Liquidity traps, debt relief, and macroprudential policy: a mechanism design approach*

Keshav Dogra†
Columbia University

Job Market Paper

December 5, 2014

Abstract

I present a unified framework to analyze debt relief and macroprudential policies in a liquidity trap when households have private information. I develop a model with a deleveraging-driven recession and a liquidity trap in which households differ in their impatience, which is unobservable. Ex post debt relief stimulates the economy, but anticipated debt relief encourages overborrowing ex ante, making savers worse off. Macroprudential taxes and debt limits prevent the recession, but can harm impatient households, since the planner cannot directly identify and compensate them. I solve for optimal policy, subject to the incentive constraints imposed by private information. Optimal allocations can be implemented either by providing debt relief to moderate borrowers up to a maximum level, combined with a marginal tax on debt above the cap, or with ex ante macroprudential policy - a targeted loan support program, combined with a tax on excessive borrowing. These policies are ex ante Pareto improving in a liquidity trap; in normal times, however, they are purely redistributive. These results extend to economies with aggregate uncertainty, alternative sources of heterogeneity, and endogenous labor supply.

JEL classification: D82, E21, E32, E44, E62

Keywords: Debt relief, liquidity trap, zero lower bound, mechanism design, macroprudential policy

*For the latest version visit http://www.columbia.edu/~kd2338/Dogra_JMP.pdf. A previous draft of this paper circulated under the title "Liquidity traps and incentive-compatible debt relief". I want to thank Ricardo Reis for his advice and support throughout this project. I also thank Jennifer La'O, Tommaso Monacelli, John Moore, Emi Nakamura, Jaromir Nosal, Pablo Ottonello, Rajiv Sethi, Jón Steinsson, Michael Woodford, Pierre Yared, and seminar participants at Columbia University and the University of Edinburgh for their advice and comments. All remaining errors are mine.

†Department of Economics, Columbia University, 1022 International Affairs Building, 420 West 118th Street, New York City, NY 10027. Email: kd2338@columbia.edu. Personal website: http://www.columbia.edu/~kd2338.

1 Introduction

The Great Recession saw an extraordinary contraction in output, employment and consumption, driven in large part by household deleveraging. There are two obvious remedies for a debt-driven recession: prevent borrowing ex ante, or write off debt ex post. A recent theoretical literature makes the case for both policies. Debt relief is an ex post optimal policy in a liquidity trap: transfers to indebted households, who have a high propensity to consume, stimulate demand, raise aggregate income, and benefit everyone, even the households who are taxed to pay for these transfers. Alternatively, macroprudential taxes or limits on borrowing can prevent the overborrowing that leads to a recession in the first place.

But both debt relief and macroprudential regulation face criticisms, which are fundamentally linked to the existence of private information. The literature mentioned above asks whether debt relief is optimal in a liquidity trap, taking as given the distribution of household debt. However, a common concern is that bailouts encourage households to borrow even more ex ante, making the recession deeper. A social planner would like to write off debt for households who would have borrowed anyway, without inducing anyone to borrow more than they would have done in the absence of policy, but this is not possible, since households' propensities to borrow are private information. Thus transfers targeted to debtors inevitably encourage even patient households to take on more debt, making the recession worse. Equally, one criticism of macroprudential limits on borrowing is that they harm households who want to borrow. This is not a concern under full information, since the planner can directly identify and compensate these households, leaving them no worse off. However, if the government cannot observe a household's type, compensating transfers are not possible, and there is an efficiency-equity tradeoff: macroprudential policy prevents a recession, but harms borrowers. Given the constraints imposed by private information, can *any* transfer policies can avert a liquidity trap and make everyone better off?

To answer this question, I take a mechanism design approach to study debt relief and macroprudential policy. I build a model with three key ingredients. First, there is a distribution of households who differ in their propensity to borrow, which can be interpreted as impatience, and which is private information. Second, interest rates are constrained by a zero lower bound (ZLB); when the ZLB binds, output is demand determined. Finally, there is an exogenous contraction in the borrowing constraint - which is perfectly anticipated in the baseline model - which can make the ZLB bind. The borrowing constraint also generates heterogeneity in households' marginal propensity to consume (MPC): highly indebted households will be liquidity constrained, and have a higher MPC than savers, who are not constrained. Thus in this economy, unanticipated transfers from savers to borrowers increase aggregate demand, and this in turn increases aggregate income when the ZLB binds, making all households better off ex post.

However, I show that the concerns raised above are valid: *anticipated* debt relief may not be ex ante Pareto improving, because it encourages overborrowing on both an intensive and an extensive margin. On the intensive margin, if the government does not commit ex ante to limit the scale of debt relief, borrowers take on more debt, since they are now richer in the future.

This means the government must tax savers more heavily to write off borrowers' debt, making savers worse off, relative to a world without debt relief. On the extensive margin, even if the government commits to a cap on debt relief, patient savers may overborrow in order to mimic impatient borrowers and qualify for the transfer. Macroprudential policy also faces constraints. Ex ante debt limits or taxes on borrowing prevent borrowers from taking on too much debt, increase aggregate demand, and mitigate the liquidity trap. However, limits on borrowing may not be ex ante Pareto improving, because they harm households who want to borrow.

To study optimal policy, I consider the problem of a social planner who chooses allocations subject to the ZLB, the borrowing constraint, and private information, which imposes incentive compatibility constraints stating that no household's allocation can be so generous that another household wants to mimic them. By varying the weight the planner puts on each agent's utility, I trace out the constrained Pareto frontier. I prove an equivalence result: any solution to the social planner's problem can be implemented as an equilibrium with transfers that depend on a household's debt level, either at date 1 (ex post redistribution) or at date 0 (ex ante macroprudential policy). Furthermore, efficient allocations can be implemented with particular simple policies. First, they can be implemented with ex post debt relief with a cap. Under such a policy, the government writes off debt up to some maximum level. Above that amount, additional borrowing is taxed, discouraging overborrowing on the intensive margin. Equivalently, efficient allocations can be implemented with ex ante targeted loan support programs, which provide a transfer to households who borrow above some minimum level, coupled with a macroprudential tax on borrowing above that level.

In fact, debt relief with a cap (equivalently, targeted loan support) can be ex ante Pareto improving relative to the competitive equilibrium. In an economy with two types, there always exists a Pareto improving debt relief policy when the ZLB binds in equilibrium. Incentive constraints eventually restrict transfers from savers to borrowers, once these transfers become too large. But in competitive equilibrium, there are no transfers, and incentive constraints are slack: each individual strictly prefers her own allocation. Starting from equilibrium, there is always *some* room to redistribute to borrowers without violating incentive constraints. However, when borrowers are not too impatient relative to savers, the ZLB does not bind, and the competitive equilibrium is constrained efficient. In this case, debt relief (equivalently, targeted loan support) has a purely redistributive role: it implements allocations which are better for borrowers, but worse for savers, relative to the competitive equilibrium.

One concern is that in a two agent economy, it may be too easy to design debt relief programs which induce no extensive margin overborrowing. To address this concern, I also study optimal policy with a continuous distribution of types, using Lagrangian methods similar to those developed by Amador et al. [2006]. In this economy, transfers targeted to highly indebted households always induce some less indebted households to borrow more; optimal policy trades off these distortions against the benefits from debt relief. Nonetheless, I show that an equilibrium in which the ZLB binds is always ex ante Pareto inefficient, and debt relief with a cap (equivalently, targeted loan support) remains ex ante Pareto efficient. However, simple linear policies may not be Pareto improving. In particular, they may make the most indebted borrowers worse off, since they impose a marginal tax on excessive levels of debt.

I then address three further concerns regarding these results. First, while in the baseline model the contraction in borrowing constraints is perfectly anticipated, a more realistic assumption is that this shock only occurs with some probability. In this case, there is an even stronger case for debt relief. The less agents anticipate the crisis, the less incentive concerns restrict debt relief: if the crisis is completely unanticipated, concerns about ex ante incentives vanish completely. Moreover, with aggregate risk and incomplete markets, there is an additional role for debt relief, namely to complete markets and insure agents against a contraction in borrowing constraints. A second concern might be that debt relief or macroprudential taxes might not be desirable if households have different motives for borrowing - perhaps if borrowers borrow because they expect high future income, rather than because they are impatient, writing off their debt will not stimulate demand. I show that even when I extend the model to include alternative motives for borrowing, all the results go through. Thirdly, while my baseline model is an endowment economy, I show that all the results go through in a more standard economy with endogenous labor supply.

The rest of the paper is structured as follows. Section 2 presents the model. Section 3 shows that inefficient overborrowing can occur in equilibrium, and ex post debt relief can be Pareto improving; however, such a policy may not be incentive compatible. Section 4 characterizes constrained efficient allocations in an economy with two types, discusses how they can be implemented, and demonstrates conditions under which debt relief can be ex ante Pareto improving. Section 5 describes how these results generize to a continuous distribution of types, and presents a numerical example. Section 6 discusses the benefits and costs of macroprudential policy, and shows how macroprudential policies can implement optimal allocations. Section 7 considers three extensions: a probability of crisis less than 1, alternative sources of heterogeneity, and endogenous labor supply. Section 8 concludes.

1.1 Related literature

Many recent contributions consider models in which deleveraging leads to a liquidity trap and debt relief is ex post optimal. Eggertsson and Krugman [2012] and Guerrieri and Lorenzoni [2011] were among the first to present models in which an exogenous shock to borrowing constraints causes a recession due to the zero lower bound. My model features the same shock, but asks a different question: what is the optimal policy in response to this shock?

There is a well-established literature on the use of monetary policy at the zero lower bound (the classic papers are Krugman [1998] and Eggertsson and Woodford [2003]; for a recent contribution, see Werning [2012]). There is also a more recent literature on government spending. Eggertsson and Krugman [2012] themselves advocated government purchases, noting that while transfers from savers to borrowers might be stimulative, such transfers are hard to target in practice. Bilbiie et al. [2013b] show that government spending is never Pareto improving in a

Eggertsson and Krugman [2012]-type model, since it hurts savers by lowering interest rates. I study transfer policy, rather than monetary policy or government spending.

My paper is also related to a growing literature on the positive effects of targeted transfers. Oh and Reis [2012] emphasize that government transfers increased much more than government spending during the Great Recession, and provide a model to understand the effects of targeted transfers. McKay and Reis [2013] assess the extent to which automatic stabilizers reduce aggregate volatility. A related empirical literature documents that the marginal propensity to consume varies across households and is correlated with debt (Misra and Surico [2014], Jappelli and Pistaferri [2014], Cloyne and Surico [2013]); Kaplan and Violante [2014] present a model which can match these facts. Relative to these authors, I focus on the effect of transfers in a liquidity trap.

A number of recent contributions discuss the role of targeted transfers in a liquidity trap. Giambattista and Pennings [2013] and Mehrotra [2013] compare the multiplier effects of targeted transfers and government spending in a liquidity trap. Bilbiie et al. [2013a] show that both balanced-budget redistribution and uniform, debt financed tax cuts are expansionary. Bilbiie et al. [2013b] show that debt-financed tax cuts are Pareto improving, as they relax borrowers' credit constraint. Rather than studying transfers in general, I focus on debt relief.

A few recent contributions discuss debt relief. Fornaro [2013] shows that debt relief is expansionary, and may be Pareto improving, at the zero lower bound. Guerrieri and Iacoviello [2013] also show numerically that debt relief can be Pareto improving in a rather different model of housing and collateral constraints. These papers study ex post debt relief, and do not consider whether the anticipation of debt relief can distort incentives ex ante. My contribution, relative to this whole literature, is to consider how ex post redistribution itself distorts incentives ex ante, and to characterize optimal policy taking these distortions into account.

In this sense, my results are most similar to those of Bianchi [2012] who considers optimal bailouts of firms in a small open economy model. Ex post, bailouts relax collateral constraints and increase output, but ex ante, bailouts induce overborrowing: optimal policy combines ex ante macroprudential policy and ex post bailouts. In Bianchi [2012]'s model, the planner can mitigate moral hazard effects by making bailouts conditional on a systemic crisis, rather than individual borrowing. Since firms are identical, there is no need to target particular firms. In my model, the central friction is that debt relief must be targeted at particular households based on observable debt, and savers can mimic borrowers if the bailout is too large. Another difference is that I consider debt relief targeted to households, rather than firms.[1]

Most closely related to my paper are Korinek and Simsek [2014] and Farhi and Werning [2013]. These authors describe in detail the macroeconomic externality that arises in models such

[1]This also differentiates my results from a few recent papers which take a mechanism design approach to study transfer policies targeted at banks. Philippon and Schnabl [2013] study efficient recapitalization in an economy with debt overhang, in which government does not observe banks' asset quality. Tirole [2012] takes a mechanism design approach to analyze how targeted purchases can rejuvenate asset markets. Farhi and Tirole [2012] analyze optimal bailouts when the government cannot perfectly observe a bank's need for liquidity. I also use a mechanism design approach to study targeted transfers, but consider different transfer policies (debt relief, rather than debt-equity swaps or asset purchases), different recipients (households, rather than banks), and a different rationale for intervention (aggregate demand externalities, rather than debt overhang or adverse selection).

as that of Eggertsson and Krugman [2012]: households overborrow and then deleverage, without internalizing that their deleveraging reduces aggregate income. They show that unanticipated ex post redistribution (debt relief) can be Pareto improving. In contrast, I study the design of ex post policies, taking into account how these policies affect borrowing ex ante. The main focus of Korinek and Simsek [2014] and Farhi and Werning [2013] is to consider how ex ante macroprudential policies can prevent overborrowing under full information.[2] I study macroprudential policy under the (realistic) assumption that preferences are private information, so it is not possible to directly target taxes and transfers to households based on their unobservable type.

A vast literature in mechanism design and optimal taxation (Mirrlees [1971]) considers the problem of a social planner who must redistribute among agents whose preferences or skills are private information.[3] The key insight from this literature is that private information reduces the ability of the planner to redistribute. However, it is still possible to achieve some redistribution, by distorting allocations away from the first-best. I apply a mechanism design approach to study optimal redistribution and macroprudential policy in an economy with macroeconomic externalities.[4] While macroeconomic externalities provide a new motive for redistribution, private information still constrains the planner's ability to redistribute.

2 Model

This section presents the baseline model and defines the equilibrium in the absence of policy.

2.1 Agents

Time is discrete and indexed by $t = 0, 1, ...$ There exists a distribution of households with total measure 1. Households have preferences over consumption

$$U(c_0^i, \theta_i) + \sum_{t=1}^{\infty} \beta^t u(c_t^i) \tag{1}$$

where $u' > 0, u'' < 0, \beta \in (0,1), U_c > 0, U_{cc} < 0$. θ_i measures household i's demand for date 0 consumption, with $U_{c\theta} > 0$. Agents with a higher θ_i are more impatient, have a more urgent need for consumption at date 0, and will be borrowers in equilibrium.[5] In all subsequent periods, agents have the same preferences (this ensures that a well-defined steady state exist). In the benchmark model, this is the only source of heterogeneity between agents. For now, I allow θ_i to

[2]There is also a much larger literature on macroprudential policy which focuses on pecuniary externalities, rather than aggregate demand externalities. The mechanism design approach I follow in this paper could also be applied to consider the equity-efficiency tradeoffs associated with these policies.

[3]Formally, my model is closest to the literature on Pareto-efficient income taxation (Werning [2007]); in particular, results for the two-type economy are similar to Stiglitz [1982], who considers a model with two agents.

[4]In this sense, my results are also related to the literature on mechanism design with externalities (Baliga and Maskin [2003]).

[5]For now, I interpret θ as a preference or discount factor shock; in Section 7, I show that it can be reinterpreted in terms of income, so high-θ households borrow because they have temporarily low income at date 0.

have a general distribution function $F(\theta)$. Later, I will focus on two special cases, which I define here.

Definition 2.1. *In the* ***two type economy****, $F(\theta)$ is a discrete distribution with probability mass $f(\theta_S) = f(\theta_B) = 1/2$, $\theta_B > \theta_S = 1$.*

In the ***continuous type economy****, θ has a continuous density $f(\theta)$ with support $[\underline{\theta}, \bar{\theta}]$.*

Agents face a standard budget constraint

$$c_t^i = y_t^i - d_t^i + \frac{d_{t+1}^i}{1 + r_t} \tag{2}$$

where d_{t+1}^i is the face value of debt agent i takes out in period t and promises to repay in period $t+1$, r_t is the real interest rate on a loan between periods t and $t+1$, and y_t^i is i's income. Each agent i can costlessly produce up to y^* of their own differentiated variety of the output good. Each agent's consumption c_t^i is an aggregate of all these varieties $y_t^j, j \in [0, 1]$, providing a motive for trade. y_t^i is not a choice variable of the household: instead, each household takes the demand for its good as given, and produces whatever is necessary to meet demand. Agents have no initial debt:

$$d_0^i = 0, \forall i \tag{3}$$

Agents also face an ad hoc borrowing constraint $\phi_t \geq 0$ in the spirit of Aiyagari [1994]:

$$d_{t+1}^i \leq \phi_t, t = 1, ... \tag{4}$$

Implicitly, ϕ_t reflects the collateralized value of durable goods such as housing, as in Kiyotaki and Moore [1997] (although this is not explicitly modelled here). As in Korinek and Simsek [2014], Eggertsson and Krugman [2012], I model a financial crisis as an exogenous tightening of the constraint. Specifically, households are unconstrained at date 0 ($\phi_0 = \infty$) but the constraint permanently falls to $\phi > 0$ at date 1: $\phi_t = \phi \geq 0, t \geq 1$. In the baseline model, this tightening is perfectly anticipated; in Section 6, I relax this assumption.

2.2 Equilibrium

First, I consider a Walrasian equilibrium, without any frictions (besides the borrowing constraint). I then add the zero lower bound constraint on interest rates. This forces me to modify the standard Walrasian equilibrium concept, as I describe later.

Definition 2.2. *A* ***Walrasian equilibrium*** *is $\{c_t^i, d_t^i, y_t, r_t\}$ such that*

1. *each household i chooses $\{c_t^i, d_t^i\}$ to maximize (1) s.t. (2), (3), (4)*

2. $\int c_t^i \, \mathrm{d}i = y^* = y_t^i, \forall i, t = 0, 1, ...$

I now characterize equilibrium in the two type economy starting in date 1, taking debt at the start of date 1 as given.[6]

Proposition 2.3. *In a Walrasian equilibrium, in the two type economy:*

1. *If $d_1^B \leq \phi$, consumption, debt and interest rates are constant in periods $t \geq 1$: $r_t = r^* := \beta^{-1} - 1$, $c_t^i = y^* - (1-\beta)d_1^i$, $d_t^i = d_1^i$.*

2. *If $d_1^B > \phi$, B is borrowing constrained in period 1: $d_2^B = \phi$. Consumption, debt and interest rates are constant in periods $t \geq 2$: $r_t = r^*, c_t^i = y^* - (1-\beta)\phi$, $d_t^i = \phi$. $r_1 = r(d_1^B)$, implicitly defined by*

$$1 + r_1 = \frac{u'\left(y^* + d_1^B - \dfrac{\phi}{1+r_1}\right)}{\beta u'\left(y^* + (1-\beta)\phi\right)}$$

$r(d_1^B)$ is decreasing in d_1^B, with $r(\phi) = r^*$.

