, making future consumption to be lower as well. This reversal in consumption suggests there could be an optimal progressivity from a welfare point of view, which addresses the fact that larger progressivity is qualitatively equivalent to lower savings in a standard intertemporal consumption problem.

[18] Notice that $(\theta_1, \theta_2) = (0, \bar{\theta})$ has the exact same progressivity as $(\theta_1, \theta_2) = (\bar{\theta}, 0)$.

Figure 8: Consumption evolution after fiscal transfer with different progressivities

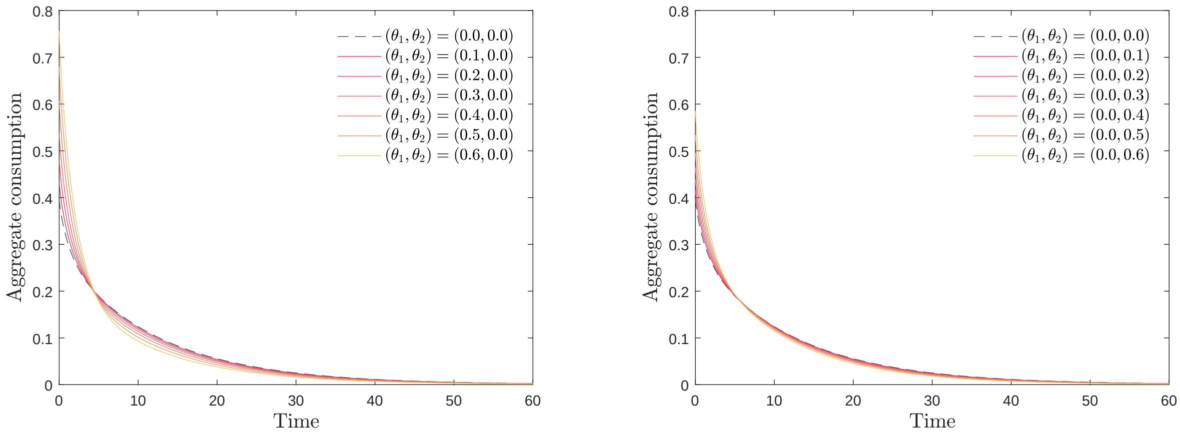

Note: Results are presented as percent deviation from the steady state.

Note: Results are presented as percent deviation from the steady state.

5.2 Comparing iso-progressive transfers

The Beta distribution function for the fiscal transfers discussed in section (4) allows to make transfers of the same aggregate size, persistence, and progressivity but using different combinations of parameters that make the progressivity of the transfer arise from different parts of the agents' wealth and income distribution. While larger θ_1 transfers concentrate a lot of the transfer in the lower part and a flatter amount in the rest of the distribution, larger θ_2 transfers concentrate very little of the transfer on the upper part, and a flatter amount in the rest of the distribution. As seen in section (3), MPCs decrease more in the lower part of both the a and z distribution. This means that the aggregate difference in response that arises from progressivity is stronger when the progressivity comes from the lower, instead of the upper part. For this reason, even though each transfer function in Figure (9) has the same level of progressivity, those in which progressivity comes from the lower part of the distribution (higher θ_2), produce a stronger reaction of consumption at the beginning, and therefore a weaker accumulation of capital.

Figure 9: Capital evolution after iso-progressive fiscal transfer

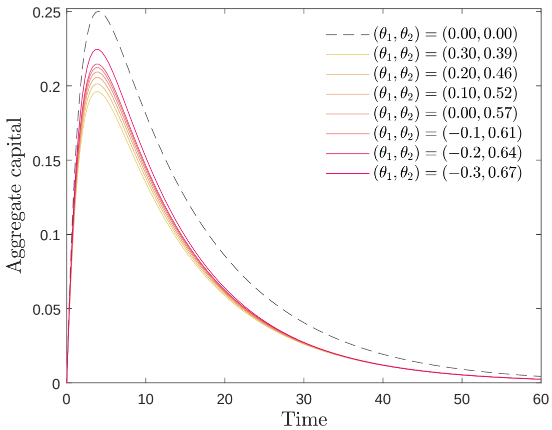

Note: Each simulation (except for the dashed one) has the same level of progressivity ($\mathcal{P} = 0.4$). Results are presented as percent deviation from the steady state.