Equilibrium interest rates are decreasing in d_1. If debt is sufficiently low, or the tightening of borrowing constraints is not too severe, the economy immediately converges to a steady state at date 1. If debt is too high, borrowers are no longer able to roll over their debt, and are forced to pay back some debt, temporarily reducing their consumption. In order for markets to clear, savers must consume more at date 1 than they do at date 2. Interest rates must fall to induce them to do so, thus r_1 is a decreasing function of d_1.

Figure 1 illustrates. Aggregate date 1 consumption is decreasing in r_1. For a given interest rate, aggregate consumption is also decreasing in borrowers' debt d_1^B, when debt is high enough that the borrowing constraint binds. When borrowers are liquidity constrained, their marginal propensity to consume is 1, and an increase in debt reduces their consumption one for one. The corresponding increase in savers' net worth increases savers' consumption, but less than one for one, because savers' MPC is much less than 1. Consequently, an increase in debt tends to reduce aggregate consumption. Interest rates must fall to keep aggregate consumption equal to y^*.

I now introduce a constraint on interest rates, $r_t \geq \bar{r}$. For simplicity, in what follows I assume $\bar{r} = 0$. The interest rate r_t required to clear markets may be negative, violating this zero lower bound (ZLB) constraint. In this case, the above equilibrium is no longer possible, and a new equilibrium concept is required. I assume that when the ZLB binds, households cannot sell their whole endowment, and output (i.e., the amount they do sell) is the variable that adjusts to clear markets. Aggregate consumption is still equal to aggregate output. However, aggregate output y_t can fall below potential output y^* when the zero lower bound binds. Formally:

Definition 2.4. *A **ZLB-constrained equilibrium** is $\{c_t^i, d_t^i, y_t, r_t\}$ such that*

1. *each household i chooses $\{c_t^i, d_t^i\}$ to maximize (1) s.t. (2), (3) , (4)*

2. $\int c_t^i \,\mathrm{d}i = 2y_t$

[6] The proof of this Proposition, and all subsequent Propositions, is in the Appendix.

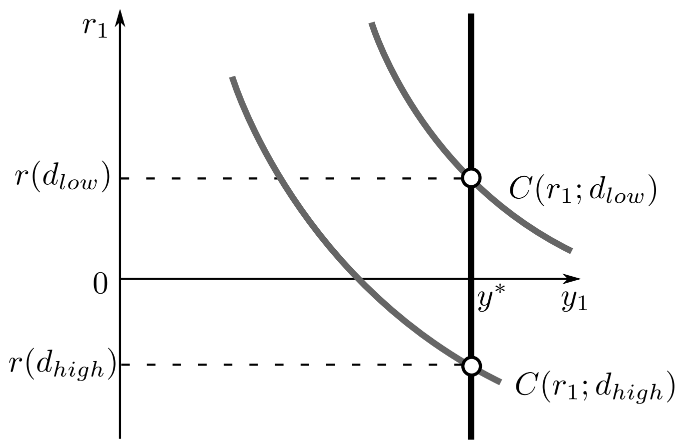

Figure 1: Walrasian equilibrium

3. $r_t \geq 0$, $y_t^i = y_t \leq y^*$, $r_t(y^* - y_t) = 0$

When interest rates can adjust to clear markets, they do, and agents sell all of their endowment. When the ZLB prevents interest rates from falling enough to clear markets, agents sell less than their total endowment, and income y_t is the variable that adjusts to clear markets.

Since this is a real model, some justification for the constraint $r_t \geq 0$ is in order. The constraint is a tractable way to model the effect of a zero lower bound on nominal interest rates, combined with a limit on the expected rate of inflation that the central bank can or will target. In Appendix A I show that this equilibrium is isomorphic to the limit of a standard New Keynesian model as prices become infinitely sticky, as in Korinek and Simsek [2014]. In Appendix B I present two alternative economies providing a microfoundation for this equilibrium concept. The first draws on the extensive literature on rationing or non-Walrasian equilibria (see e.g. Benassy [1993] for a survey, and Kocherlakota [2013], Caballero and Farhi [2014] for two recent papers employing a similar concept). The second is an economy with downward nominal wage rigidity drawing on Schmitt-Grohé and Uríbe [2011].

The following Proposition characterizes ZLB-constrained equilibria in the two-type economy.

Proposition 2.5. *In the two-type economy:*

1. *If $r(d_1) \geq 0$, the Walrasian equilibrium is also the ZLB-constrained equilibrium.*

2. *If $r(d_1) < 0$, then in a ZLB-constrained equilibrium:*

$$\begin{aligned} &r_1 = 0 \\ &u'(c_1^S) = \beta u'\left(y^* + (1-\beta)\phi\right) \\ &y_1 = c_1^S - d_1 + \phi < y^* \\ &c_1^B = c_1^S - 2d_1 + 2\phi \end{aligned}$$

The economy enters a steady state in period 2, as in the Walrasian equilibrium.

Proof. The first part of the proposition is obvious. To prove the second part, first note that if $r(d_1) < 0$, we cannot have a Walrasian equilibrium satisfying the ZLB. We must have $r_1 = 0$ and $y_1 < y^*$. Since borrowers are constrained at $t = 1$, $c_1^B = y_1 - d_1 + \phi$. Substituting this into the market clearing condition $c_1^S + c_1^B = 2y_1$, we get the above result. □

When debt is not too high, interest rates are positive, the ZLB does not bind, and markets clear. When debt is too high, borrowers' consumption falls sharply in period 1, and the ZLB prevents interest rates from falling enough to induce savers to consume the remaining output. The fall in income forces borrowers to reduce spending further. Figure 2 illustrates.

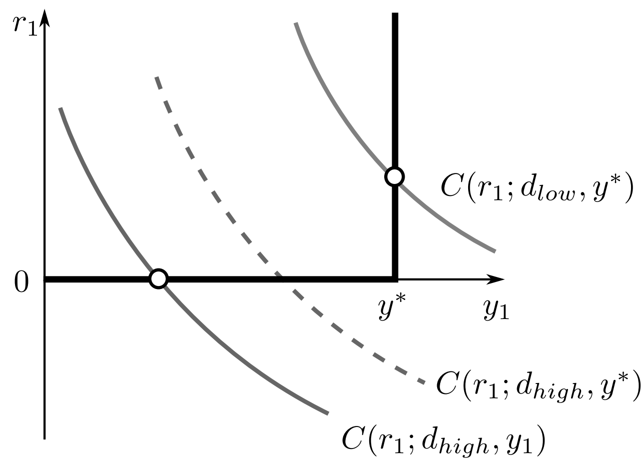

Figure 2: ZLB-constrained equilibrium

3 Liquidity traps and overborrowing

In this section, I first restate two results from the recent literature on liquidity traps and debt relief (in particular, Korinek and Simsek [2014]). First, overborrowing can occur in equilibrium: if the motive for borrowing is sufficiently strong, borrowers may take on so much debt that they trigger the ZLB, even though they know this will happen. Second, any equilibrium in which the ZLB binds is Pareto inefficient. Ex post, taxing savers and writing off borrowers' debt restores output to potential, makes borrowers better off, and leaves savers no worse off.

I then show that such a policy, while ex post optimal, may have adverse incentive effects ex ante. First, lump sum redistribution induces overborrowing, as borrowers take out more debt, anticipating that they will be richer in the future. Second, even if the government commits to a cap on debt relief, a transfer large enough to restore full employment may not be incentive compatible, since savers may take on debt in order to qualify for the transfer.[7]

In this section, and in Section 4, I focus on the two type economy to build intuition. In Section 5, I show how these results generalize to a continuum of types.

[7]In section 6, I show that private information introduces similar problems for macroprudential policy.

3.1 Overborrowing and the potential gains from transfers

A natural question, posed by Korinek and Simsek [2014], is whether borrowers may take on so much debt that the ZLB binds at date 1 even though they anticipate this happening. The next proposition shows that this is indeed possible, under the following assumption:

Assumption 3.1.

$$u'(2y^*) < \beta u'\left(y^* + (1-\beta)\phi\right)$$

The market clearing interest rate is decreasing in d_1. As borrowers become more impatient, they borrow more: in the limit as they become infinitely impatient, they promise to repay all their income at date 1, so savers must consume the whole aggregate endowment. Assumption (3.1) ensures that they would only do so if the date 1 interest rate was negative. This guarantees that if borrowers are sufficiently impatient, they will take on so much debt that $r(d) < 0$.

Proposition 3.2. *There exists $\theta^{ZLB} \in (1, \infty)$ such that $r(d_1) < 0$ if $\theta_B > \theta^{ZLB}$.*

If the ZLB binds at date 1, the recession makes households poorer, and ceteris paribus they want to borrow more at date 0. But if borrowers are sufficiently impatient, their impatience outweighs this wealth effect, and they take on so much debt that the ZLB binds.

An equilibrium in which the ZLB binds is ex post Pareto inefficient, since unanticipated redistribution from savers to borrowers can be Pareto improving. Suppose savers' income is unexpectedly reduced to $y_t - T$, and borrowers' income is unexpectedly increased to $y_t + T$. This is identical to an unanticipated reduction in the borrowers' debt. Redistribution directly reduces savers' income by T. However, since output is decreasing one for one in borrowers' debt, writing off T debt increases savers' income by T, leaving them no worse off. Borrowers, meanwhile, benefit twice from the redistribution: their debt falls by T, and their income rises by T due to the multiplier effect of their own spending. So we have a Pareto improvement. Given a large enough redistribution T, it is possible to restore full employment, as the following Proposition states:

Proposition 3.3. *If the borrowers receive an unanticipated increase in income $T^{FE}(d_1) = d_1 - (c_1^B + \phi - y^*)$ and the savers face an unanticipated fall in income $T^{FE}(d_1)$:*

1. *There is full employment: $y_1 = y^*$*

2. *Borrowers' consumption increases to $2y^* - c_1^S > c_1^S - 2d_1 + 2\phi$*

3. *Savers' consumption is unchanged.*

If the redistribution is equal to $T < T^{FE}(d_1)$, $y_1 = c_1^S + T + \phi - d$.

Proof. The unanticipated redistribution is equivalent to a change in d_1^B. The result follows by replacing d_1 with $d_1 - T$ in the equilibrium described in Proposition 2.5. □

3.2 Equilibrium with transfers

The government may attempt to implement a lump sum redistribution by writing off borrowers' debt and taxing savers. However, if redistribution is implemented through policy, it will be anticipated, and may distort decisions ex ante. To consider this possibility, it is necessary to define an equilibrium with policy.

I now replace the budget constraint (2) with the following:

$$
\begin{aligned}
c_1^i &= y_1^i - d_1^i + \frac{d_2^i}{1+r_1} + T(d_1^i, \theta_i) - \bar{T} \\
c_t^i &= y_t^i - d_t^i + \frac{d_{t+1}^i}{1+r_t} \text{ for } t \neq 1
\end{aligned}
\tag{5}
$$

where for any debt level d_1^i, the transfer to agent i in period 1 is $T(d_1^i, \theta_i) - \bar{T}$.[8] For now, I allow the government to observe households' type, and target transfers directly. The bulk of this paper will consider the case where θ is private information, and transfers can only depend on d_1^i.

The planner cannot make any taxes or transfers to agents starting in date 2, and must run a balanced budget:

$$\int T(d_1^i, \theta_i)\,\mathrm{d}i = \bar{T} \tag{6}$$

This assumption is crucial. If the government could impose taxes and transfers forever, it would be possible to completely undo the effect of the liquidity constraint, for example through a deficit-financed transfer to all households (Woodford [1990], Yared [2013], Bilbiie et al. [2013b]). I rule out such policies in order to isolate the effects of debt relief. Government credit policies may be a powerful tool in responding to recessions caused by a contraction in private credit. My goal is to evaluate whether non-credit policies, such as debt relief, can also be effective (not to seriously compare debt relief to other fiscal or monetary policies, a task which I leave to future work).[9]

I now define a ZLB-constrained equilibrium with debt-contingent date 1 transfers.[10]

Definition 3.4. *A **ZLB-constrained equilibrium with transfers** is $\{c_t^i, d_t^i, y_t, r_t, \bar{T}\}$ such that, given a transfer function $T(d, \theta)$:*

1. *each household i chooses $\{c_t^i, d_t^i\}$ to maximize (1) s.t. (5), (3), (4)*
2. $\int c_t^i\,\mathrm{d}i = y_t$
3. $r_t \geq 0$, $y_t \leq y^*$, $r_t(y^* - y_t) = 0$
4. *the government budget constraint (6) is satisfied.*

[8] It is always possible to normalize $\bar{T} = 0$. I write the transfer in this general form to ensure that balanced-budget equilibrum is defined for any transfer function $T(d, \theta)$.

[9] One motivation for employing non-credit policy in addition to deficit-financed transfers might be a concern with the distortionary effects of non-lump sum transfers in the long run, which is not modelled here. Non-credit polices such as debt relief are also feasible even when the government faces borrowing constraints, in addition to the private sector.

[10] Equilibria with debt-contingent date 0 transfers are defined analogously, and are considered in Section 6.

3.3 Debt relief induces overborrowing

Having defined ZLB-constrained equilibrium with transfers, I describe how the prospect of debt relief distorts decisions ex ante. In this stylized model, there are two ways in which this can happen: the intensive and the extensive margin. Along the intensive margin, the anticipation of debt relief causes borrowers to borrow more, because they will be richer in period 1, and want to borrow against that wealth at date 0. I show that the ex post optimal policy is never Pareto improving ex ante: some commitment is necessary if we are to obtain a Pareto improvement.

Suppose first that the government does not commit ex ante to a particular level of transfers. Instead, after observing the equilibrium level of debt d_1^*, it makes whatever transfer $T^{FE}(d_1^*)$ to borrowers restores full employment, and finances this with a lump sum tax on savers.[11] Note that since an individual borrower is measure zero, the transfer she receives does not depend on her own debt, but only on aggregate debt d_1^*, which she is too small to affect.

Proposition 3.5. *Consider a constant transfer function $T(d, \theta_B) = T^*$, $T(d, \theta_S) = -T^*, \forall d$. Suppose $T^* = T^{FE}(d^*)$, where d^* is the level of d_1^B in the equilibrium with transfers, given T^*. This equilibrium is not Pareto improving relative to the equilibrium without transfers. Savers are strictly worse off.*

Proof. Combining the savers' and borrowers' Euler equations and using market clearing

$$\frac{U_c(c_0^S, \theta_S)}{U_c(2y^* - c_0^S, \theta_B)} = \frac{u'(c_1^S)}{u'(2y^* - c_1^S)} > \frac{u'(c_1^S)}{u'(\hat{c}_1^B)}$$

where $\hat{c}_1^B < 2y^* - c_1^S$ denotes the equilibrium without policy. It follows that $c_0^S < \hat{c}_0^S$, and since $c_0^S = \hat{c}_0^S$, savers are worse off. □

If the government makes a transfer to borrowers, they will be richer at date 1. Anticipating this, they borrow against this future income to consume more at date 0, and in equilibrium, savers consume less. In this sense, debt relief encourages overborrowing.

Note that this result holds even if borrowers do not perceive that their individual debt will be written off: if they did, there would be even more overborrowing ex ante. Suppose borrowers expect the government to write off a fraction τ of their debt: that is, $T_B(d) = \tau d$. Then they face an effective gross interest rate of $(1 + r_0)(1 - \tau)$, while savers face an interest rate $1 + r_0$. This wedge between the borrowers' and savers' Euler equations makes the borrowers' consumption even higher at date 0 and makes savers even worse off.

[11] Throughout this paper, I consider debt relief policies in which the government assumes responsibility for borrowers' debt, and makes payments to savers on the borrowers' behalf, financing these payments with lump sum taxes on the savers. Crucially, since an individual saver is measure zero, her lending decision does not affect the lump sum tax required to pay for the debt relief. One could consider an alternative debt relief policy, in which the government decrees that borrowers no longer have to make some promised payments to savers (effectively, legislating a mass default). Under this alternative policy, debt relief would create default risk, which will be priced into the interest rates charged by savers at date 0. In contrast, under the policy considered in this paper, interest rates are not directly affected by debt relief, since there is no default.

3.4 Is debt relief incentive compatible?

Above, I showed that debt relief without commitment induces overborrowing, and is not ex ante Pareto improving. In this section, I show that even if the government commits to limit the amount of debt relief, the full employment transfer may not be incentive compatible. If debt relief is too generous, savers switch to become borrowers, and the equilibrium breaks down. This problem arises if borrowers are sufficiently impatient, so they take on so much debt that the required amount of debt relief would tempt the savers to borrow.

One way to limit debt relief is as follows. Take the level of debt in the equilbrium without policy, $\hat{d}$. Let the government give a transfer $T^{FE}(\hat{d})$ to borrowers with exactly $\hat{d}$ debt.[12] Since borrowers only receive a transfer if they borrow exactly $\hat{d}$, this policy obviously cannot induce overborrowing. Clearly, this transfer function is unrealistic. My goal is to show that even if I allow the government to completely avoid overborrowing in this way, another problem remains.

Assume that household type, θ_i, is private information: the government cannot directly distinguish type S and type B agents. In this case, transfers must be anonymous: $T(d,\theta_S) = T(d,\theta_B) = T(d)$. Formally:

Definition 3.6. *A* ***balanced budget equilibrium with anonymous transfers*** *(henceforth, an* ***equilibrium with transfers****) is a balanced budget equilibrium with transfers in which $T(d,\theta) = T(d), \forall\theta$.*

With anonymous transfers, the only way to transfer funds to borrowers is to reward agents who take on more debt, and tax savers, i.e. to make $T(d)$ positive for some $d > 0$, and negative for some $d < 0$. If the transfer $T(\hat{d})$ is sufficiently large, savers will receive strictly higher utility by mimicking borrowers, taking on debt $\hat{d}$ instead of saving $\hat{d}$. Then the ex post optimal debt relief policy is not incentive compatible, and cannot be implemented. If the government offered such generous debt relief, at least some savers would mimic borrowers, and borrow $\hat{d}$. Then the government will be forced to raise taxes at date 1, and the net transfer to borrowers will not be enough to restore full employment

Figure 3 illustrates a case in which the full employment transfer is not incentive compatible. Date 1 output is below potential, $y_1 < y^*$. Savers' date 1 consumption is constrained by the ZLB. Through taxes and transfers, the government could essentially transfer all the surplus output, $y^* - y$, to the borrowers, increasing their date 1 consumption and leaving everything else unchanged. However, then borrowers' allocation would give strictly higher utility to savers than their own allocation, and savers would rather mimic borrowers than choose their own allocation. At least some savers will take on the same debt as borrowers in order to qualify for the transfer at date 1, and the government will be forced to reduce the net transfer to borrowers.

When will the FE transfer violate incentive compatibility? That is, when will savers prefer borrowers' allocation (including the date 1 transfer) to their own allocation? Borrowers consume less than savers at date 1, despite the transfer, but consume more at date 0. If borrowers are sufficiently impatient (θ_B is large enough), they consume so much more at date 0 that savers

[12] Formally, this corresponds to the transfer functions $T_B(d) = T^{FE}(\hat{d})\mathbf{1}(d = \hat{d})$, $T_S(d) = -T^{FE}$.

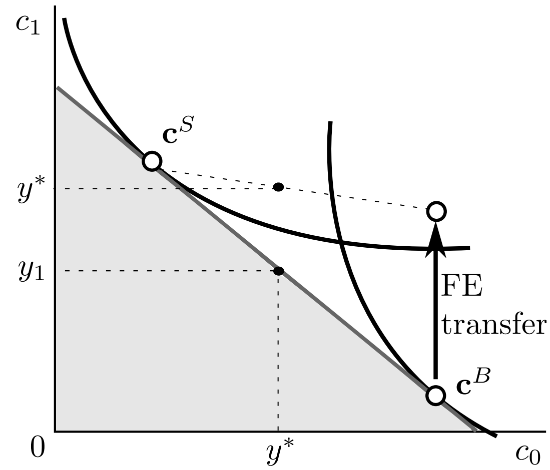

Figure 3: Full employment transfer is not incentive compatible

would want to mimic them and do the same, if the government writes off their debt at date 1. As $\beta \to 1$, steady state interest rates go to zero, so savers' steady state consumption converges to their income y^*; at the same time, as savers become more patient, they become unwilling to consume more at date 1 than they do in the steady state, so their date 1 consumption also converges to y^*. The cost of mimicking borrowers - lower consumption at date 1 and in steady state - converges to zero, so savers become increasingly willing to mimic borrowers.

Proposition 3.7. *Consider the transfer function $T(d) = T^{FE}(\hat{d})$ if $d = \hat{d}$, $T(d) = -T^{FE}(\hat{d})$ if $d \neq \hat{d}$, where $\hat{d}$ is the debt level in the equilibrium without policy. There exists a continuous function $\theta^{FE}(\beta, \phi)$, which may equal ∞, such that:*

1. *If $\theta_B \leq \theta^{FE}$, the transfer is incentive compatible. There exists an equilibrium with transfers with full employment which is a Pareto improvement over the equilibrium without policy.*

2. *If $\theta_B > \theta^{FE}$, the transfer is not incentive compatible.*

3. *$\theta^{FE}(\beta, \phi) \geq \theta^{ZLB}$. If the FE transfer is not incentive compatible, the ZLB must bind in equilibrium.*

4. *$\theta^{FE}(\beta, \phi)$ is increasing in β and decreasing in ϕ, with $\theta^{FE}(1, 0) = 1$. That is, if $\beta = 1, \phi = 0$, the transfer is not incentive compatible for any $\theta_B > 1$. For any $\theta_B \in [1, \infty)$, there exist $\bar{\beta}, \bar{\phi}$ sufficiently close to $1, 0$ such that the FE transfer is not incentive compatible if $\beta > \bar{\beta}, \phi < \bar{\phi}$.*

4 Optimal policy in the two type economy

In the previous section, I showed that poorly designed debt relief policies, while optimal ex post, can have harmful incentive effects ex ante. The question remains: what is optimal policy in this economy? Can sophisticated debt relief programs avoid these adverse incentives?

To answer these questions, I now analyze optimal policy. First I solve the Pareto problem, subject to the constraints imposed by incentive compatibility and the zero lower bound. I let the government choose *any* system of taxes and transfers which depend only on an agent's observable debt, not on her unobservable type, and show that solutions to the Pareto problem can be implemented with such debt-contingent transfers. I show that debt relief always implements some constrained efficient allocations (in particular, those which are relatively favorable for borrowers). When the ZLB binds, some debt relief policy is always Pareto improving. When the ZLB does not bind, debt relief is purely redistributive, taking from savers and giving to borrowers.

4.1 Pareto problem

To characterize constrained efficient allocations, I consider the problem of a social planner who puts weight α on savers and $1-\alpha$ on borrowers, and faces four sets of constraints. First, resource feasibility. Second, the liquidity constraint at date 1, which (combined with the assumption of no transfers at date 2) puts a lower bound on the date 2 consumption of borrowers. Third, the zero lower bound, which imposes that the savers' Euler equations must be satisfied with a nonnegative interest rate. And fourth, incentive compatibility, which states that neither agent can strictly prefer the other agent's allocation. As discussed above, since my focus is on non-credit policies, I assume the planner cannot make any taxes or transfers to agents starting in date 2, and must run a balanced budget. Consequently, the economy always enters a steady state at date 2.

$$\max \alpha \left\{ U(c_0^S, \theta_S) + \beta u(c_1^S) + \frac{\beta^2}{1-\beta} u(c_2^S) \right\} + (1-\alpha) \left\{ U(c_0^B, \theta_B) + \beta u(c_1^B) + \frac{\beta^2}{1-\beta} u(c_2^B) \right\} \tag{7}$$

$$\text{s.t. } c_0^S + c_0^B \leq 2y^* \tag{RC0}$$

$$c_1^S + c_1^B \leq 2y^* \tag{RC1}$$

$$c_2^S + c_2^B = 2y^* \tag{RC2}$$

$$c_2^B \geq y^* - (1-\beta)\phi \tag{BC}$$

$$u'(c_1^S) \geq \beta u'(c_2^S) \tag{ZLB}$$

$$U(c_0^S, \theta_S) + \beta u(c_1^S) + \frac{\beta^2}{1-\beta} u(c_2^S) \geq U(c_0^B, \theta_S) + \beta u(c_1^B) + \frac{\beta^2}{1-\beta} u(c_2^B) \tag{ICS}$$

$$U(c_0^B, \theta_B) + \beta u(c_1^B) + \frac{\beta^2}{1-\beta} u(c_2^B) \geq U(c_0^S, \theta_B) + \beta u(c_1^S) + \frac{\beta^2}{1-\beta} u(c_2^S) \tag{ICB}$$

Constrained efficient allocations solve (7) for some $\alpha \in (0,1)$. Varying α traces out the constrained Pareto frontier. The following proposition characterizes constrained efficient allocations.

Proposition 4.1. *There are ten classes of constrained efficient allocation.*

1. (RC0) *and* (RC2) *always bind.*

2. *Either* (ICS) *binds,* (ICB) *binds, or no incentive constraints bind. There exist* α_B, α_S *with* $1 > \alpha_B > \alpha_s > 0$ *such that* (ICB) *binds iff* $\alpha > \alpha_B$ *and* (ICS) *binds iff* $\alpha < \alpha_S$.

3. *Either neither* (BC) *nor* (ZLB) *bind, only* (BC) *binds, or* (BC) *and* (ZLB) *both bind.*

4. (RC1) *binds unless* (ICS), (BC) *and* (ZLB) *all bind. In this case,* (RC1) *may be slack.*

Figure 4 illustrates point 2.[13] It shows which incentive constraints bind, as a function of α and θ_B. The dark grey region on the left shows the set of (α, θ_B) for which (ICS) binds; the unshaded middle region shows the set where neither constraint binds; and the light grey region on the right shows the set where (ICB) binds. Setting $\alpha = 0$ selects an allocation which maximizes B's utility, subject to the remaining constraints, putting no weight on S. Absent incentive compatibility constraints, this allocation would have B consuming everything and S nothing. Incentive compatibility rules this out, since S would want to mimic B. So when $\alpha = 0$, (ICS) always binds. Increasing α, we go from left to right, moving along the Pareto frontier towards allocations that are better for S and worse for B. Eventually, S's utility increases so much that (ICS) no longer binds, and each individual strictly prefers his own allocation. [14] Increasing α further, eventually B's utility falls so much that he prefers S's allocation, and (ICB) binds.

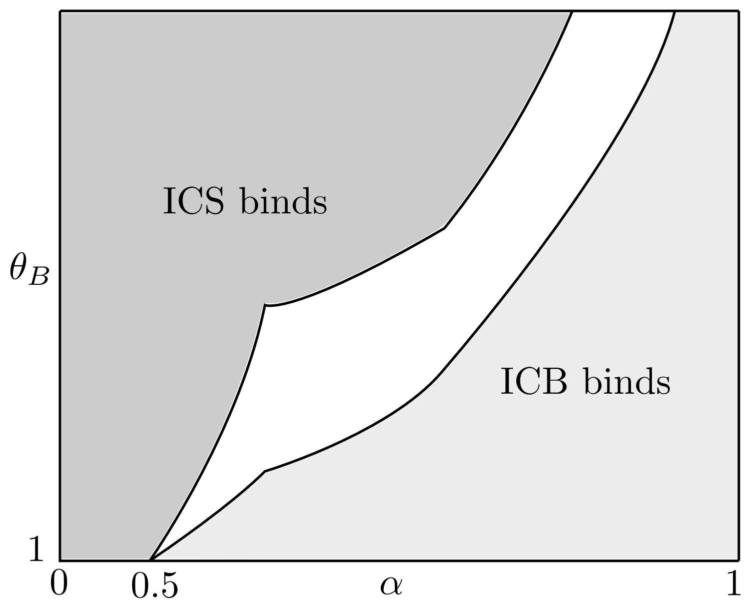

Figure 4: Optimal allocations in (α, θ_B)-space

Figure 5 illustrates point 3. It shows whether the borrowing constraint is slack, the borrowing constraint binds, or both the borrowing constraint and the ZLB binds, as a function of α and θ_B. When θ_B is close to $\theta_S = 1$, borrowers are almost as patient as savers, and have similar consumption profiles. The borrowing constraint does not bind at date 1, and the economy enters steady state immediately. As we raise θ_B, borrowers become more impatient, consuming more at date 0 and less at dates 1 and 2. Eventually, the borrowing constraint binds: the planner would like to give reduce borrowers' steady state consumption, but would increase steady state debt above ϕ. Increasing θ_B further, borrowers become yet more impatient, and the planner gives less

[13] This figure is not to scale. In particular, note that the two shaded regions meet at the point $\alpha = 0.5, \theta = 1$, which should lie in the middle of the horizontal axis.

[14] This region only exists when agents have different preferences, i.e. $\theta_B > \theta_S = 1$.

to borrowers and more to savers at date 1. Date 2 allocations, however, remain fixed. Savers must tolerate an increasingly steep decline in consumption between dates 1 and 2; interest rates fall to induce them to do so. Eventually, (ZLB) binds, and we enter the light grey region at the top of Figure 5. Savers consume $\bar{c}_1^S$, the maximum they can be induced to consume with zero interest rates. Unlike in the competitive equilibrium, this does not generally mean that aggregate consumption falls below output: the planner recognizes the ZLB constraint, gives the remaining consumption to the borrowers, and distributes date 0 consumption between B and S.

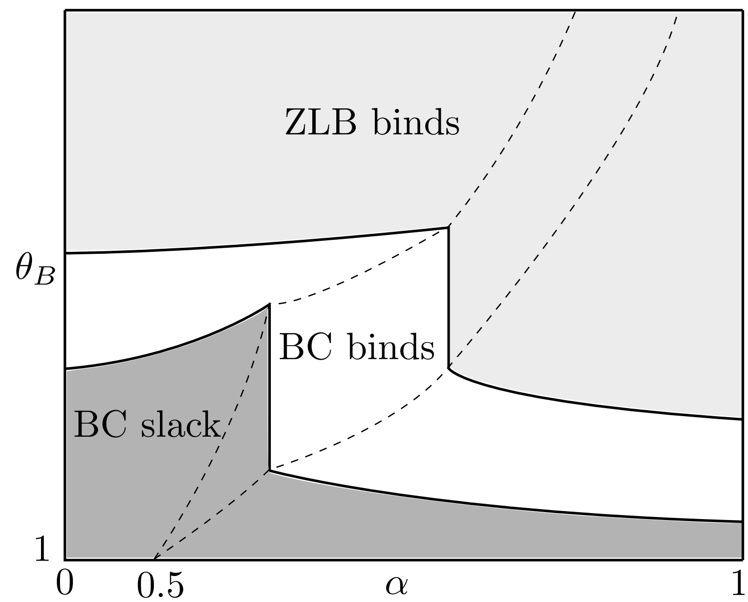

Figure 5: Optimal allocations in (α, θ_B)-space

Finally, Figure 6 illustrates point 3. It shows the region of parameter space (shaded black) in which the date 1 resource constraint is slack, so aggregate consumption is less than potential output. The figure shows that this may be constrained efficient, but only if (ZLB) and (ICS) both bind. This will be the case if θ_B is high enough, in allocations which are relatively favorable for borrowers (corresponding to a low α). Suppose borrowers are very impatient, so θ_B is very high, and suppose (ICS) and (ZLB) bind. Borrowers would like more date 0 consumption, but that would tempt savers to choose the borrowers' allocation, violating incentive compatibility. To get more date 0 consumption, borrowers can throw away some of their date 1 consumption. (They cannot give it to savers because (ZLB) binds.) This makes their allocation less attractive to savers, who put a high weight on date 1 consumption, which in turn means that borrowers can get away with higher date 0 consumption, which they value more.So if θ_B is sufficiently high, there are some constrained efficient allocations in which (ZLB) and (ICS) bind but (RC1) is slack.

While I will discuss implementation below, note for now that the difference between potential output and consumption, $2y^* - c_1^S - c_1^B$, can be interpreted as unproductive government spending. Point 5 then states that unproductive spending may be optimal, provided not only that the economy is in a demand-driven slump (as in the standard Keynesian argument), but also

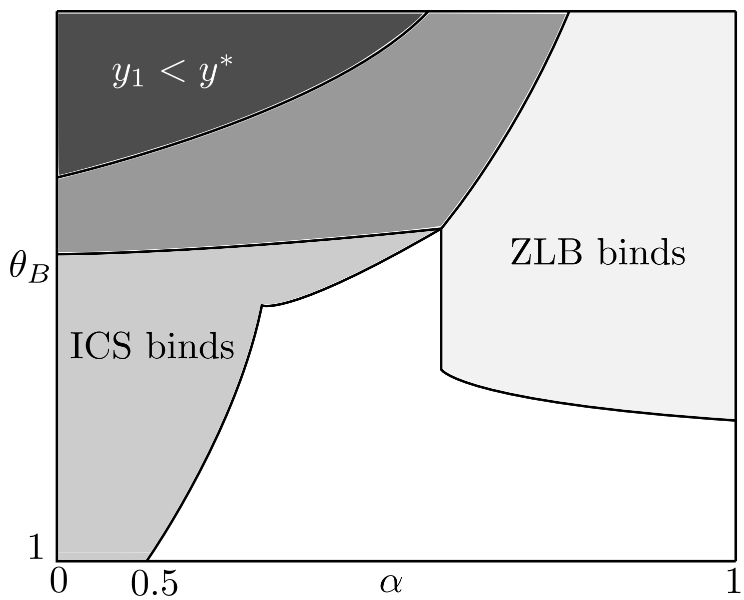

Figure 6: Optimal allocations in (α, θ_B)-space

that incentive constraints prevent the government from achieving full employment with targeted transfers alone. More generally, if I introduced government spending explicitly and allowed it to have some value for households, it would be optimal to increase government spending above the normal efficient level when incentive constraints bind. Intuitively, one advantage of spending on pure public goods is that they benefit all agents equally, and do not create incentive problems.

To summarize, the optimal allocation has three important properties. First, when moving along the Pareto frontier, in allocations which are relatively favorable for borrowers, (ICS) binds; in intermediate allocations, neither incentive constraint binds; and in allocations which are better for savers, (ICB) binds. Second, when differences in preferences between the two households are large, creating a motive for borrowing, it is optimal for (ZLB) to bind. Unlike in the equilibrium without policy, however, even when (ZLB) binds, the planner generally uses targeted transfers to prevent unemployment. Finally, it may be constrained optimal to allow some unemployment when (ZLB) binds, but only if (ICS) also binds, so the planner cannot give the surplus output to B without making S prefer B's allocation.

4.2 Implementation

Next, I consider how constrained efficient allocations can be implemented. First I show that any solution to the Pareto problem can be implemented as an equilibrium with debt-contingent transfers. Then I ask when it is optimal for these transfers to take the form of debt relief.

The next proposition states that any constrained efficient allocation can be implemented as an equilibrium with debt-contingent transfers at date 1.[15]

[15] In Section (6), I show that efficient allocations can also be implemented with date 0 transfers.

Proposition 4.2. *Any solution to (7) can be implemented as an equilibrium with transfers.*

Intuitively, every transfer function $T(d)$ maps out a nonlinear budget constraint in consumption space. In any constrained efficient allocation, each agent prefers her own consumption allocation to the other agent's allocation. The implementability problem is to construct a nonlinear budget set such that each agent prefers her own allocation to all other allocations in the budget set.[16] Figure 7 provides an illustration. c^B, c^S are arbitrary consumption allocations in (c_0, c_1)-space satisfying incentive compatibility, so S's allocation lies weakly below B's indifference curve, and vice versa. The gray line shows one particular nonlinear budget constraint which implements this allocation. Graphically, it is clear that for any allocation, we can find some nonlinear budget set which lies below the lower envelope of both agents' indifference curves, and intersects the indifference curves at their intended allocations.

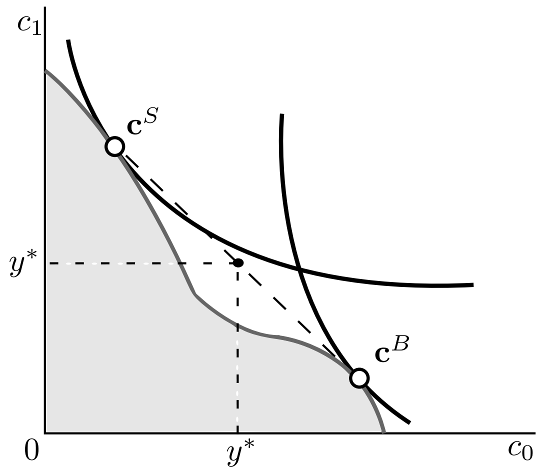

Figure 7: Implementation with debt-contingent transfers

4.3 Debt relief implements constrained efficient allocations

Next, I ask whether constrained efficient allocation can be implemented with particular simple transfer functions.