Figure 10: Consumption evolution after iso-progressive fiscal transfer

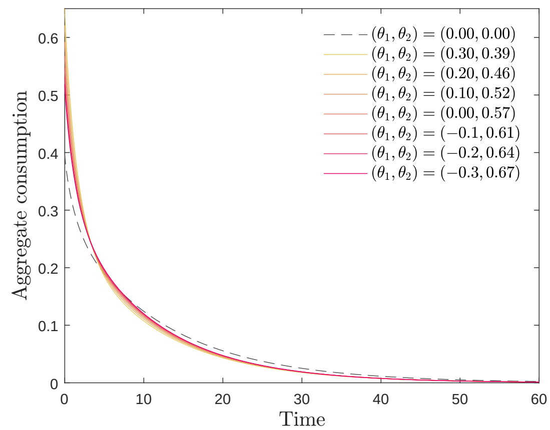

Note: Each simulation (except for the dashed one) has the same level of progressivity ($\mathcal{P} = 0.4$). Results are presented as percent deviation from the steady state.

5.3 Initial size and persistence of the windfall trade-off

The choice of ψ_F is relevant because of two reasons:

- First, the consumption reaction of each household to a windfall, is flatter across the distribution when the windfall is larger. This happens because, as it was discussed in section (3), the precautionary saving component of consumption decreases more on the lower part of the distribution. This implies that the reduction of precautionary saving as fraction of the windfall decreases with the windfall size. Because of this, the fraction of the windfall that is consumed is smaller when the windfall is larger.

- Second, because of consumption smoothing, a larger fraction of a windfall is saved when it is less persistent.

These two reasons imply that aggregate savings will be larger with less persistent windfalls, and therefore less capital will be accumulated, as shown in Figure (11). Notice in Figure (12), that even though savings are larger when persistence is lower, consumption is also larger. This is just because in this setting, lower persistence implies larger $\mathcal{F}_0$.

Figure 11: Capital evolution after transfers with different persistences

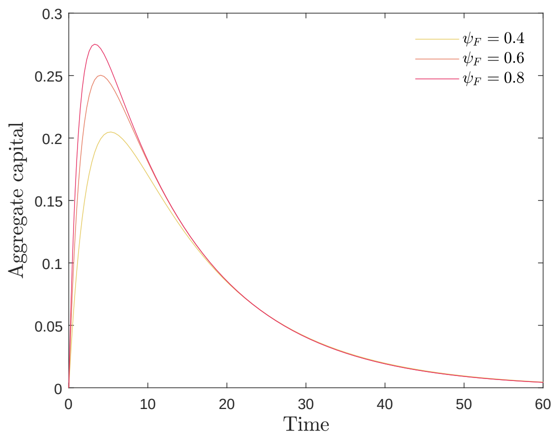

Note: $(\theta_1, \theta_2) = (0, 0)$. Results are presented as percent deviation from the steady state.

Figure 12: Consumption evolution after transfers with different persistences

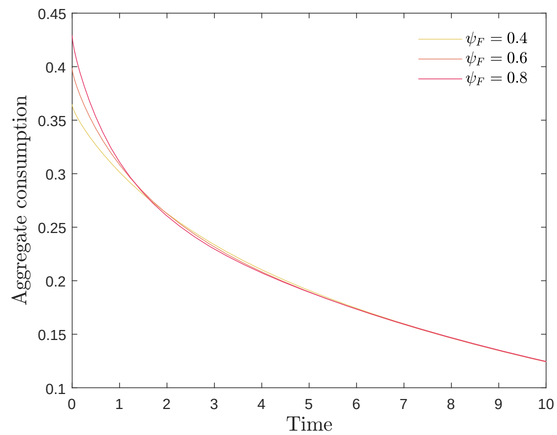

Note: $(\theta_1, \theta_2) = (0, 0)$. Results are presented as percent deviation from the steady state.