It turns out that any constrained efficient allocation can be implemented with a piecewise linear transfer function with at most three segments. The transfer function is one of two kinds. I call the first a debt relief transfer function:

Definition 4.3. *$T(d)$ is a **debt relief transfer function** if it has the form*

$$\begin{aligned} T^{DR}(d; \bar{T}, \underline{d}, \bar{d}) &= -\bar{T} \text{ if } d < \underline{d} \\ &= -\bar{T} + (d - \underline{d}) \text{ if } d \in [\underline{d}, \bar{d}] \\ &= -\bar{T} + (\bar{d} - \underline{d}) - \tau(d - \bar{d}) \text{ if } d > \bar{d} \end{aligned}$$

[16]There are multiple ways to do this. A trivial solution is to offer only two points in the budget set, corresponding to the desired debt levels of S and B, and set $T(d) = -\infty$ for all off-equilibrium levels of debt.

where $\bar{T} > 0$, $\underline{d}$, $\bar{d} > \underline{d}$, τ are parameters.

There is a lump sum tax on all households with debt below a certain level $\underline{d}$. The government writes off all debt above $\underline{d}$ up to a maximum level $\bar{d}$. Above that point, further borrowing is penalized: the transfer falls by τ dollars for each dollar of debt above the maximum level $\bar{d}$.[17] Debt relief is offered only to borrowers with a moderate level of debt: excessive borrowing is discouraged. Figure 8 shows the budget constraint induced by a debt relief transfer function.

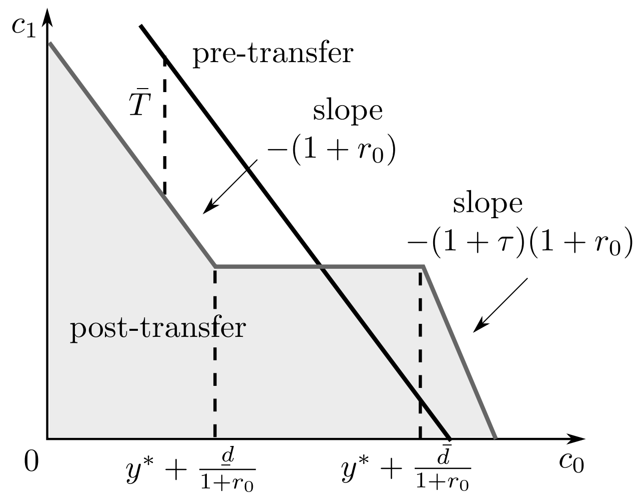

Figure 8: A debt relief transfer function

The next proposition states conditions under which debt relief is Pareto optimal.

Proposition 4.4. *There exists $\alpha(\theta_B)$ such that*

1. *debt relief implements the optimal allocation iff $\alpha < \alpha(\theta_B)$.*
2. *If (ICS) binds, $\alpha < \alpha(\theta_B)$ and debt relief implements the optimal allocation.*

First, consider a constrained efficient allocation in which S's incentive compatibility constraint binds. Figure 9 shows that debt relief implements such an allocation. In equilibrium, borrowers all choose exactly the maximum level of debt. While there are other transfer schemes that might implement allocations in which (ICS) binds, in any such scheme, $T(d)$ must be nondifferentiable at d_1^B, B's equilibrium debt level. In this case, S and B's indifference curve intersect at $\mathbf{c}^B$. Because the two households have different preferences, the indifference curves have different slopes at this point. It follows that the budget set must have a kink at $\mathbf{c}^B$.

[17] One could imagine a situation in which borrowers can verify that they have a debt, but can hide part of their debt - if their true debt level is $d > 0$, they can claim to have any debt level $\tilde{d} \in [0, d]$. In this case, it would not be possible to offer a transfer function where $T(d)$ is decreasing over some range. Instead, it would be necessary to combine debt relief with a 'macroprudential' debt limit $d \leq \bar{d}$ at date 0.

Crucially, I assume that all asset trades are observable. If agents could engage in secret asset trades, savers could reduce their tax burden. For example, consider two savers who save $a > 0$ in equilibrium, and both pay the lump sum tax $\bar{T}$. Consider the following deviation: one saver saves $3a$ and pays $\bar{T}$, the other saver borrows a and receives a transfer equal to $\bar{T}$, and they pool their resources at each date. This deviation reduces their tax burden to zero. I rule out such deviations by prohibiting secret asset trades.

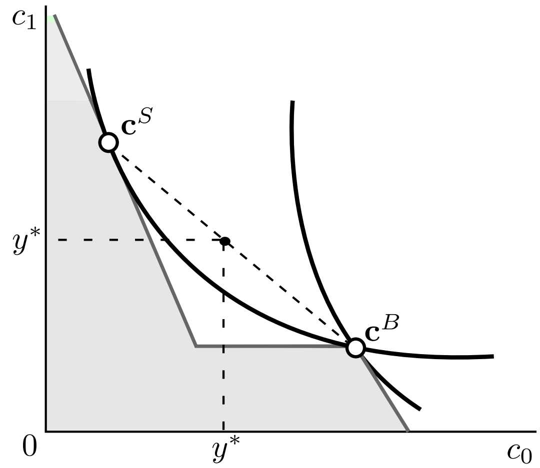

Figure 9: Debt relief implements allocations in which (ICS) binds

Debt relief can also be used to implement allocations in which (ICS) does not bind, but which are still relatively favorable to borrowers. Figure 10 provides an example. In this case, it is not strictly necessary for the transfer function to be nondifferentiable, since both agents will locate at interior points. If the ZLB does not bind, borrowers and savers face the same interest rates in the first period, and we can set the tax on excessive debt, τ, equal to zero. If the ZLB binds, it will in general be necessary to tax either borrowers or savers.

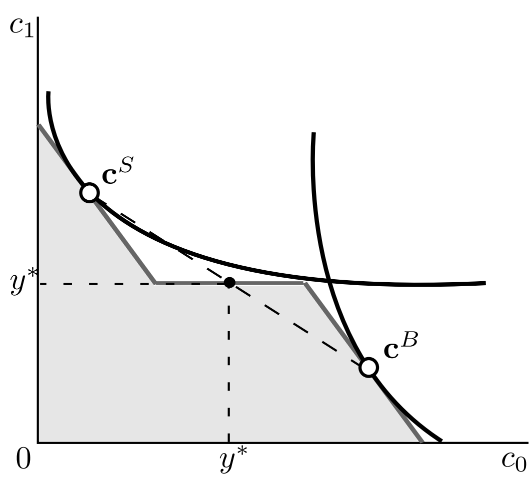

Figure 10: Debt relief implements some allocations in which no ICs bind

4.4 Savings subsidies

Intuitively, it seems unlikely that debt relief implements every constrained efficient allocation. Consider the extreme case in which $\alpha = 1$, so we look for the constrained efficient allocation which is best for the saver, ignoring the borrower's welfare altogether. Clearly, savers have little to gain by offering debt relief to the borrowers. If anything, they would prefer to tax borrowers, and transfer resources to themselves - subject to the *borrowers'* incentive compatibility constraints. In this section, I show that efficient allocations which are relatively favorable for savers can be implemented with a subsidy to savers, as this intuition suggests.

Define savings subsidies as follows:

Definition 4.5. *$T(d)$ is a **savings subsidy** if it has the form*

$$\begin{aligned} T^{SS}(d; T_s, T_b, d^*, \tau) &= T_b \text{ if } d > d^* \\ &= T_s - \tau(d^* - d) \text{ if } d \leq d^* \end{aligned}$$

where T_s, T_b, d^*, τ are parameters.

There is a lump sum tax on households with debt above a certain level d^*. There is a lump sum subsidy to households with debt equal to d^*. Excessive saving is penalized by a tax τ on saving above this level.

Proposition 4.6. *1. A savings subsidy implements the optimal allocation iff $\alpha > \alpha(\theta_B)$.*

2. If (ICB) binds, $\alpha > \alpha(\theta_B)$ and a savings subsidy implements the optimal allocation.

Consider a constrained efficient allocation in which B's incentive compatibility constraint binds. Figure 11 shows that savings subsidies implement such an allocation. In equilibrium, savers all choose exactly the maximum level of savings. In any transfer scheme implementing an optimal allocation of this kind, $T(d)$ must be nondifferentiable at d_1^S, S's equilibrium debt level. As in the case when (ICS) binds, S and B's indifference curve intersect, this time at c^S. Again, because the two households have different preferences, the indifference curves have different slopes at this point, and the budget set must be kinked at c^S.

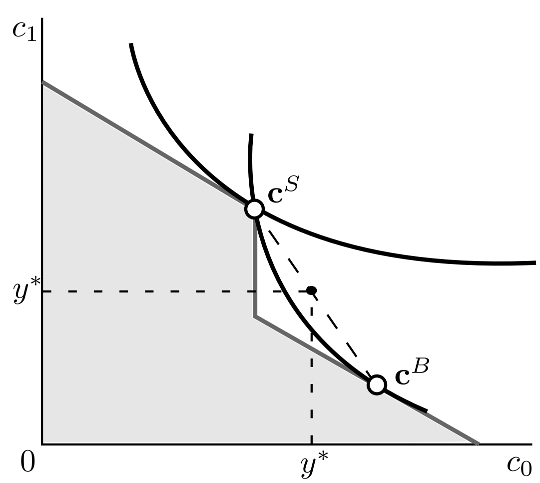

Figure 11: Savings subsidies implement allocations in which (ICB) binds

4.5 Pareto improving debt relief

While debt relief always implements some Pareto optimal allocations, we have just seen that savings subsidies - the opposite of debt relief - always implement efficient allocations. In what sense is debt relief a desirable policy?

In this section, I ask whether debt relief is Pareto improving relative to the competitive equilibrium without policy. The following proposition states that this is always the case, provided the zero lower bound binds in equilibrium. Even if $\alpha = 1$, so the planner only values S's welfare, there exists a debt relief policy which is welfare improving at the zero lower bound. However, if the ZLB does not bind, debt relief is purely redistributive: it increases utility for borrowers but reduces utility for savers.

Proposition 4.7. *1. If $\theta_B \leq \theta^{ZLB}$, the competitive equilibrium is Pareto optimal (it is the solution to the planner's problem with $\alpha = \alpha(\theta_B)$. Debt relief is not Pareto improving.*

2. If $\theta_B > \theta^{ZLB}$, the competitive equilibrium is Pareto inefficient. Debt relief is always Pareto improving.

Intuitively, suppose the zero lower bound binds in equilibrium, and date 1 consumption is below potential output. In any competitive equilibrium, each agent strictly prefers their own allocation to the other agent's allocation, and both incentive constraints are slack. Suppose we attempt to increase date 1 consumption. The savers cannot consume more, since the zero lower bound binds. However, we can increase borrowers' consumption by some amount before we make savers' incentive compatibility constraint bind. This is a Pareto improvement.

If the zero lower bound does not bind in equilibrium, the competitive equilibrium is Pareto optimal, for the usual reasons. We already know that debt relief remains constrained efficient in this case: relative to the competitive equilibrium, it provides higher utility to borrowers and lower utility to savers. But it does not offer a Pareto improvement over the competitive equilibrium.

4.6 Characterizing optimal debt relief

In Section 4.5, I showed that there always exists some Pareto improving debt relief policy at the zero lower bound. In this section, I focus on one particular Pareto improving policy, namely the one which is most favorable to borrowers. I explain what determines the amount of debt relief which is optimal. In Section 3 I showed that poorly designed debt relief causes overborrowing; in this section, I explain how optimal policy avoids overborrowing.

As we have seen above, there are a continuum of constrained efficient allocations, indexed by α, the Pareto weight on savers. I focus one one particular allocation, namely, the allocation which maximizes borrowers' utility, subject to the constraint that savers are no worse off than in the equilibrium without policy. That is, the allocation solves

$$\max_{c_0^S, c_1^S, c_2^S, c_0^B, c_1^B, c_2^B} U(c_0^B, \theta_B) + \beta u(c_1^B) + \frac{\beta^2}{1-\beta} u(c_2^B) \tag{8}$$

$$\text{s.t. } U(c_0^S, \theta_S) + \beta u(c_1^S) + \frac{\beta^2}{1-\beta} u(c_2^S) \geq \bar{U}(\theta_S, \theta_B, \phi) \tag{US}$$

$$\text{(RC0), (RC1), (RC2), (BC), (ZLB), (ICS)}$$

where $\bar{U}(\theta_S, \theta_B, \phi)$ is the savers' utility in the equilibrium without policy.[18] I call the solution to this program the **borrower-optimal** allocation. To be clear, this always solves our original Pareto problem (7) for some α. All the results in Sections 4 and 4.2 therefore apply.

To guarantee that borrower-optimal allocations are continuous in θ_B, we make the following assumption.

Assumption 4.8. *$U_{cc\theta} < 0$.*

This is satisfied by $U(c, \theta) = \theta u(c)$ with u concave, and by $U(c, \theta) = u(c - \theta)$ if $u''' > 0$.

The following proposition characterizes borrower-optimal allocations.

Proposition 4.9. *The solution to the borrower-optimal problem is in one of five classes.*

1. *If the ZLB does not bind in equilibrium, the optimal allocation is the equilibrium without policy.*

2. *If the ZLB binds in equilibrium, and the full employment transfer is incentive compatible, i.e.*

$$U(c_0^S, \theta_S) + \beta u(\bar{c}_1^S) + \frac{\beta^2}{1-\beta} u(\bar{c}_2^S) \geq U(c_0^B, \theta_S) + \beta u(2y^* - \bar{c}_1^S) + \frac{\beta^2}{1-\beta} u(\bar{c}_2^B)$$

then the full employment transfer is optimal. That is, the optimal allocation is identical to the equilibrium without policy, except that c_1^B is increased to $2y^* - c_1^B$. (US) and (RC1) bind.

If the full employment transfer is not incentive compatible, there are three possibilities. Let $\hat{c}(\theta_B)$ solve

$$U_c(\hat{c}, \theta_B) = U_c(\hat{c}, \theta_S) + U_c(2y^* - \hat{c}, \theta_S) \tag{9}$$

And let $\underline{c}(\phi)$ be the smallest level of c_0^B which satisfies (ICS) and (RC1) with equality:

$$U(2y^* - \underline{c}, \theta_S) + \beta u(\bar{c}_1^S) + \frac{\beta^2}{1-\beta} u(\bar{c}_2^S) = U(\underline{c}, \theta_S) + \beta u(2y^* - \bar{c}_1^S) + \frac{\beta^2}{1-\beta} u(\bar{c}_2^B)$$

Then:

3. *If $\hat{c}(\theta_B) < \underline{c}(\phi) < c_0^B$, the optimal allocation is $\underline{c}(\phi)$, and (ICS) and (RC1) bind.*

4. *If $\underline{c}(\phi) < \hat{c}(\theta_B) < c_0^B$, the optimal allocation is $\hat{c}(\theta_B)$, and (ICS) binds.*

[18] Note that we can omit (ICB), since we are trying to maximize the borrowers' utility.

5. *If* $\underline{c}(\phi) < c_0^B < \hat{c}(\theta_B)$*, the optimal allocation is* c_0^B*, and* (ICS) *and* (US) *bind.*

When the ZLB is slack, this allocation is identical to the equilibrium, since the equilibrium is already efficient, and there is no policy. When the ZLB binds, Proposition 4.7 implies that the borrower-optimal allocation involves debt relief.

Even when the full employment transfer is not incentive compatible, it is still possible to increase borrowers' date 1 consumption until (ICS) binds. In fact, it is sometimes possible to do even better, as Figure 12 illustrates. When (ICS) binds, we can increase c_1^B further by reducing c_0^B and increasing c_0^S, to keep savers indifferent between their own allocation and the borrowers' allocation. This can be implemented by reducing the borrowers' debt, i.e. setting the cap on debt relief, $\bar{d}$, below the level of debt in the equilibrium without policy.

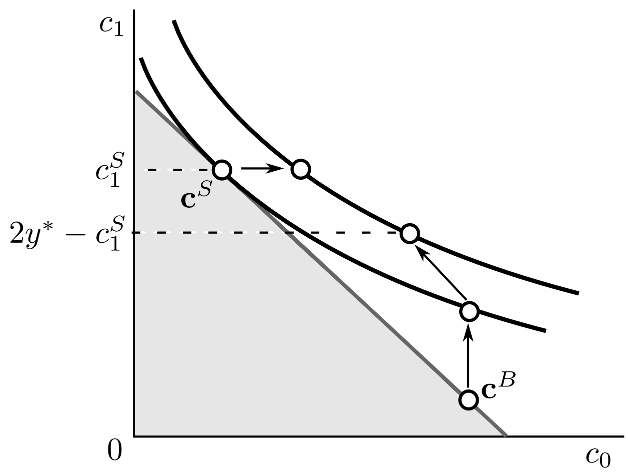

Figure 12: Borrower-optimal policy reduces c_0^B

When is it optimal to reduce debt in this way? When borrowers are not too impatient (θ_B is not too high), the gains from higher date 1 consumption outweigh the cost of lower date 0 consumption. When θ_B is slightly higher, it is optimal to reduce debt to some extent, but not all the way to full employment. Finally, if borrowers are sufficiently impatient, it is never optimal to reduce debt in return for higher date 1 consumption.

In section 3, I showed that debt relief can induce overborrowing, on both the intensive and extensive margins. On the intensive margin, debt relief increases borrowers' lifetime wealth. Absent any change in interest rates, this increase in wealth would induce borrowers to consume more at date 0, which would mean that savers consume less in equilibrium, making savers worse off. On the extensive margin, debt relief might induce some households who would otherwise save, to borrow instead, in order to qualify for this transfer.

The optimal debt relief policy avoids both these pitfalls. To prevent overborrowing on the intensive margin, the optimal policy requires a wedge τ between the shadow interest rates faced by borrowers and savers, as the following proposition states.

Proposition 4.10. *Define the wedge*

$$\tau := \frac{U_c(c_0^B, \theta_B)}{U_c(c_0^S, \theta_S)} \frac{u'(c_1^S)}{u'(c_1^B)} - 1.$$

If c^S, c^B is Pareto improving relative to the equilibrium without policy, $\tau > 0$.

A positive wedge $\tau > 0$ increases the effective marginal interest rate faced by borrowers at date 0, making them borrow less. As I showed in Section 4.2, this wedge can be implemented by making debt relief conditional on having a relatively moderate level of debt, below some maximum $\bar{d}$. Above $\bar{d}$, an additional dollar of debt reduces the transfer that households receive by τ dollars.[19] Making debt relief conditional prevents borrowers from borrowing more than they would have done in the equilibrium without policy.