As the consumption reaction of households is flatter across the distribution when the windfall is larger, higher persistence of the windfall (and therefore lower $\mathcal{F}_0$) also implies that the progressivity trade-off gets amplified. In other words, high persistence windfalls mean a smaller amount of transfer at the beginning, for which a more heterogeneous response is given by households, and therefore a more relevant difference is observed between a progressive and a uniform case. This can be seen in Figure (13), where capital difference is larger in the higher persistence case, as well as consumption in Figure (14)

Figure 13: Capital evolution difference

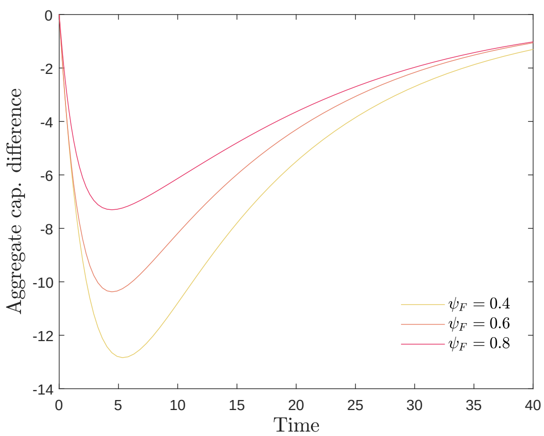

Note: Each line represents the difference between capital with $(\theta_1, \theta_2) = (0.4, 0.0)$ and $(\theta_1, \theta_2) = (0.0, 0.0)$ transfers. Results are presented as percent of fiscal shock present value.

Figure 14: Consumption evolution difference

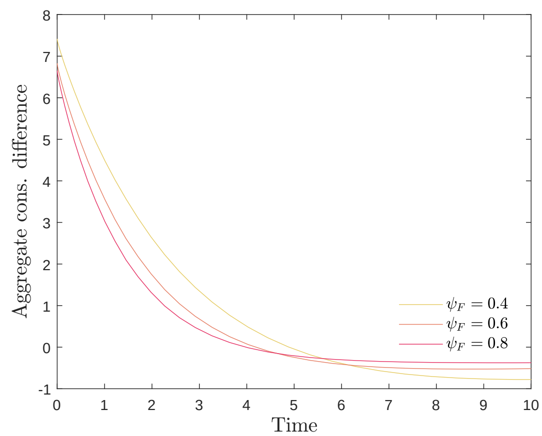

Note: Each line represents the difference between consumption with $(\theta_1, \theta_2) = (0.4, 0.0)$ and $(\theta_1, \theta_2) = (0.0, 0.0)$ transfers. Results are presented as percent of fiscal shock present value.

5.4 Welfare trade-off and optimal progressivity

We have previously shown that aggregate consumption faces a reversal with higher progressivity, meaning there is an intertemporal consumption trade-off. However, in terms of welfare, there is not yet a trade-off, because gains from redistribution always outweigh the gains from capital accumulation. This is because the setting as it has been presented so far, lacks some permanent way of redistribution, making redistribution during the windfall monotonically better in terms of welfare. To maximize discounted aggregate welfare, with an interior solution for progressivity, the setting must be modified to include transfers during both the steady state and the whole transition back to it. If we do so, the direct welfare gains coming from redistribution are at some point countered by the lower level of capital accumulation. For this variation, we repeat the same exercise, but add a tax on production of 0.15, that finances in each point of time a transfer with parameters $(\theta_1, \theta_2) = (0.4, 0.0)$. Within this alternative set-up, we simulate windfalls varying θ_1, and show there is actually an optimal level.

Also, when a cost for redistribution is introduced, and the aggregate size of the windfall is affected by progressivity, the optimal θ_1 can become significantly reduced. In our experiment, we introduce a cost of redistribution that is a fraction $0.2|\theta_1|$ of the aggregate size the transfer would have in a $\theta_1 = 0$ case, which is the same 0.02 production we simulated before. Notice that for this exercise we only move the value of the first parameter of the fiscal transfer (θ_1), and therefore the cost function only depends on that parameter. Results are shown in Figure (15), that maps each level of θ_1 to a measure of welfare. The measure of welfare we use is the percent increase in steady-state consumption that would be equivalent to the windfall in terms of discounted aggregate welfare.