Optimal policy avoids overborrowing on the extensive margin in two ways: by keeping T at a moderate level, and by lowering debt. Conditional debt relief can reduce equilibrium debt by setting $\bar{d}$ lower than the equilibrium level of debt, and charging a high tax on debt above $\bar{d}$. However, this is not optimal for borrowers if they are too impatient. In this case, the only way to prevent overborrowing is to keep T below the level required for full employment.

5 Optimal policy with a continuous distribution of types

A concern with the model presented above is that incentive compatibility conditions are not too demanding when there are only two types. The planner can always design allocations in which no agent strictly prefers to mimic another agent's allocation. In this sense, it is possible to provide debt relief without encouraging any agents to overborrow. With a distribution of types, any debt relief policy will always induce some agent to borrow more. Debt relief may still be optimal, but the planner must now trade off the benefits of debt relief against the cost of distorting incentives towards overborrowing. Does the striking result presented above - that some debt relief is always Pareto improving when the zero lower bound binds - still hold?

To answer this question, I modify the model to include a continuous distribution of types. Agents have date 0 preferences $U(c, \theta) = \theta u(c)$, with $u(c) = \dfrac{c^{1-\sigma}}{1-\sigma}$. θ has a continuous density $f(\theta)$ with support $\Theta = [\underline{\theta}, \bar{\theta}]$. Equilibrium is defined as before.

As in the discrete type economy, the ZLB binds in equilibrium if agents are sufficiently impatient. Index households by $i \in [0, 1]$, and let $\theta(i) = F^{-1}(i)$ be the type of household i, so the most patient agent is 0, and he has type $\underline{\theta}$. The following proposition states that if we make each of the remaining agents $i > 0$ sufficiently impatient, then the ZLB binds in equilibrium: the impatient households borrow so much that the patient households accumulate large savings, and it would take a negative real interest rate at date 1 to make them consume all their wealth.

[19] Any tax greater than or equal to τ would suffice.

Proposition 5.1. *Take any sequence of functions $\theta_N : [0,1] \to [1,\infty)$ such that for all $N = 0,1,...$ $\theta_N(0) = \underline{\theta}$, θ_N is increasing, and for all $i \in (0,1]$, $\theta_N(i) \to \infty$ as $N \to \infty$. There exists N^* such that $r_1 < 0$ if $N > N^*$. Informally, the ZLB binds if agents' types are high enough.*

When the ZLB binds, equilibrium consumption falls below potential output, and it would be ex post Pareto improving to transfer wealth from the most patient household, whose consumption is limited by the ZLB, to an impatient household, who is liquidity constrained. As in the discrete type economy, such a transfer may not be incentive compatible. In fact, incentive compatibility is a much stronger constraint in the continuous type economy. Any transfer targeted at high θ individuals will induce some households with slightly lower θ to borrow more. To put this another way, any debt relief policy induces some overborrowing on the extensive margin, as well as overborrowing on the intensive margin.

5.1 Pareto problem

I now proceed to set up the Pareto problem in the continuous type economy. It is useful to write this problem in terms of households' compensated demand functions, which I now define.

Define the date 1 expenditure function to be

$$\begin{aligned} E(v_1, r_1) = \min_{c_1, c_2} c_1 + \frac{c_2}{(1+r_1)(1-\beta)} \\ \text{s.t. } u(c_1) + \frac{\beta}{1-\beta} u(c_2) \geq v_1 \\ c_2 \geq \underline{c}_2 \phi \end{aligned} \tag{EMP}$$

where $\underline{c}_2 = y^* - (1-\beta)\phi$. Let $C_1(v_1, r_1), C_2(v_1, r_1)$ be the solutions to this cost minimization problem.

These compensated demand functions will enter the Pareto problem: the planner will choose the value $v_1(\theta)$ to provide to each type, and the resource cost of providing this value at dates 1 and 2 will be $C_1(v_1, r_1)$, $C_2(v_1, r_1)$. It will be more intuitive, however, to relate the first order conditions of the planner's problem to the Marshallian uncompensated demand functions, which I now define. Let the date 1 value function to be

$$\begin{aligned} V(a_1, r_1) = \max_{c_1, c_2} u(c_1) + \frac{\beta}{1-\beta} u(c_2) \\ \text{s.t. } c_1 + \frac{c_2}{(1+r_1)(1-\beta)} \leq a_1 \\ c_2 \geq \underline{c}_2 \end{aligned} \tag{UMP}$$

Let $X_1(a_1, r_1), X_2(a_1, r_1)$ be the solutions to this utility maximization problem. The following results are, for the most part, standard.

Lemma 5.2. 1. ***Convexity.*** *E is convex in v_1. V is concave in a_1.*

2. ***Duality.*** $E(V(a_1, r_1), r_1) = a_1$. $V(E(v_1, r_1), r_1) = v_1$. $C_t(v_1, r_1) = X_t(E(v_1, r_1), r_1), t = 1, 2$.

3. ***Envelope theorem.*** $V_a(a_1, r_1) = u'(c_1)$.

4. ***Marginal cost of utility and MPC.*** $\dfrac{\partial C_t(v_1, r_1)}{\partial v_1} = \dfrac{\partial X_t(E(v_1, r_1), r_1)/\partial a_1}{u'(c_1)}$, $t = 1, 2$.

5. ***Borrowing constraint eventually binds.*** *There exists* $\bar{v}_1(r_1)$ *such that*

$$C_2(v_1, r_1) = \underline{c}_2 \text{ if } v_1 \leq \bar{v}_1(r_1).$$

   There exists $\bar{a}_1(r_1) = E(\bar{v}_1(r_1), r_1)$ *such that*

$$X_2(a_1, r_1) = \underline{c}_2, \frac{\partial X_1}{\partial a_1} = 1, \frac{\partial X_2}{\partial a_1} = 0 \text{ if } a_1 < \bar{a}_1(r_1)$$

6. ***Unconstrained CRRA households have constant MPC.*** *If* $u(c) = \dfrac{c^{1-\sigma}}{1-\sigma}$, *then if* $a_1 > \bar{a}_1(r_1)$,

$$\frac{\partial X_1}{\partial a_1} = \frac{1-\beta}{1 - \beta + \beta^{1/\sigma}(1+r_1)^{\frac{1}{\sigma}-1}}$$

7. ***Convex savings function.*** *If* $X_1(a_1, r_1)$ *is concave in* a_1, $C_2(v_1, r_1)$ *is convex in* v_1. *A sufficient condition for this is that* $u(c) = \dfrac{c^{1-\sigma}}{1-\sigma}$.

With these definitions in hand, I consider the social planner problem. The social planner puts weight $a(\theta)$ on type θ households. As before, the planner faces resource and incentive compatibility constraints. Now, however, the borrowing constraint and the constraints imposed by liquidity unconstrained households' Euler equations are embodied in the Hicksian demand functions in constraints (RC1) and (RC2). I write the zero lower bound constraint (ZLB) explicitly as a constraint on the real interest rate.

$$\max_{c_0, v_1, r_1} \int a(\theta)\left[\theta u(c_0(\theta)) + \beta v_1(\theta)\right] \mathrm{d}\theta \tag{PP}$$

$$\text{s.t.} \int c_0(\theta) f(\theta)\, \mathrm{d}\theta \leq y^* \tag{RC0}$$

$$\int C_1(v_1(\theta), r_1) f(\theta)\, \mathrm{d}\theta \leq y^* \tag{RC1}$$

$$\int C_2(v_1(\theta), r_1) f(\theta)\, \mathrm{d}\theta \leq y^* \tag{RC2}$$

$$\theta u(c_0(\theta)) + \beta v_1(\theta) \geq \theta u(c_0(\hat{\theta})) + \beta v_1(\hat{\theta}), \forall \theta, \hat{\theta} \tag{IC}$$

$$r_1 \geq 0 \tag{ZLB}$$

First I transform this problem into an equivalent, concave programming problem; I then characterize solutions to the problem using Lagrangian theorems.[20] It is convenient to work

[20] The Lagrangian optimization approach used here follows that of Amador et al. [2006].

in terms of utilities, rather than consumption allocations. Define the convex, increasing cost of utility function $C_0(u_0) = u^{-1}(u_0)$, and the date 0 value function $v(\theta) = \theta u_0(\theta) + \beta v_1(\theta)$. It is possible to express the incentive compatibility constraint (IC) as an integral condition, using the result of Milgrom and Segal [2002]. u_0, v_1 satisfies (IC) if and only if

$$v(\theta) = v(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z \tag{10}$$

and u_0 is nondecreasing.

Instead of choosing functions u_0 and v_1, we can equivalently choose a function u_0 and a scalar $\underline{v} := v(\underline{\theta})$, subject to the constraint that $u_0 \in \Omega$, the space of nondecreasing functions. v, v_1 are then implicitly defined by

$$\begin{aligned} v(\theta) &= \underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z \\ v_1(\theta) &= \beta^{-1}[\underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z - \theta u_0(\theta)] \end{aligned}$$

The objective function can be rewritten as

$$\int a(\theta) v(\theta)\,\mathrm{d}\theta = \underline{v} + \int (1 - A(\theta)) u_0(\theta)\,\mathrm{d}\theta$$

where $A(\theta) := \int_{\underline{\theta}}^{\theta} a(z)\,\mathrm{d}z$, and I normalize $\int_{\underline{\theta}}^{\bar{\theta}} a(\theta)\,\mathrm{d}\theta = 1$.

Putting all this together, we can rewrite the social planner's problem as

$$\mathcal{W}^* = \max_{u_0 \in \Omega, \underline{v}, r_1} \underline{v} + \int (1 - A(\theta)) u_0(\theta)\,\mathrm{d}\theta \tag{PP'}$$

$$\text{s.t} \int C_0(u_0(\theta)) f(\theta)\,\mathrm{d}\theta \leq y^* \tag{RC0'}$$

$$\int C_1\left(\beta^{-1}\left[\underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z - \theta u_0(\theta)\right], r_1\right) f(\theta)\,\mathrm{d}\theta \leq y^* \tag{RC1'}$$

$$\int C_2\left(\beta^{-1}\left[\underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z - \theta u_0(\theta)\right], r_1\right) f(\theta)\,\mathrm{d}\theta \leq y^* \tag{RC2'}$$

$$r_1 \geq 0 \tag{ZLB}$$

5.2 Necessary and sufficient conditions

It is now almost possible to apply Lagrangian theorems to (PP'). There are two remaining problems. First, the date 1 consumption function $C_1(v_1, r_1)$ is not convex in v_1: it has a kink at $\bar{v}_1(r_1)$. Second, the consumption functions need not be convex in r_1. In the following proposition, I show that it is nonetheless possible to use Lagrangian methods.

Proposition 5.3. *$u_0, \underline{v}, r_1$ solves (PP') if and only if there exist Lagrange multipliers $\lambda_0, \lambda_1, \lambda_2$ such that*

$u_0, \underline{v}$ solve

$$\mathcal{W}^* = \max_{u_0 \in \Omega, \underline{v}} \underline{v} + \int (1 - A(\theta)) u_0(\theta) \, \mathrm{d}\theta - \lambda_0 \int C_0(u_0(\theta)) f(\theta) \, \mathrm{d}\theta \qquad (11)$$
$$- \int M \left(\beta^{-1} \left[\underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z) \, \mathrm{d}z - \theta u_0(\theta) \right] | \lambda_1, \lambda_2, r_1 \right) f(\theta) \, \mathrm{d}\theta$$

where

$$M(v_1 | \lambda_1, \lambda_2, r_1) := \lambda_1 C_1(v_1, r_1) + \lambda_2 C_2(v_1, r_1)$$

and the Lagrange multipliers satisfy the following conditions. If the ZLB is slack, $\frac{\lambda_1}{\lambda_2} = (1 + r_1)(1 - \beta)$. If the ZLB binds, $\frac{\lambda_1}{\lambda_2} < 1 - \beta$.

Before providing a sketch of the proof, I explain how to interpret the function M. $M(v_1 | \lambda_1, \lambda_2, r_1)$ represents the total social cost of providing date 1 utility v_1 to a household, given that the household will choose its spending at dates 1 and 2 according to the interest rate r_1, and given that the social planner's shadow price of date 1 and date 2 output are λ_1 and λ_2 respectively. When $\frac{\lambda_1}{\lambda_2} = (1 + r_1)(1 - \beta)$, the relative shadow price of date 1 and 2 output, from the planner's perspective, is the same as the relative price of output faced by agents. In this case, M is simply a rescaled version of the expenditure function $M = \lambda_1 E(v_1, r_1)$. Proposition 5.3 states that the planner sets interest rates to equalize the private and social relative price of output, whenever this is not prevented by the ZLB.

However, when $\frac{\lambda_1}{\lambda_2} < (1 + r_1)(1 - \beta)$, the planner perceives that date 1 output is socially cheaper than date 2 output - the economy is in recession at date 1 - but private agents do not internalize this, because the relative price of date 1 output is still too high, because of the ZLB. This provides a motive for the planner to redistribute utility towards households with a higher propensity to spend at date 1 consumption, when consumption is socially cheap. In this economy, households with relatively low date 1 utility and wealth (i.e. with $v_1 \leq \bar{v}_1$) have a higher propensity to spend at date 1. In fact, their consumption functions are kinked at $\bar{v}_1$, which means that M is kinked at $\bar{v}_1$. Given that some households have date 1 utility $v_1(\theta)$ below $\bar{v}_1$, the planner would like $v_1(\theta)$ to be relatively high, since it is relatively cheap to supply this utility. As a result, it may be optimal to redistribute towards households with a higher propensity to consume date 1 consumption (which is socially cheap), or to give those households incentives to save at date zero so they have more wealth to spend at date 1.[21]

The proof of Proposition 5.3 has six steps. First we show that solutions to (PP') also solve a modified problem in which we replace the date 1 resource constraint with the aggregate expenditure function. Second, the modified problem can be solved in two stages: first maximize social welfare given r_1, yielding welfare $\mathcal{W}(r)$, and then choose r to maximize $\mathcal{W}(r)$ subject to the ZLB.

[21]This is exactly the result of Farhi and Werning [2013]. Relative to their framework, however, here the social planner faces additional incentive compatibility constraints resulting from private information, which make it harder to redistribute.

Third, the first stage of this problem is concave, and Lagrangian theorems (Luenberger [1969]) apply. Fourth, we can also express the expenditure functions as maximized sub-Lagrangians. Substituting these sub-Lagrangians into the main Lagrangian, we see that $W(r)$ is also the maximum of an expanded Lagrangian. Fifth, returning to our two stage problem, we can switch the order of maximization, first choosing r to minimize a certain function, subject to the ZLB, and then choosing utilities to maximize social welfare. Sixth, and finally, I show that when the ZLB is slack, one constraint in the planner's problem becomes slack, and the expanded Lagrangian is equivalent to (11), with $\frac{\lambda_1}{\lambda_2} = (1+r_1)(1-\beta)$. When the ZLB binds, we have $\frac{\lambda_1}{\lambda_2} < 1-\beta$.

It is possible to express necessary and sufficient conditions for an optimum in terms of Gateaux differentials of the Lagrangian.[22] Before doing so, it is first necessary to show that these differentials can be computed.

Lemma 5.4. *The Gateaux differential of the Lagrangian* (11) *is*

$$\begin{aligned}\delta\mathcal{L}(u_0, \underline{v}; \Delta_0, \underline{\Delta}) = &\underline{\Delta} + \int (1 - A(\theta))\Delta_0(\theta)\,\mathrm{d}\theta - \lambda_0 \int C_0'(u_0(\theta))\Delta_0(\theta)f(\theta)\,\mathrm{d}\theta \\ &- \int_{\Theta_+} M_+'(v_1(\theta)|\lambda_1, \lambda_2, r_1)\beta^{-1}\left[\underline{\Delta} + \int_{\underline{\theta}}^{\theta} \Delta_0(z)\,\mathrm{d}z - \theta\Delta_0(\theta)\right] f(\theta)\,\mathrm{d}\theta \\ &- \int_{\Theta_-} M_-'(v_1(\theta)|\lambda_1, \lambda_2, r_1)\beta^{-1}\left[\underline{\Delta} + \int_{\underline{\theta}}^{\theta} \Delta_0(z)\,\mathrm{d}z - \theta\Delta_0(\theta)\right] f(\theta)\,\mathrm{d}\theta\end{aligned}$$

where $v_1(\theta) = \beta^{-1}\left[\underline{v} + \int_{\underline{\theta}}^{\theta} u_0(z)\,\mathrm{d}z - \theta u_0(\theta)\right]$, *and*

$$\begin{aligned}\Theta_+ &= \left\{\theta \in \Theta : \underline{\Delta} + \int_{\underline{\theta}}^{\theta} \Delta_0(z)\,\mathrm{d}z - \theta\Delta_0(\theta) > 0\right\} \\ \Theta_- &= \left\{\theta \in \Theta : \underline{\Delta} + \int_{\underline{\theta}}^{\theta} \Delta_0(z)\,\mathrm{d}z - \theta\Delta_0(\theta) < 0\right\}\end{aligned}$$

Putting all these results together, we can now characterize constrained efficient allocations in terms of first order conditions.

Lemma 5.5. $u_0, \underline{v}, r_1$ *solves* (PP′) *if and only if there exist Lagrange multipliers* $\lambda_0, \lambda_1, \lambda_2$ *such that, for all* $\underline{\Delta}, \Delta_0$ *such that* $u_0 + \Delta_0 \in \Omega$,

$$\delta\mathcal{L}(u_0, \underline{v}; \Delta_0, \underline{\Delta}) \le 0$$

where $\delta\mathcal{L}(u_0, \underline{v}; \Delta_0, \underline{\Delta})$ *is defined as in Lemma 5.4, and the Lagrange multipliers satisfy the following conditions. If the ZLB is slack,* $\frac{\lambda_1}{\lambda_2} = (1+r_1)(1-\beta)$. *If the ZLB binds,* $\frac{\lambda_1}{\lambda_2} < 1-\beta$.

[22]Given a real valued functional f defined on a vector space X, if the limit

$$\lim_{\alpha \downarrow 0} \frac{1}{\alpha}[f(x + \alpha h) - f(x)]$$

exists, then it is called the Gateaux differential of f at x with increment h and is denoted by $\delta f(x; h)$.

5.3 Constrained efficient debt relief

I now show that the main results from the two type economy considered above generalize to a continuous distribution of types. When the ZLB binds in the equilibrium without policy, the equilibrium is constrained inefficient. Somewhat surprisingly, the simple, piecewise linear debt relief transfer functions which implemented optimal allocations in the two type economy also implement some (but by no means all) optimal allocations with a continuous distribution of types. In particular, if a debt relief transfer function has a positive marginal tax above the cap $\tau > 0$, and implements full employment, then it is constrained efficient.