Figure 15: Welfare trade-off

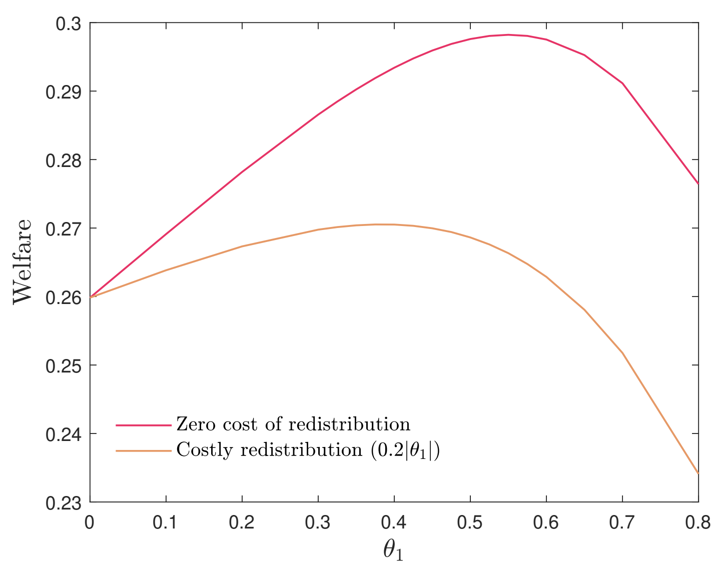

6 Heterogeneous-agents models over *types* models

In this section, we show how some valuable features of heterogeneous agents are lost in models in which different MPCs arise from types. The first lost feature from HA models, is the dependence of the main trade-off intensity, on the size of the windfall. Secondly, HA models have an endogenous distribution of MPCs, that is determined by the combination of (σ^2, η). In types models, the MPCs distribution must be calibrated exogenously, giving lower-income agents a higher MPC just out of *ad hoc* assumptions.

6.1 The TA model

There are two types of agents, that differ both in their labor productivity and in their discount rate. We will call agents with lower z and higher ρ *spenders* and the other type of agents will be called *savers*. Spenders represent a fraction $\Lambda \in (0,1)$ of the population and have labor productivity z_L, while savers have $z_H > z_L$. Spenders have a discount rate ρ_L and savers have a discount rate $\rho_H < \rho_L$. Aggregate labor productivity is assumed to be equal to 1, so $z_H = \frac{1-\Lambda z_L}{1-\Lambda}$. In this economy, steady-state is achieved by imposing $r = \rho_H$. With $r = \rho_H$, spenders have no assets at the steady-state, while the stock of assets of savers is determined by the capital demand of the firms, which is $K = \left(\frac{\alpha}{\rho_H+\delta}\right)^{\frac{1}{1-\alpha}}$.

We choose Λ to match the fraction of the HA economy that has an average of zero assets at steady-state, which is 0.1349. Then we choose (ρ_L, ρ_H) to be consistent with the steady-state MPCs of the HA model, which are 0.1760 and 0.0723. Recalling that, for the riskless models, agents had $\text{MPC}_1(a,z) = \left[r - \left(\frac{r-\rho}{\gamma}\right)\right] \frac{\mathrm{e}^{\frac{r-\rho}{\gamma}}-1}{\frac{r-\rho}{\gamma}}$, discounts rates are calibrated to be $(\rho_L, \rho_H) = (0.2800, 0.0723)$. Finally, $(z_L, z_H) = (0.3370, 1.1034)$.

Within this setting, we introduce a fiscal *MIT* shock, that is analogous to the HA case, in the sense that its aggregate size follows the same evolution, and it is not equal for each type of agent. Particularly, we simulate a fiscal transfer in which a fraction $\Lambda + \theta$ of the aggregate transfer is received by the Λ fraction of agents that are spenders. With that transfer rule, the level of progressivity is $\mathcal{P} = \theta$.[19]

We simulate a fiscal shock with the same size with respect to capital[20] and persistence, and varying θ to match the progressivities in Figure (7). The results of these simulations (see Figure (16)) reveal the same trade-off between progressivity and capital accumulation that we saw in the

[19]Notice that $\Lambda + \theta \in [0,1]$, so $\theta \in [-\Lambda, 1-\Lambda]$. This means that progressivity is bounded by the value of Λ. This does not happen in the HA case, because the distribution is continuous.

[20]Notice, that the absence of idiosyncratic risk in this model leads to a lower level of capital. As the production function remains the same, a 0.02 fraction of the product represents a larger fraction of the capital in the TA setting. Numerical details are available on the `replication materials`.

HA case. However, even with the same calibrations for the distribution of a, z, and MPCs, two key quantitative differences are worth remarking:

1. A same-size uniform fiscal transfer leads to lower levels of capital accumulation in the TA setting.
2. The progressivity and capital accumulation trade-off is stronger in the TA setting.

Figure 16: TA model simulations

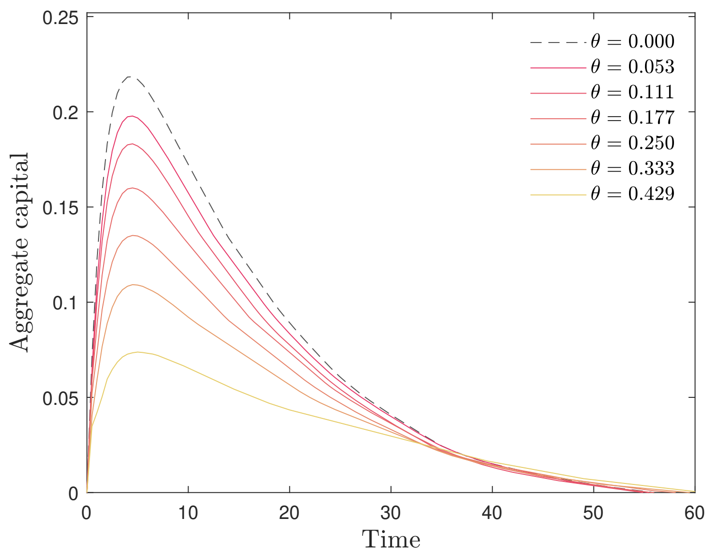

These two differences are explained for the same reason, which is that in the HA setting, as opposed to the TA case, MPCs heterogeneity is neutralized with the size of the transfer, as it was shown in subsection (5.3). In the TA model, MPCs represent not an approximation of the fraction of a windfall that would be consumed, but the actual fraction that will not depend on the size of the windfall. For this reason, the lower part of the distribution in the HA model will have a slightly lower reaction of its consumption compared to the TA model, as the windfall is not a marginal change of disposable wealth, but a measurable increase, whose size does not have any effect on the TA model whatsoever. The consequence of this is that a lower level of capital is accumulated in the TA model, and, as the MPCs distribution does not get neutralized by the windfall size within the TA setting, the trade-off remains more relevant.

Besides the problem we have just exposed, MPCs in the TA setting were calibrated to match our HA calibration but did not arise endogenously from the characteristics of the HA model that could have been taken from micro-data, which are (σ^2, η). This means that, on a TA model, MPCs are chosen with *ad hoc* assumptions that associate lower income agents with an MPC arbitrarily higher than the higher income agents MPC, for no rational reason. There is a huge theoretical advantage of HA modeling, as there is no need to assume an MPCs distribution if the characteristics of the income process that generates them are correctly calibrated.

Finally, the fiscal rule for the distribution of the windfall is highly limited for the TA case, compared to the HA case, where a very flexible rule can be used, that allows, as we have already discussed, for rich analysis about the sources of progressivity within the transfer. The TA case only allows for limited transfer rules because the distribution of agents is discrete, while the Beta function presented in section (4) and used in section (5) has a continuous input.[21]

7 Conclusion

In this paper, we have presented a heterogeneous-agents incomplete markets model in continuous time, that allows for a fiscal transfer, that is non-uniform across the wealth and income distribution, as opposed to what is typically assumed in these kinds of models when fiscal policy is included. Because of the MPCs heterogeneity, that arises from the precautionary behavior of households induced by the idiosyncratic risk and market incompleteness, the departure from the uniform transfers assumption has large consequences on the aggregate response of the economy to the fiscal policy, that are exposed in this paper. We depart from the uniform transfers assumption using a Beta function, that is both flexible and rather simple to compute in continuous time settings. This function allows us to explore different levels of progressivity, but also explore the role that is played by the specific part of the distribution from which that progressivity arises. Having presented the fiscal transfers function, we do several exercises simulating windfalls that are distributed according to that function, and show how they differ substantially from the uniform case. Particularly, the fact that MPCs are larger in the lower part of both the assets and income distribution, induces lower savings from the windfall when transfers are more progressive, so lower levels of capital and thereby lower wages are achieved during the windfall.