Proposition 5.6. *If the ZLB binds in the equilibrium without policy:*

1. *The equilibrium is constrained inefficient.*
2. *Debt relief transfer functions with $\tau > 0$ implementing full employment allocations are constrained efficient.*

If the ZLB is slack, the equilibrium is constrained efficient.

A sketch of the proof of part 1 is as follows. Suppose by contradiction that a competitive equilibrium in which the ZLB binds, and the date 1 resource constraint is slack, is constrained efficient. Consider the following deviation: make a small transfer to households whose date 0 consumption is such that they are 'just' borrowing constrained. Households with slightly higher date 0 consumption will consume less at date 0, and more at date 1, to qualify for the transfer; households with slightly lower date 0 consumption will consume more at date 0, and less at dates 1 and 2. To first order, the effect on date 0 consumption cancels out, but aggregate date 1 consumption increases, and date 2 consumption falls. This is clearly feasible, since the date 1 resource constraint is slack, and it increases utility for the households who change their behavior. Thus the original allocation cannot have been constrained efficient.

To show that debt relief transfer functions with $\tau > 0$ implement constrained efficient allocations, two steps are necessary. First, it is necessary to show that such allocations exist. The proof proceeds by showing that aggregate consumption demand at dates 0,1 and 2 is a continuous function of the parameters of the debt relief transfer function, $\bar{T}$, $\underline{d}$, $\bar{d}$ and τ. A fixed point argument then shows that there exists a transfer function implementing a full employment allocation. Second, it is necessary to show that such allocations, if they exist, are constrained efficient. The proof proceeds by showing that the allocations satisfy the condition in Lemma 5.5.

5.4 How the optimal policy prevents overborrowing

Debt relief encourages overborrowing on the intensive margin, inducing borrowers who had enough debt to qualify for the transfer, even in the equilibrium without policy, to borrow and consume more through a wealth effect. In addition, any debt relief policy induces some overborrowing on the extensive margin. That is, if households with higher debt receive a higher transfer, some households will take on more debt in order to benefit from this transfer. This was not true in the two type economy: in that economy, there was always some room to give borrowers a transfer without encouraging savers to overborrow.

Figure 13 illustrates. Although there are a continuum of types, to simplify the figure I only show the indifference curves of two types, θ_M and $\theta_B > \theta_M$. Suppose households' debt is written off one-for-one in some range, shifting the budget set to the right as shown in the figure. This induces some households to increase their date zero consumption and borrowing. For high types like θ_B, debt relief acts through a wealth effect, and these types increase their date 0 and date 1 consumption. Intermediate types like θ_M, however, will borrow up to the kink in the budget constraint. Thus there is overborrowing on both an intensive and an extensive margin. Finally, there are some more patient types who are not affected by the policy.

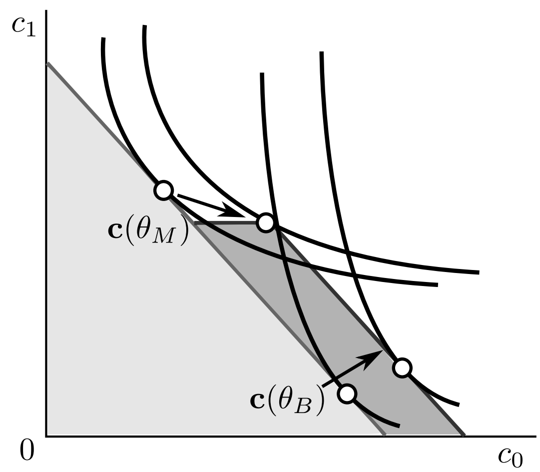

Figure 13: Debt relief induces intensive and extensive margin overborrowing

The budget set shown in 13 does not implement a feasible allocation, since aggregate date 0 consumption has increased and markets do not clear. To prevent overborrowing, and clear markets, the government can tax borrowing above the cap at rate $\tau > 0$. This reduces the date 0 consumption and borrowing of the most impatient households, such as θ_B. That is, this debt relief policy reduces the debt of extreme borrowers, through a marginal tax on debt, to balance out the overborrowing of moderate borrowers such as θ_M. Figure 14 illustrates.

By increasing the transfer to moderate borrowers, and increasing the tax on extreme borrowers to ensure markets clear at date 0, it is always possible to return the economy to full employment at date 1. However, the further the economy is from full employment, the larger the transfer required to return it to full employment, and the higher the marginal tax on debt required to balance out the overborrowing induced by this transfer. It is possible that the tax on debt required to prevent overborrowing is so high that some impatient types (such as θ_B) are made worse off by this policy. Debt relief transfer functions which implement full employment are always constrained efficient, but they may not be Pareto improving. Figure 15 illustrates such a case.

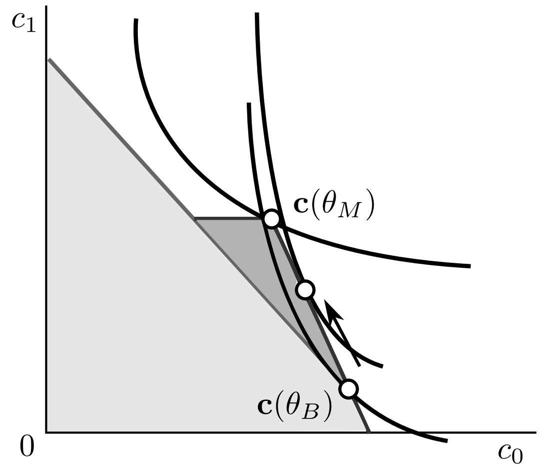

Figure 14: τ prevents overborrowing

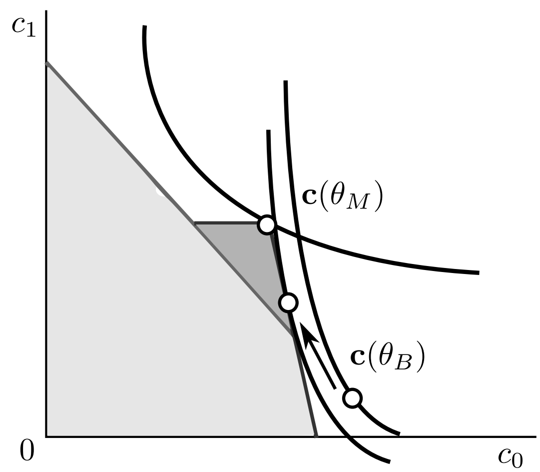

Figure 15: Debt relief transfer function may not be Pareto improving

5.5 Numerical exercise

How large are the optimal transfer and the marginal debt tax τ likely to be? Is linear debt relief policy likely to be Pareto improving? In this section I present a numerical exercise to give a rough answer to these questions.

I interpret 1 period as 5 years. θ is lognormal, with $\ln \theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$. I set $\beta = 0.975$, corresponding to a steady state real interest rate of 0.5%. I set $\phi = 0.2$. I then vary σ, the inverse of the intertemporal elasticity of substitution, and choose $\mu_\theta, \sigma_\theta^2$ to roughly match the 2008-2012 fall in output and deleveraging, and the 2007 distribution of debt, in the United States.

Potential output y^* is normalized to 1. The first variable to compare to the data is y_1, the shortfall in output relative to potential output. I measure potential output by fitting either a linear trend or an exponential trend to 1984-2007 real GDP. I then measure y_1 as the average value of annual output divided by potential output over 2008-2012. Using a linear trend, this yields $y_1 = 0.95$; using an exponential trend yields $y_1 = 0.90$. In the table below, I report a simple average of these, $y_1 = 0.92$.

Data on aggregate household debt comes from the Financial Accounts of the United States. I divide total household debt by 5 times trend annual GDP. Since I interpret the crisis period $t = 1$ as 2008-2012, I interpret aggregate debt at date 1, $D_1 = \int d_1^i \mathbb{1}\{d_1^i \geq 0\}\,\mathrm{d}i$ as household debt in 2008, and aggregate debt at date 2, $D_2 = \int d_2^i \mathbb{1}\{d_2^i \geq 0\}\,\mathrm{d}i$ as household debt in 2012. This yields $D_1 = 0.19, D_2 = 0.14$.

Data on the distribution of debt comes from the 2007 Survey of Consumer Finances. I restrict the sample to heads of household aged between 25 and 65 who are not students and who are not retired. I interpret d_t^i in the model as total household debt minus financial assets, divided by 5 times the average family income for households in the sample.[23]. As explained above, whether simple linear debt relief policies are Pareto improving depends crucially on the right tail of the distribution of debt. I therefore attempt to roughly match the 90th percentile of the distribution of the debt.

To solve the model, I approximate the distribution of types as a discrete type economy with 500 types, using an unequally spaced grid for θ. I truncate the distribution of θ, setting $\underline{\theta} = 0.05$, $\bar{\theta} = 20$. For each set of parameter values, I first verify that the ZLB binds in equilibrium, then search for the debt relief transfer policy that ensures full employment at date 1, while keeping the date 0 rate of interest, $1 + r_0$, the same as in the equilibrium without policy. In what follows, I call this the 'optimal policy', but it is important to bear in mind that there are a large set of constrained efficient policies, and this is only one of them.

Table 1: Optimal policy

Parameters			**Moments**				**Policy**			
σ	μ_θ	σ_θ	y_1	D_1	D_2	p_{90}	$\bar{T}$	$\underline{d}$	$\bar{d}$	τ
0.5	0.1	0.3	0.94	0.16	0.11	0.46	0.00	0.15	0.17	0.16
1	0.5	0.5	0.92	0.16	0.11	0.44	0.02	0.16	0.20	0.53
2	1	1	0.93	0.14	0.10	0.40	0.03	0.16	0.23	1.66
			Data							
			0.92	0.19	0.14	0.42				

Table 1 presents the results. The maximum transfer is $2 - 4\%$ of 5 years' income, or $\$8,000 - 16,000$. The marginal tax above the cap is a $3 - 33\%$ annual spread: it is higher if the IES, $\frac{1}{\sigma}$, is low.

Next, I evaluate the welfare effect of these policies. Following Lucas [1987], define the welfare benefit of debt relief for type θ, $\lambda(\theta)$, as the percentage increase in consumption, in every period t, such that household θ would be indifferent between the equilibrium without policy (plus the percentage increase in consumption) and the equilibrium with policy. Let c_t^* denote allocations

[23]The SCF collects data from two samples: a standard multistage area-probability sample selected from the 48 contiguous US states, and a list sample designed to disproportionately sample wealthy families. The SCF provides probability weights which account for the sample design, and also for differential patterns of non-response among families with different characteristics. All SCF data presented here is weighted using these probability weights.

in the equilibrium without policy, and c_t^{**} allocations in the equilibrium with policy. For each type θ, $\lambda(\theta)$ solves

$$\theta u(c_0^{**}(\theta)) + \sum_{t=1}^{\infty} \beta^t u(c_t^{**}(\theta)) = \theta u((1+\lambda(\theta))c_0^*(\theta)) + \sum_{t=1}^{\infty} \beta^t u((1+\lambda(\theta))c_t^*(\theta))$$

Figure 16 plots $\lambda(\theta)$, for $\sigma = 0.5, 1, 2$. Since the distribution of types is different under the three experiments, I plot $\lambda(\theta)$ against $F(\theta)$, the rank of type θ in the whole distribution. In each case, the optimal debt relief policy increases welfare for most households, particularly for 'moderate' types with intermediate values of θ. However, the optimal policy is never Pareto improving: extremely high types are made worse off by debt relief with a cap. These households have an extremely high demand for date 0 consumption; anything that reduces date 0 consumption, such as a marginal tax on debt, makes them worse off.

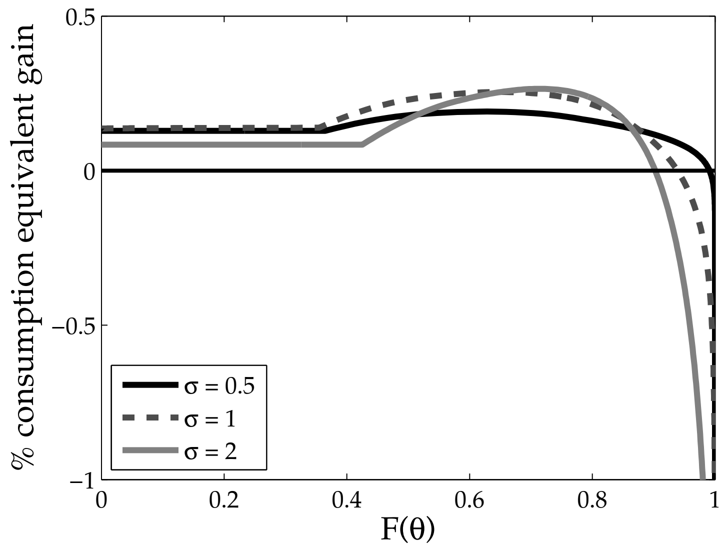

Figure 16: Consumption equivalent gains under optimal debt relief

This result should be interpreted with caution, for two reasons. First, the exercise here is to interpret the empirical distribution of debt in 2007 through the lens of a model in which all households started with zero debt in 2002, and in some cases took out large amounts of debt, knowing with certainty that the financial crisis would happen in 2008, and they would be forced to deleverage. Clearly, such a model requires substantial dispersion in preferences θ to generate the significant wealth inequality observed in the data. In reality, households did not expect the crisis to happen with probability one; had they done so, they would surely have taken out less debt. Second, under the more realistic assumption that households do not expect the crisis to happen for sure (which I consider in Section 7.1 below), the incentive problems associated with debt relief are less severe, and it is easier to find an ex ante Pareto improving policy. In the extreme case, if the crisis is a zero probability event, there are no incentive concerns associated with debt relief. For both these reasons, one should not necessarily conclude that it is hard to design Pareto improving debt relief policies.

However, Figure 16 does highlight that any policy to reduce overborrowing runs the risk of harming households who need to borrow.[24] This risk is even greater for alternative macroprudential policies, which attempt solely to prevent overborrowing, as I discuss in section 6 below. Debt relief with a cap is less of a culprit in this regard, since it combines a transfer to indebted households with a marginal tax to reduce their borrowing.

6 Macroprudential policy

In this section I compare debt relief to macroprudential policies. Korinek and Simsek [2014] and Farhi and Werning [2013] have discussed the role of taxes on overborrowing, debt limits, and insurance requirements in preventing liquidity traps. Such macroprudential policies prevent overborrowing ex ante, while debt relief corrects overborrowing ex post. Korinek and Simsek [2014] and Farhi and Werning [2013] both consider optimal policy in the case when the planner can observe households' type. In contrast, I assume that type is private information.

6.1 Macroprudential policy with full information

Is macroprudential policy superior or inferior to debt relief, or are the two policies equivalent? To answer this question, I define equilibrium with two macroprudential policies considered by Korinek and Simsek [2014] and Farhi and Werning [2013] which are an alternative to debt relief: debt limits with compensating transfers at date 0, and linear debt taxes and transfers at date 1. I allow transfers and taxes to depend on an agent's type.

Definition 6.1. *An* ***equilibrium with macroprudential taxes*** *is a ZLB-constrained equilibrium with transfers in which transfer functions have the form* $T(d,\theta) = T_1(\theta) - \tau(\theta)d$.

An ***equilibrium with debt limits*** *is* $\{c_t^i, d_t^i, y_t, r_t\}$ *such that, given a date 0 debt limit* ϕ_0 *and date 0 transfers* $T_0(\theta)$*,*

1. *each household i chooses* $\{c_t^i, d_t^i\}$ *to maximize* (1) *s.t.* (2), (4)*, and*

$$
\begin{aligned}
d_0^i &= -T_0(\theta_i) \\
d_1^i &\leq \phi_0
\end{aligned}
$$

2. $c_t^B + c_t^S = 2y_t$

3. $r_t \geq 0$, $y_t \leq y^*$, $r_t(y^* - y_t) = 0$

Farhi and Werning [2013] and Korinek and Simsek [2014] show that these macroprudential policies implement first-best allocations when household type is observable to the planner. The following proposition merely restates their result.

[24] As Zinman [2014] emphasizes, we still understand little about why some households borrow as much as they do.

Proposition 6.2. *Consider the relaxed Pareto problem without constraints* (ICS) *and* (ICB).

1. *Every solution to the relaxed Pareto problem can be implemented as an equilibrium with macroprudential taxes. If the ZLB does not bind,* $\tau(\theta_S) = \tau(\theta_B) = 0$*. If the ZLB binds,*

$$\frac{1 + \tau(\theta_B)}{1 + \tau(\theta_S)} = \frac{u'(c_1^S)}{u'(c_1^B)} \frac{U_c(c_0^B, \theta_B)}{U_c(c_0^S, \theta_S)} > 1 \tag{12}$$

   In particular, the optimal allocation can be implemented with a tax on debt targeted only at borrowers ($\tau(\theta_B) > 0, \tau(\theta_S) = 0$)*.*

2. *Every solution to the relaxed Pareto problem can be implemented as an equilibrium with debt limits. If the ZLB binds, then the debt limit binds, and it is equal to*

$$\phi_0 = \bar{d}_1 := \bar{c}_1^S - y^* + \phi \tag{13}$$

Since these transfers depend on a household's type θ, they cannot be implemented when θ is private information. Next, I study macroprudential policies under private information, and compare them to the ex post policies considered throughout the paper. Once we restrict the simple linear taxes and transfers considered in Proposition 6.2 to be anonymous, macroprudential taxes are sub-optimal, and debt limits only implement particular optimal allocations. However, modified macroprudential policies with nonlinear taxes and transfers are equivalent to ex post policies, and implement optimal allocations. Thus the results in Sections 4 and 4.2 can also be interpreted as describing optimal macroprudential policy under private information.

6.2 Macroprudential taxes

First, I study macroprudential taxes under private information. Under full information, linear debt taxes implement optimal allocations, and satisfy equation (12) (only borrowers face a tax on debt). Under private information, linear debt taxes are not optimal. Nonlinear debt taxes implement optimal allocations, since they are equivalent to the debt-contingent transfers considered throughout this paper; however optimal marginal tax rates may not satisfy equation (12).

Anonymous *linear* taxes on debt which do not depend on a household's type are redundant: they change the equilibrium interest rate, but implement the same allocations as an equilibrium without policy.[25] To see this, note that taxes only enter households' Euler equations:

$$(1 + r_0)(1 + \tau) = \frac{U_c(c_0^i, \theta_i)}{\beta u'(c_1^i)}$$

If τ is increased, r_0 falls to clear markets, and the equilibrium is unchanged. When this equilibrium is inefficient, anonymous linear taxes fail to implement optimal allocations.

[25] This is noted by Farhi and Werning [2013] who emphasize that only the relative financial taxes faced by different agents affect the allocation: the average level of taxes is indeterminate.

If anonymous macroprudential taxes are to implement optimal allocations, they must be nonlinear. Nonlinear taxes on debt are equivalent to the debt-contingent transfers considered throughout this paper, and implement the same allocations. All the results above about optimal debt-contingent transfers can be interpreted in terms of nonlinear macroprudential taxes. Equation (12), which describes the optimal marginal debt tax under full information, must now be modified to take into account incentive constraints. We can interpret $-T'(d)$ as the analogue of τ, the marginal tax on debt. (The negative sign is present because $T(d) > 0$ denotes a positive transfer, whereas $\tau > 0$ denotes a positive tax on debt.)

Proposition 6.3. *1. If* (ZLB), (ICS) *and* (ICB) *do not bind,* $T'(d_1^S) = T'(d_1^B) = 0$.[26]

2. *If* (ZLB) *binds and neither incentive constraint binds,* $T'(d_1^S) > T'(d_1^B)$.

3. *If* (ZLB) *does not bind and either* (ICS) *or* (ICB) *binds,* $T'(d_1^S) < T'(d_1^B)$.

4. *If* (ZLB) *binds and either* (ICS) *or* (ICB) *binds, we may have* $T'(d_1^S) < T'(d_1^B)$, $T'(d_1^S) > T'(d_1^B)$, *or* $T'(d_1^S) = T'(d_1^B) = 0$.

If neither incentive constraint binds, and (ZLB) binds, we want to induce borrowers to take on less debt, since excessive debt imposes a macroeconomic externality. To do this, we impose a marginal tax on debt, as in the case with linear macroprudential taxes and perfect information. However, when incentive constraints bind, marginal taxes play a different role: they make the allocations of savers and borrowers different, so that savers do not want to mimic borrowers (or vice versa). If for example (ICS) binds, we distort allocations so that $\frac{\beta u'(c_1^S)}{U_c(c_0^S, \theta_S)} < \frac{\beta u'(c_1^B)}{U_c(c_0^B, \theta_B)}$. Absent incentive constraints, it would be Pareto improving to have B consume more at date 1, and have S consume more at date 0. But this would make B's allocation more attractive to S, who values date 1 consumption more. To deter S, who prefers later consumption, from mimicking B, optimal policy front-loads B's consumption by offering him a *lower* marginal interest rate than S, or a *negative* marginal tax on debt. When both (ZLB) and one incentive constraint bind, both motives are in play, and the sign of the marginal tax on debt is ambiguous.

6.3 Debt limits and date 0 transfers

Next, I consider debt limits under private information. Under full information, debt limits together with compensating transfers implement optimal allocations, and satisfy equation (13) (debt is low enough to ensure full employment). Under private information, debt limits implement one particular optimal allocation, which may not be Pareto improving. A modified version of the debt limit policy, with debt-contingent date 0 transfers, implements any optimal allocation. However, the optimal debt limit may not satisfy equation (13).

If transfers do not depend on a household's type, they must be equal to zero, under our maintained assumption that the government runs a balanced budget. A debt limit without any

[26] $T'(d)$ denotes either the left-hand derivative or the right-hand derivative of $T(d)$; as noted above, allocations in which an incentive constraint binds cannot be implemented with a differentiable value function.

compensating transfer implements one particular constrained efficient allocation. This policy is better for savers, and worse for borrowers, relative to debt relief. It is always Pareto efficient, but it may not be Pareto improving relative to the equilibrium without policy. If borrowers are sufficiently impatient, they will be worse off with debt limits than in the equilibrium without policy: the cost of lower date 0 consumption outweighs the benefit of higher date 1 consumption.

Debt limits therefore appear to be more restrictive than debt-contingent transfers. However, if we allow the date 0 compensating transfers to depend on debt, debt limits implement the same set of allocations as date 1 debt-contingent transfers.

Proposition 6.4. *1. Every constrained efficient allocation can be implemented as an allocation with debt limits and date 0 debt-contingent transfers. The debt limit may be greater than $\bar{d}_1$ if (ICS) binds.*

2. A debt limit $\phi_0 = \bar{d}_1$, together with no transfer at date 0 ($T_0(d) = 0$), implements constrained efficient allocations corresponding to a weight of $\bar{\alpha}(\theta_B)$ in the social planner's problem.

3. There exists $\bar{\theta}$ such that this allocation is not Pareto-improving, relative to the equilibrium without policy, if $\theta_B > \bar{\theta}$.

4. If the ZLB binds, the debt limit is always binding and equal to $\bar{d}_1$, unless (ZLB) and (ICS) both bind. In this case, $d_1^B > \bar{d}_1$, and there is underemployment at date 1.

Part 1 of this Proposition states that date 0 transfers and date 1 transfers are equivalent: either policy defines a nonlinear mapping from date 0 consumption to date 1 consumption. Part 2 states that a debt limit without compensating transfers is Pareto efficient at the ZLB. However, part 3 states that this allocation is not Pareto improving if θ_B is sufficiently large. A binding debt limit reduces borrowers' date 0 consumption, but increases date 1 consumption. If B is sufficiently impatient, a fall in date 0 consumption is very costly, and the binding debt limit reduces her welfare (although it increases S's welfare). Similarly, part 4 notes that some Pareto efficient allocations in which the ZLB binds cannot be implemented with a debt limit equal to $\bar{d}_1$. These are the efficient allocations in which there is underemployment, and $c_1^S + c_1^B < 2y^*$.[27]

Date 0 transfers are equivalent to date 1 transfers, and can be used to induce exactly the same budget sets. In particular, we can construct date 0 transfer functions which induce exactly the same budget sets as debt relief transfer functions. These date 0 transfers can be interpreted as targeted loan support programs, combined with macroprudential taxes on excessive borrowing.

Definition 6.5. *$T_0(d)$ is a **targeted loan support program** if it has the form*

$$\begin{aligned} T^{LS}(d) &= -\bar{T} \text{ if } d < d^* \\ &= T^* - \tau d \text{ if } d \geq d^* \end{aligned}$$

where $T^*, \bar{T} > 0, d^*, \tau$ are parameters.

[27] As discussed in Section 4, when (ICS) and (ZLB) both bind, some date 1 unemployment may be optimal.

There is a lump sum tax $\bar{T}$ on all households who borrow less than a certain amount d^*. The government gives a subsidy of T^* to households who borrow exactly d^*. Above that point, further borrowing is penalized: the transfer falls by τ dollars for each dollar of debt above the minimum level d^*.

Given that targeted loan support programs are isomorphic to debt relief transfer functions, the following result follows immediately from Propositions 4.4 and 4.7.

Proposition 6.6. *1. Targeted loan support implements the optimal allocation iff $\alpha < \alpha(\theta_B)$.*

2. *If (ICS) binds, $\alpha < \alpha(\theta_B)$ and targeted loan support implements the optimal allocation.*

3. *If $\theta_B \leq \theta^{ZLB}$, the competitive equilibrium is Pareto optimal (it is the solution to the planner's problem with $\alpha = \alpha(\theta_B)$. Targeted loan support is not Pareto improving.*

4. *If $\theta_B > \theta^{ZLB}$, the competitive equilibrium is Pareto inefficient. Targeted loan support is always Pareto improving.*

Targeted loan support programs, like ex post debt relief with a cap, implement efficient allocations which are relatively favorable for borrowers. Unlike ex ante debt limits without compensating transfers, targeted loan support compensates borrowers for the reduction in their ability to borrow, while the tax on debt still discourages overborrowing.

6.4 Debt limits with a distribution of types

Debt limits, like other forms of rationing, force agents who would otherwise borrow different amounts to borrow the same amount. This inefficiency is absent in an economy with only one type of borrower. To properly compare ex ante debt limits and ex post debt relief, I return to the economy with a distribution of types considered in Section 5, and ask whether a binding debt limit without compensating transfers is constrained efficient.[28] The following proposition states that the debt limit is inefficient if some households are sufficiently impatient.

Proposition 6.7. *Take any allocation in which a debt constraint binds for agents in some interval $[\theta^*, \bar{\theta}]$ at date 0. Suppose $\dfrac{f(\theta)}{u'(c_0(\theta))}$ is increasing. Then there exists θ^* such that if $\bar{\theta} > \theta^*$, the allocation is Pareto inefficient.*

Figure 17 illustrates a case where a debt limit is Pareto inefficient. The lightly shaded region shows the budget set with a debt limit, which binds for types $\theta \in [\theta_M, \theta_H]$. Construct a Pareto improving deviation as follows. Increase date 1 consumption for types $\theta \leq \theta_M$, rearranging their date 0 consumption so as to reduce it on average. If the date 1 resource constraint binds, pay for this increase in date 1 consumption by decreasing date 1 consumption for some high θ types (such as θ_H), compensating them by increasing their date 0 consumption. If these households

[28] Even with only two agents, we have already seen that a debt limit is not always Pareto improving, since it may make impatient borrowers worse off, by preventing them from borrowing.

are impatient enough, even a small increase in date 0 consumption compensates for a large fall in date 1 consumption, and our deviation is feasible; thus the original allocation was Pareto inefficient. The dark shaded area shows one transfer function implementing this deviation.

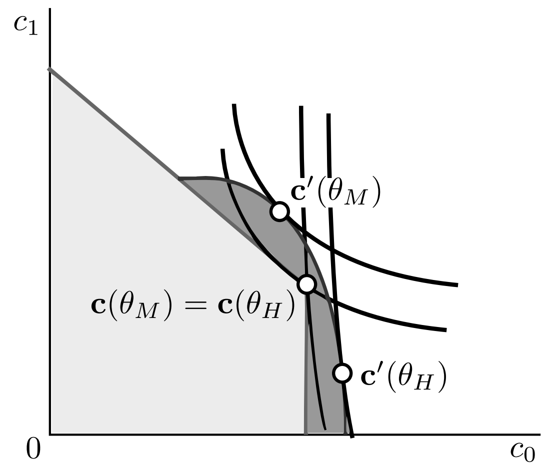

Figure 17: Inefficient debt limit and Pareto-improving deviation

While debt limits prevent overborrowing, they are a blunt instrument, preventing even the most impatient households from borrowing. This is unnecessary, since to prevent a debt-driven recession, it is only necessary to limit borrowers' *aggregate* debt. It may be better to allow impatient households to borrow, but tax them at a high rate. This is exactly what targeted loan support programs do. Again, the next result follows immediately from Proposition (5.6), given the equivalence of targeted loan support programs and debt relief transfer functions.

Proposition 6.8. *If the ZLB binds in the equilibrium without policy:*

1. *The equilibrium is constrained inefficient.*
2. *Targeted loan support programs with $\tau > 0$ implementing full employment allocations are constrained efficient.*

If the ZLB is slack, the equilibrium is constrained efficient.

7 Further questions

In this section, I consider three additional questions. First, how does optimal policy change if the crisis does not occur with probability one? Second, are the conclusions above robust to making borrowers and savers differ in their income, rather than their preferences? Third, are the conclusions robust to introducing endogenous labor supply?

Before considering these substantive extensions of the baseline model, I discuss how it can be reinterpreted. In the model presented above, borrowers and savers differ in their preferences. Consider the following three alternatives. First, suppose that individuals' date 0 income is $y^* - \theta_i$, where $\theta_i \in \mathbb{R}$ is a transitory income shock, unobservable to the planner. Agents with a negative shock will want to borrow; agents with a positive shock will save. Second, suppose agents have initial debt θ_i, which is unobservable to planner. Again, agents with high debt will seek to roll over some of this debt, paying it off gradually over time. Finally, suppose that in addition to purchasing nondurable consumption, agents have inelastic demand for a certain quantity of a 'necessary' consumption good (which could be housing, healthcare, etc.). I now show that all these cases are isomorphic to the economy considered above.

In all three cases, households' date 0 consumption is $c_0^i = y^* - \theta_i + \frac{d^i}{1+r_0}$. The planner only observes households' debt and can no longer infer their period 0 consumption, as in the model with purely preference-based heterogeneity. Consequently, the planner cannot choose consumption allocations for workers. Instead, the planner chooses $\hat{c}_0^i := c_0^i + \theta$, which can be inferred from a household's debt. Preferences over this object are $u(\hat{c}_0^i - \theta_i)$. This is equivalent to an economy in which θ_i is a taste shock, and agents have preferences

$$u(c_0^i - \theta_i) + \sum_{t=1}^{\infty} \beta^t u(c_t^i),$$

which is a special case of the baseline model with $U(c, \theta) = u(c - \theta)$.

I now extend the baseline model with two types to answer the questions raised above.[29] Recall the main conclusions from Section 4: efficient allocations have the structure described in Proposition 4.1; debt-contingent transfers (in particular, debt relief) implement efficient allocations; and debt relief is Pareto-improving when the ZLB binds. In each of the three extensions below, I ask whether these results still hold.

7.1 Probability of crisis less than one

I now extend the results in Section 4 to the case where the probability of crisis is less than 1. With probability π, the borrowing constraint permanently falls to ϕ at date 1, as before. With probability $1-\pi$, the borrowing constraint never binds.[30] If the crisis does not occur, the economy immediately converges to steady state. Letting hats denote variables in the non-crisis state, households have preferences

$$\theta_i u(c_0^i) + \pi \left[\beta u(c_1^i) + \frac{\beta^2}{1-\beta} u(c_2^i) \right] + (1-\pi) \frac{\beta}{1-\beta} u(\hat{c}_1^i)$$

Since there is now aggregate risk in the economy, it is necessary to specify the financial assets available to households. I consider two cases. In the **incomplete markets economy**, households trade a riskless bond, as before. In the **complete markets economy**, they trade in a complete set

[29] For simplicity, I focus on implementation with date 1 transfers; the results from Section 6 regarding implementation with macroprudential policy would carry over in the same way.

[30] I assume agents have the same, model-consistent expectations regarding the probability of crisis. Korinek and Simsek [2014] consider in detail the case where households have different expectations, and borrowers are more optimistic than savers.

of Arrow-Debreu securities. Let q_0 and $\hat{q}_0$ denote the price of securities paying one unit of the consumption good at date 1 in the crisis state and in the non-crisis state, respectively, and let d_1^i and $\hat{d}_1^i$ denote the securities of each type issued by household i.

At date 1, the transfers provided by the government now depend on the aggregate state, in addition to household borrowing. In the incomplete markets economy, the government offers two transfer functions, $T(d)$ in the crisis state and $\hat{T}(d)$ in the non-crisis state. In the complete markets economy, these transfers may depend on households' issuance of each security, so the functions have the form $T(d, \hat{d}), \hat{T}(d, \hat{d})$.

Naturally, since the crisis may not occur, debt relief distorts ex ante incentives less. The following proposition is elementary:

Proposition 7.1. *If $\pi = 0$, the full employment transfer is always incentive compatible.*

As before, I define a constrained efficient allocation as the solution a Pareto problem.[31] The following proposition states that solutions to the Pareto problem have the same structure as before, and can be implemented as equilibria with transfers.

Proposition 7.2. *In the economy with $\pi < 1$:*

1. *Constrained efficient allocations have the structure described in Proposition 4.1.*
2. *Every constrained efficient allocation can be implemented as an equilibrium with transfers in both the incomplete markets economy, and in the complete markets economy.*

As in Section 4.2, I now ask when debt relief implements optimal allocations. With $\pi < 1$ it is no longer possible to implement optimal allocations with simple piecewise linear transfer functions (the debt relief transfer functions defined above), so a broader definition of debt relief is necessary. I will say that there is debt relief in the crisis state if $d_1^B > 0 > d_1^S$ (B takes on debt while S saves) and $T(d_1^B) > T(d_1^S)$ (B receives a larger transfer than S in equilibrium). With these redefinitions, Proposition 4.4 remains valid.

Proposition 7.3. *There exists $\bar{\alpha}(\theta, \pi)$ such that*

1. *A transfer function with debt relief implements the optimal allocation if $\alpha < \bar{\alpha}(\theta, \pi)$*
2. *If (ICS) binds, $\alpha < \bar{\alpha}(\theta, \pi)$ and debt relief implements the optimal allocation.*

Finally, I ask whether debt relief is Pareto-improving, relative to the equilibrium without policy. The following Proposition confirms that Proposition 4.7 still holds in the complete markets economy: debt relief is generally purely redistributive, but is always Pareto-improving when the ZLB binds. However, in the incomplete markets economy, debt relief is Pareto improving even when the ZLB does not bind, provided that the borrowing constraint binds.

Proposition 7.4. 1. *If $\theta_B \leq \theta^{BC}$, the competitive equilibrium is Pareto optimal. Debt relief is not Pareto improving.*

[31] The full Pareto problem with $\pi < 1$ is presented in the Appendix.

2. *If $\theta_B \in (\theta^{BC}, \theta^{ZLB}]$, the competitive equilibrium is Pareto optimal under complete markets, and Pareto inefficient under incomplete markets. Debt relief is always Pareto improving in the incomplete markets economy.*

3. *If $\theta_B > \theta^{ZLB}$, the competitive equilibrium is Pareto inefficient. Debt relief is always Pareto improving.*

When markets are incomplete, debt relief can be Pareto improving even if the ZLB does not bind. The incomplete markets economy has aggregate risk: with probability π the borrowing constraint will tighten, but agents can only invest in a riskless bond ex ante, and cannot insure against this shock. This creates an additional rationale for transfers to borrowers, to help them smooth consumption when credit markets do not allow them to do so. The first-best allocation smooths consumption across states of the world: $c_1^i = \hat{c}_1^i$, $i = S, B$. When the borrowing constraint binds in the crisis state, this cannot be the equilibrium allocation, since for any level of debt d^1, $c_1^B = y^* + \phi - d_1$, while $\hat{c}_t^B = y^* - (1-\beta)d_1 > c_1^B$. In order to implement the first best, it is necessary to write off part of borrowers' debt. Debt relief, like bankruptcy, completes markets (Zame [1993]). In the complete markets economy this inefficiency disappears, and the competitive equilibrium is Pareto inefficient only when the ZLB binds.

Targeted transfers have two *distinct* macroeconomic roles (in addition to their purely redistributive role). First, when output is demand constrained, transfers can stimulate demand by redistributing resources to agents with a higher propensity to consume. Secondly, public transfers can substitute for private insurance opportunities (such as borrowing and lending in credit markets) in times when these opportunities are not available, helping agents to smooth consumption. There are many reasons why it may be harder for individuals to smooth consumption in recessions: household wealth is depleted, credit constraints are tighter, and lifetime income falls more after job loss (Davis and von Wachter [2011]). Public transfers (debt relief, unemployment insurance, or stimulus payments) targeted at individuals who lack other insurance mechanisms can be Pareto improving, irrespective of whether they increase aggregate output.

7.2 Persistent types

So far, I have assumed that borrowers and savers only differ in their income or preferences at date 0. Borrowers are initially more impatient or have lower income at date 0, but starting at date 1 they are identical to savers. Alternatively, borrowers might borrow because they expect higher future income than savers, or because they have fewer necessary expenditures to make in the future and have less need to save. One might worry that in this case, redistribution from the poor savers to the rich borrowers might no longer increase aggregate demand.