This result is more relevant when the progressivity of the windfalls comes from the lower part of the distribution (i.e., households at the bottom of the distribution receive a significantly larger amount of resources, compared with the rest of the distribution, which receives a lower and more flat amount). Because of this, for any given level of progressivity, capital accumulation is smaller if the progressivity comes from the lower part of the distribution. Besides the above results, we find that the size and persistence of the windfall matter from the aggregate point of view, even when the present value of the windfall is maintained. Concretely, the effect of progressivity on the level of capital accumulated becomes attenuated when the windfall is larger at the beginning but less persistent. This happens because MPCs heterogeneity is neutralized with the size of the windfalls. Finally, we show how these differences in the aggregate response of capital imply a welfare trade-off

[21] Technically, that difficulty would be solved by introducing more types than just two of them. However, working with the Beta function as we do in the continuous case would require a large enough number of types, which would undermine the advantage of assuming types of agents, which is simplicity.

from which an optimal level of progressivity can be observed.

We also show the difference between HA models and TA models in which there is no idiosyncratic risk, but heterogeneity comes instead from *ad hoc* assumptions of types of agents (in this case, labor productivity and intertemporal discount rates.). The main difference is that MPCs heterogeneity in the TA model does not become weaker when the windfall becomes larger. For this reason, the main trade-off found in HA is substantially larger, driving to a wrong measurement of the true magnitude of the windfall effects. Also, there is an observational and theoretical advantage for the HA setting, which is the fact that MPCs are not assumed for a level of income or assets but are implied by the parameters of the income process, which can be observed in microdata.

These large differences between HA models compared to TA, and then the implications of departing from the uniform transfer assumption, make it relevant for further macroeconomic research, to take into account more sophisticated rules for fiscal policy, that can give much more precise predictions about its effects.

References

Achdou, Y., Han, J., Lasry, J. M., Lions, P. L., & Moll, B. (2022). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Review of Economic Studies*, 89(1), 45-86.

Bakış, O., Kaymak, B., & Poschke, M. (2015). Transitional dynamics and the optimal progressivity of income redistribution. *Review of Economic Dynamics*, 18(3), 679-693.

Bassetto, M. (2014). Optimal fiscal policy with heterogeneous agents. *Quantitative Economics*, 5(3), 675-704.

Bhandari, A., Evans, D., Golosov, M., & Sargent, T. J. (2017). Fiscal policy and debt management with incomplete markets. *The Quarterly Journal of Economics*, 132(2), 617-663.

Bhandari, A., Evans, D., Golosov, M., Sargent, T. J. (2021). Inequality, Business Cycles, and Monetary-Fiscal Policy. *Econometrica*, 89(6), 2559-2599.

Brinca, P., Holter, H. A., Krusell, P., & Malafry, L. (2016). Fiscal multipliers in the 21st century. *Journal of Monetary Economics*, 77, 53-69.

Broer, T., Krusell, P., & Öberg, E. (2021). *Fiscal Multipliers: A Heterogenous-Agent Perspective* (No. w28366). National Bureau of Economic Research.

Cantore, C., & Freund, L. B. (2021). Workers, capitalists, and the government: Fiscal policy and income (re) distribution. *Journal of Monetary Economics.*

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics, 112(1), 1-55.*

Guerrieri, V., & Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics, 132(3), 1427-1467.*

Heathcote, J. (2005). Fiscal policy with heterogeneous agents and incomplete markets. *The Review of Economic Studies*, 72(1), 161-188.

Heathcote, J., Storesletten, K., Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4), 1693-1754.

Jappelli, T., & Pistaferri, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4), 107-36.

Kakwani, N. C. (1977). Applications of Lorenz curves in economic analysis. *Econometrica: Journal of the Econometric Society*, 719-727.

Kaplan, G., Violante, G. L., & Weidner, J. (2014). *The wealthy hand-to-mouth* (No. w20073). National Bureau of Economic Research.

KAPLAN, G., MOLL, B., & VIOLANTE, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3), 697-743.

KAPLAN, G., & VIOLANTE, G. L. (2022). *The Marginal Propensity to Consume in Heterogeneous Agent Models* (No. w30013). National Bureau of Economic Research.

LAIBSON, D., MAXTED, P., & MOLL, B. (2021). *Present bias amplifies the household balance-sheet channels of macroeconomic policy* (No. w29094). National Bureau of Economic Research.

OH, H., & REIS, R. (2012). Targeted transfers and the fiscal response to the great recession. *Journal of Monetary Economics*, 59, S50-S64.

SPECTOR, M. (2020). Redistributive Fiscal Policy and Marginal Propensities to Consume. *Unpublished Manuscript, MIT.*

A

Figure 17: Transfers and subsidies as (%) of government expense

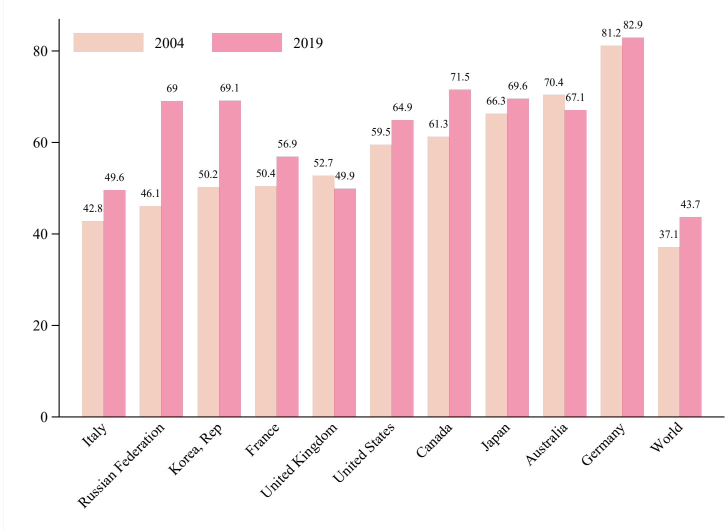

Source: Government Finance Statistics, International Monetary Fund.

B Proofs

B.1

Assuming $L = 1$, without incomplete markets or uncertainty, the optimal consumption for a household given (a, z, w, r), comes just from a deterministic consumption smoothing (CS) problem, which would be:

$$v(a_0, z_0) = \max_{\{c_t\}_{t \in \mathbb{R}_+}} \int_0^\infty \mathrm{e}^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} \mathrm{d}t; \text{ s.t. : } \dot{a}_t = w\mathrm{e}^{z_t} + r a_t - c_t. \tag{19}$$

Where, as there is no stochastic component on the z_t trajectory, z_t just converges to its mean at a constant rate. As we aim to compare consumption with the noninsurable idiosyncratic risk case, we use the same mean for z, which is $-\frac{\sigma^2}{4\eta}$. The trajectory is therefore $z_t = z_0 - \mathrm{e}^{-\eta t}\left(z_0 + \frac{\sigma^2}{4\eta}\right)$.

The Hamiltonian of the problem in current form is:

$$\mathcal{H}(c_t, a_t) = \frac{c_t^{1-\gamma}}{1-\gamma} + \lambda_t \left(w\mathrm{e}^{z_t} + r a_t - c_t\right). \tag{20}$$

First order conditions:

$$\frac{\partial \mathcal{H}(c_t, a_t)}{\partial c_t} = 0 \Longleftrightarrow \left(\frac{1}{c_t}\right)^{\gamma} = \lambda_t \Longrightarrow \dot{\lambda}_t = -\gamma \left(\frac{1}{c_t}\right)^{\gamma+1} \dot{c}_t$$
$$\frac{\partial \mathcal{H}(c_t, a_t)}{\partial a_t} = \rho\lambda_t - \dot{\lambda}_t \Longleftrightarrow (\rho - r)\lambda_t = \dot{\lambda}_t. \tag{21}$$