To address this concern, I augment my baseline model to allow a household's type to have a persistent effect on preferences and income at all dates, not just date 0. I continue to assume the planner only observes households' debt, borrowing and lending, and not their income or other consumption needs. Let agents have preferences $U(\boldsymbol{x}, \theta) = \sum_{t=0}^{\infty} \beta^t u_t(x_t, \theta)$, where x_t represents net financial plus public transfers: $x_t = \frac{d_{t+1}}{1+r_t} - d_t + T(d_t)$. This specification allows θ to capture differences in preferences or income in any period t. Borrowers may be lucky individuals who borrow against their long-run income, unlucky individuals facing temporary income losses, or simply more impatient. The main conclusions of the model go through as long as preferences satisfy the following assumptions. First, preferences are concave and satisfy an Inada condition:

Assumption 7.5. *For all t, θ, there exists $x_t(\theta)$ such that $u_t(\cdot, \theta)$ is C^2 on $(x_t(\theta), \infty)$, with $u_t' > 0$, $u_t'' < 0$, $\lim_{x \to x_t(\theta)} u_t'(x, \theta) = +\infty$, $\lim_{x \to x_t(\theta)} u_t(x, \theta) = -\infty$.*

Second, preferences satisfy a standard Spence-Mirlees condition: higher θ agents want to borrow more in period 0.

Assumption 7.6. $\frac{u_0'(x_0, \theta)}{\beta u_1'(x_1, \theta)}$ *is increasing in θ.*

Finally, I make the following assumption ensuring that the economy could reach steady state in period 1, were it not for the borrowing constraint.

Assumption 7.7. $\frac{u_{t+1}'(x, \theta)}{u_t'(x, \theta)} = 1$, *for all θ, x, $t \geq 1$.*

Under these three assumptions, the main results from Section 4 and 4.2 go through. As before, I characterize constrained efficient allocations as the solution to a modified Pareto problem.[32]

Proposition 7.8. *Suppose Assumptions 7.5, 7.6, 7.7 are satisfied. Then:*

1. *Constrained efficient allocations have the structure described in Proposition 4.1.*

2. *Every constrained efficient allocation can be implemented as an equilibrium with transfers.*

3. *There exists $\alpha(\theta_B)$ such that debt relief implements the optimal allocation if $\alpha < \alpha(\theta_B)$, and a savings subsidy implements the optimal allocation if $\alpha > \alpha(\theta_B)$. If (ICS) binds, $\alpha < \alpha(\theta_B)$; if (ICB) binds, $\alpha > \alpha(\theta_B)$.*

4. *Debt relief is Pareto-improving relative to the equilibrium without policy if and only if $\theta_B > \theta_{ZLB}$.*

Above, I raised the concern that if B's borrowing is motivated by higher future income (rather than low current income), transfers from S to B may not increase aggregate demand. This concern turns out to be unfounded. Even in this more general setting, household S is never liquidity constrained at date 1, whereas B may be constrained. Intuitively, type B households who borrow to consume more than their income at date 0, whatever their motive, must consume less than their income at date 1. Conversely, type S households who save at date 0 must consume more

[32]The full Pareto problem is presented in the Appendix. There is one relatively minor difference, which only applies when the borrowing constraint does not bind in the planner's problem, and when one incentive constraint binds. Since I focus on optimal policy in the regime where the borrowing constraint and ZLB both bind, this makes no difference to any of the main results. It is described in the Appendix.

than their income at date 1. If any household is constrained at date 1, it must be B, since B is consuming less than her income (even if she earns more than S).

In this simple model, the only difference in propensity to consume comes from binding liquidity constraints. Since B is sometimes constrained while S is never constrained, it follows that B must have a (weakly) higher propensity to consume than S; the zero lower bound restricts S's consumption, but not B's; and redistribution from S to B can stimulate aggregate demand. The same result would go through if type B households were permanently more impatient than type S households: clearly, this would only increase type B's propensity to consume.

7.3 Endogenous labor supply

So far, I have discussed the optimality of debt relief in an endowment economy. I now modify the model to include endogenous labor supply in the simplest possible way.

Households have concave preferences $U(c,h)$ over consumption and hours worked:

$$\theta_i U(c_0^i, h_0^i) + \sum_{t=1}^{\infty} \beta^t U(c_t^i, h_t^i)$$

Firms hire labor from households at a real wage w_t, and produce output using a linear technology, $y = h$. Each household receives an equal share of firms' real profit, $\pi_t = (1 - w_t)h_t$. In a Walrasian equilibrium, $w_t = 1$ and labor supply is efficient. As we have seen, this equilibrium may imply a negative real interest rate. In order to introduce the zero lower bound in this economy, I assume that firms are demand constrained in the market for final goods, borrowing from the literature on non-Walrasian equilibria (Benassy [1993]). Let a firm's desired sales be $y_t^* = \arg\max_{y \geq 0}(1 - w_t)y$. We have $y_t^* = 0$ for $w_t > 1$, $y_t = \infty$ if $w_t < 1$, and $y_t = [0, \infty)$ if $w_t = 1$. Output is less than or equal to desired sales: $y_t \leq y^*$. As before, I assume that interest rates clear the goods market whenever this does not violate the ZLB: $y_t = y_t^*$ if $r_t > 0$. When hiring labor, firms take into account the quantity constraints they face on the goods market as well as the real wage, so their demand for labor is $h_t = y_t$ (not $h_t = y_t^*$).[33] Households supply labor freely at the market-clearing real wage. In this economy, recessions occur when a fall in demand makes the real interest rate fall to zero. Firms become rationed in the goods market, and the real wage falls so that demand equals supply in the labor market.

I now consider optimal policy in this economy. I modify the definition of equilibrium with transfers to allow for a linear tax on labor income at date 1, which may depend on debt. Household i's budget constraint at date 1 is

$$c_1^i + d_1^i = T(d_1^i) + (1 - \tau(d_1^i))w_1 h_1^i + \pi_1 + \frac{d_2^i}{1 + r_1}$$

where $\tau(d_1^i)$ is the debt-contingent tax on labor income, and π_1 is the representative firm's profit

[33] As emphasized by Benassy [1993], quantity constraints in the goods market can cause underemployment in the labor market, even if the labor market itself is flexible.

at date 1.

I modify the Pareto problem to allow for endogenous labor supply.[34] As before, I define a constrained efficient allocation as the solution to the Pareto problem: constrained efficient allocations solve this Pareto problem.

Proposition 7.9. *Any solution to the Pareto problem with endogenous labor supply can be implemented as an equilibrium with transfers.*

The following proposition characterizes constrained efficient allocations.

Proposition 7.10. *Suppose the borrowing constraint binds in the Pareto problem. Then in any optimal allocation:*

1. *For any θ_B, there exist $1 > \alpha_B > \alpha_S > 0$ such that ICB binds iff $\alpha > \alpha_B$ and ICS binds iff $\alpha < \alpha_S$.*
2. *If the ZLB binds, S faces a positive labor wedge unless utility is quasilinear ($U_{ch}/U_{cc} = U_h/U_c$)*
3. *B faces a zero labor wedge unless ICS binds.*
4. *If preferences are separable ($U_{ch} = 0$), B always faces a zero labor wedge.*

Finally, as in the endowment economy, debt relief is always optimal, and is Pareto improving when the ZLB binds in equilibrium.

Proposition 7.11.
1. *There exists $\alpha(\theta_B) \in (\alpha_S, \alpha_B)$ such that the optimal allocation can be implemented with debt relief if $\alpha < \alpha(\theta_B)$, and the optimal allocation can be implemented with a savings subsidy if $\alpha > \alpha(\theta_B)$.*
2. *If the ZLB binds in equilibrium, the competitive equilibrium is Pareto inefficient. Debt relief is Pareto improving.*

8 Conclusion

I present a model in which both debt relief and macroprudential policy have costs and benefits. Debt relief redistributes towards households with a high propensity to consume, stimulating the economy at the zero lower bound, but encourages overborrowing ex ante. Macroprudential policies prevent the overborrowing that leads to a recession, but can make borrowers worse off. Naive debt relief and macroprudential policies may not be Pareto improving, because the costs outweigh the benefits. However, it is possible to design sophisticated ex post or ex ante transfer policies which are Pareto improving, because the benefits outweigh the costs.

I conclude by comparing optimal debt relief and macroprudential policy to some practical policy proposals. I forbid the planner from relaxing the borrowing constraint. This rules out direct lending to households, deficit-financed lump sum transfers, (Bilbiie et al. [2013b]), and

[34] The full Pareto problem with endogenous labor supply is presented in the Appendix.

postponement of debt payments (a feature of most debt restructurings). More subtly, it rules out converting mortgages into shared appreciation mortgages (Caplin et al. [2008]), which give lenders a share of future house price appreciation. Such a policy compensates lenders for writing down principal, and prevents moral hazard, since applicants lose out if prices rise. It is ruled out in my model, since promising more payments from borrowers to lenders (if prices rise) violates the borrowing constraint. Relaxing the borrowing constraint would always be optimal, if possible. In order to evaluate credit policies, it is necessary to have a model in which the government can circumvent borrowing constraints, while the private sector cannot.[35]

In this model, households are identical except for their unobservable preferences and observable debt. In reality, households differ in many other observable characteristics; as in any optimal taxation problem, it will in general be optimal to target debt relief and macroprudential taxes based on all these observables. Governments should target transfers to high debt households only insofar as they are liquidity constrained, and have a higher MPC than low-debt households. For example, if MPCs depend on debt relative to income, debt relief (or macroprudential taxes) should be aimed at households with a high debt-income ratio, not high debt per se.[36]

Many debates concern households with an existing stock of debt. If lenders offer restructuring to households who miss payments, then even financially healthy households will have an incentive to miss payments (Mayer et al. [2014]). My model can be interpreted along these lines. Suppose all agents have some inital debt, and owe some payments at dates 0 and 1. Type S households are financially healthy, and can make payments at date 0. Type B households face a negative shock at date 0, and must delay payments until date 1, when they repay everything and are liquidity constrained. Transfers to B would stimulate aggregate demand, but would encourage S to mimic B by delaying payments. My model suggests that optimal policy can avoid this problem in two ways. First, transfers should not be too large. Second, borrowers should only qualify for restructuring if they make some minimum level of payments. Such a policy is similar to Hockett et al. [2012]'s proposal for contingent principal reduction.

My model abstracts from housing, secured lending, and default. In reality, most household debt is mortgage debt, and most debates about debt relief focus on the mortgage market, where the benefits and cost of debt relief are subtly different from those considered here. As well as stimulating the overall economy, targeted mortgage debt relief could support the housing market, reducing fire sales and foreclosure externalities.[37] However, as discussed above, mortgage relief targeted to delinquent borrowers might induce even financially 'healthy' homeowners to delay payments. While these benefits and costs differ from those studied here, my main result - cleverly designed debt relief can be welfare improving - still applies. The core intuition is that

[35] Further, any credit policy must be relatively protracted. If the government was to lend to households for one period and then stop lending, that would merely postpone the liquidity trap.

[36] If low income households, rather than indebted households, have the highest propensity to consume, transfers should be targeted based on income, not wealth. In that case, optimal policy would balance the macroeconomic benefits from redistribution against the incentive effects of higher transfers.

[37] While the question is controversial, recent research tends to confirm that foreclosure reduces the value of nearby homes, and this externality seems to come from physical effects (foreclosed homes are not maintained, making the neighborhood less attractive) rather than a direct effect on prices (Gerardi et al. [2012], Fisher et al. [2014]).

in an equilibrium without policy, 'healthy' homeowners do not miss payments, as 'precarious' households do - presumably because they face some cost of delinquency (stigma, deterioration of their credit score, or risk of foreclosure). Therefore there is room to give *some* transfers to precarious households, without inducing healthy households to mimic them.

The model can also be reinterpreted to apply to sovereign debt relief. Outright debt relief is sometimes proposed as a solution to the European debt crisis; just as frequently, it is met with the criticism that debt relief represents a pure redistribution from creditor countries to debtor countries, and encourages overborrowing, sowing the seeds of future crises. My results suggest that one can design sovereign debt relief policies so that the benefits outweigh these costs. However, the model lacks several important characteristics of sovereign debt, especially default. A full extension of the model to cover sovereign debt relief is left to future work.

References

S Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–84, August 1994.

Manuel Amador, Iván Werning, and George-Marios Angeletos. Commitment vs. flexibility. *Econometrica*, 74(2):365–396, 2006.

Sandeep Baliga and Eric Maskin. Mechanism design for the environment. In K. G. Mäler and J. R. Vincent, editors, *Handbook of Environmental Economics*, volume 1 of *Handbook of Environmental Economics*, chapter 7, pages 305–324. Elsevier, June 2003.

Jean-Pascal Benassy. Nonclearing markets: Microeconomic concepts and macroeconomic applications. *Journal of Economic Literature*, 31(2):732–61, 1993.

Javier Bianchi. Efficient bailouts? Working Paper 18587, National Bureau of Economic Research, December 2012.

Florin O. Bilbiie, Tommaso Monacelli, and Roberto Perotti. Public debt and redistribution with borrowing constraints. *The Economic Journal*, 123(566):F64–F98, 2013a.

Florin O. Bilbiie, Tommaso Monacelli, and Roberto Perotti. Tax cuts vs. government spending: Welfare and the zero lower bound. Working paper, November 2013b.

Ricardo J. Caballero and Emmanuel Farhi. The safety trap. Working Paper 19927, National Bureau of Economic Research, February 2014.

Andrew Caplin, Noel Cunningham, Mitchell Engler, and Frederick Pollock. *Facilitating shared appreciation mortgages to prevent housing crashes and affordability crises*. Brookings Institution, 2008.

James Cloyne and Paolo Surico. Household debt and the dynamic effects of income tax changes. CEPR Discussion Papers 9649, C.E.P.R. Discussion Papers, September 2013.

Steven J. Davis and Till von Wachter. Recessions and the costs of job loss. *Brookings Papers on Economic Activity*, 43(2 (Fall)):1–72, 2011.

Gauti B. Eggertsson and Paul Krugman. Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach. *The Quarterly Journal of Economics*, 127(3):1469–1513, 2012.

Gauti B. Eggertsson and Michael Woodford. The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity*, 34(1):139–235, 2003.

Emmanuel Farhi and Jean Tirole. Collective moral hazard, maturity mismatch, and systemic bailouts. *American Economic Review*, 102(1):60–93, 2012.

Emmanuel Farhi and Iván Werning. A theory of macroprudential policies in the presence of nominal rigidities. Working Paper 19313, National Bureau of Economic Research, August 2013.

Lynn M. Fisher, Lauren Lambie-Hanson, and Paul S. Willen. The role of proximity in foreclosure externalities: Evidence from condominiums. Working Paper 20080, National Bureau of Economic Research, April 2014.

Luca Fornaro. Debt deleveraging, debt relief and liquidity traps. Working paper, May 2013.

Kristopher Gerardi, Eric Rosenblatt, Paul S. Willen, and Vincent Yao. Foreclosure externalities: Some new evidence. Working Paper 18353, National Bureau of Economic Research, September 2012.

Eric Giambattista and Steven Pennings. When is the government transfer multiplier large? Working paper, February 2013.

Luca Guerrieri and Matteo Iacoviello. Collateral constraints and macroeconomic asymmetries. International Finance Discussion Papers 1082, Board of Governors of the Federal Reserve System (U.S.), 2013.

Veronica Guerrieri and Guido Lorenzoni. Credit crises, precautionary savings, and the liquidity trap. NBER Working Papers 17583, National Bureau of Economic Research, November 2011.

Robert C Hockett, Daniel Alpert, and Nouriel Roubini. The way forward: Moving from the post-bubble, post-bust economy to renewed growth and competitiveness. *New America Foundation (2011)*, pages 12–01, 2012.

Tullio Jappelli and Luigi Pistaferri. Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–36, 2014.

Greg Kaplan and Giovanni L. Violante. A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239, July 2014.

Nobuhiro Kiyotaki and John Moore. Credit cycles. *Journal of Political Economy*, 105(2):211–48, April 1997.

Narayana R. Kocherlakota. Impact of a land price fall when labor markets are incomplete. Working paper, April 2013.

Anton Korinek and Alp Simsek. Liquidity trap and excessive leverage. Working Paper 19970, National Bureau of Economic Research, March 2014.

Paul R. Krugman. It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, 29(2):137–206, 1998.

Robert Lucas. *Models of Business Cycles*. Oxford: Basil Blackwell, 1987.

David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, Inc., 1969.

Christopher Mayer, Edward Morrison, Tomasz Piskorski, and Arpit Gupta. Mortgage modification and strategic behavior: Evidence from a legal settlement with Countrywide. Kreisman Working Papers Series in Housing Law and Policy 4, 2014.

Alisdair McKay and Ricardo Reis. The role of automatic stabilizers in the U.S. business cycle. Working Paper 19000, National Bureau of Economic Research, April 2013.

Neil R. Mehrotra. Fiscal policy stabilization: Purchases or transfers? Working paper, July 2013.

Paul Milgrom and Ilya Segal. Envelope Theorems for Arbitrary Choice Sets. *Econometrica*, 70(2): 583–601, March 2002.

James A Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(114):175–208, April 1971.

Kanishka Misra and Paolo Surico. Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106, 2014.

Hyunseung Oh and Ricardo Reis. Targeted transfers and the fiscal response to the great recession. *Journal of Monetary Economics*, 59(S):S50–S64, 2012.

Thomas Philippon and Philipp Schnabl. Efficient recapitalization. *The Journal of Finance*, 68(1): 1–42, 2013.

Stephanie Schmitt-Grohé and Martin Uríbe. Pegs and pain. Working Paper 16847, National Bureau of Economic Research, March 2011.

Joseph E. Stiglitz. Self-selection and Pareto efficient taxation. *Journal of Public Economics*, 17(2): 213–240, March 1982.

Jean Tirole. Overcoming adverse selection: How public intervention can restore market functioning. *American Economic Review*, 102(1):29–59, 2012.

Iván Werning. Pareto efficient income taxation. Working paper, MIT, April 2007.

Iván Werning. Managing a liquidity trap: Monetary and fiscal policy. Working paper, MIT, April 2012.

Michael Woodford. Public Debt as Private Liquidity. *American Economic Review*, 80(2):382–88, May 1990.

Pierre Yared. Public Debt Under Limited Private Credit. *Journal of the European Economic Association*, 11(2):229–245, 04 2013.

William R Zame. Efficiency and the Role of Default When Security Markets Are Incomplete. *American Economic Review*, 83(5):1142–64, December 1993.

Jonathan Zinman. Household debt: Facts, puzzles, theories, and policies. Working Paper 20496, National Bureau of Economic Research, September 2014.