Combining both FOCs:

$$(\rho - r)\left(\frac{1}{c_t}\right)^{\gamma} = -\gamma \left(\frac{1}{c_t}\right)^{\gamma+1} \dot{c}_t \Longleftrightarrow \dot{c}_t = \left(\frac{r-\rho}{\gamma}\right) c_t. \tag{22}$$

Solving the differential equation is straightforward, and gives the following expression: $c_t = c_0 \mathrm{e}^{\left(\frac{r-\rho}{\gamma}\right)t}$. The transversality condition implies that the present value of consumption must be equal to the value of assets plus the present value of income flow:

$$\int_0^{\infty} \mathrm{e}^{-rt} c_0 \mathrm{e}^{\left(\frac{r-\rho}{\gamma}\right)t} \mathrm{d}t = a_0 + \int_0^{\infty} \mathrm{e}^{-rt} w \mathrm{e}^{z_0 - \mathrm{e}^{-\eta t}\left(z_0 + \frac{\sigma^2}{4\eta}\right)} \mathrm{d}t \equiv a_0 + \mathrm{PV}(z_0)$$
$$\Longrightarrow \frac{c_0}{r - \left(\frac{r-\rho}{\gamma}\right)} = a_0 + \mathrm{PV}(z_0). \tag{23}$$

The trajectory of consumption is thereby:

$$c_t = \left[r - \left(\frac{r-\rho}{\gamma}\right)\right] \left[a_0 + \mathrm{PV}(z_0)\right] \mathrm{e}^{\left(\frac{r-\rho}{\gamma}\right)t}. \tag{24}$$

Plugging the consumption trajectory into equation (12) gives:

$$\mathrm{MPC}_1(a_0, z_0) = \partial_{a_0} \int_0^1 \left[r - \left(\frac{r-\rho}{\gamma}\right)\right] \left[a_0 + \mathrm{PV}(z_0)\right] \mathrm{e}^{\left(\frac{r-\rho}{\gamma}\right)t} \mathrm{d}t$$
$$= \left[r - \left(\frac{r-\rho}{\gamma}\right)\right] \int_0^1 \mathrm{e}^{\left(\frac{r-\rho}{\gamma}\right)t} \mathrm{d}t = \left[r - \left(\frac{r-\rho}{\gamma}\right)\right] \frac{\mathrm{e}^{\frac{r-\rho}{\gamma}} - 1}{\frac{r-\rho}{\gamma}}. \tag{25}$$

Notice that if $r < \rho$, then $\mathrm{MPC}_1(a, z) < \rho$ and $\lim_{r \longrightarrow \rho} \mathrm{MPC}_1(a, z) = \rho$. Finally linearizing for r around ρ gives a pretty much accurate approximation which would be:

$$\mathrm{MPC}_1(a, z) \approx \frac{\rho + (\gamma - 1)r}{\gamma}. \tag{26}$$

B.2

To prove $T \in \mathbb{T} \Longleftrightarrow T(p|\alpha, \beta) = C \frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}$, with $(\alpha, \beta, C) \in \mathbb{R}^3_{++}$, notice that $T \in \mathbb{T}$ means there is some $(\alpha, \beta) \in \mathbb{R}^2_{++}$ such that $\frac{\int_0^x T(p)\mathrm{d}p}{\int_0^1 T(p)\mathrm{d}p} = I_x(\alpha, \beta)$. By definition, $I_x(\alpha, \beta) = \int_0^x \left[\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}\right] \mathrm{d}p$. $\int_0^1 T(p)\mathrm{d}p$ is a constant C (that is interpreted as the aggregate size of the transfer), and therefore, $\int_0^x T(p)\mathrm{d}p = \int_0^x C\left[\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}\right]\mathrm{d}p$, which is equivalent to $T(p) = C\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}$. In the other direction, for every constant C, $T(p) = C\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)} \implies \int_0^1 T(p) = C$. Knowing that, the LCC is $\frac{\int_0^x C\left[\frac{p^{\alpha-1}(1-p)^{\beta-1}}{\mathrm{B}(\alpha,\beta)}\right]\mathrm{d}p}{C} = I_x(\alpha,\beta)$, which means that $T \in \mathbb{T}$, and we conclude the proof.