October 2010

Public versus Private Risk Sharing[1]

Dirk Krueger

University of Pennsylvania, CEPR and NBER
dkrueger@econ.upenn.edu

Fabrizio Perri

University of Minnesota, Minneapolis FED, NBER and CEPR
fperri@umn.edu

Abstract

Can public income insurance through progressive income taxation improve the allocation of risk in an economy where private risk sharing is incomplete? The answer depends crucially on the fundamental friction that limits private risk sharing in the first place. If risk sharing is limited because insurance markets are missing for model-exogenous reasons (as in [8]) publicly provided risk sharing improves on the allocation of risk. If instead private insurance markets exist but their use is limited by limited enforcement (as in [24] then the provision of public insurance interacts with equilibrium private insurance, as, by providing risk sharing, the government affects the value of exclusion from private insurance markets and thus the enforcement mechanism of these contracts. We characterize consumption allocations in an economy with limited enforcement and a continuum of agents facing plausible income risk and tax systems with various degrees of progressivity (public risk sharing). We provide conditions under which more publicly provided insurance actually reduces total insurance for agents (excess crowding-out), or under which more public insurance increases total insurance (partial crowding-out).

KEYWORDS: Incomplete Markets, Progressive Taxation, Insurance, Limited Enforcement

JEL CODES: D52, E62, H31

[1]Corresponding Author: Dirk Krueger [email: dkrueger@econ.upenn.edu]. We would like to thank Andrew Atkeson, Tim Besley, V.V. Chari, Hal Cole, Dean Corbae, Joanne Feeney, Fatih Guvenen, Ken Judd, Patrick Kehoe, Tim Kehoe, Narayana Kocherlakota, Ramon Marimon, Giuseppe Moscarini, Edward Prescott, Tom Sargent, eight anonymous referees and seminar participants at numerous institutions for helpful comments and suggestions, as well as the NSF (under grant SES 0820494) for financial support. All remaining errors are our own. A previous version of this paper has circulated under the title "Risk Sharing: Private Insurance Markets or Redistributive Taxation?

1 Introduction

Should the government provide public insurance against idiosyncratic income risk by implementing a progressive tax system in which households with higher income realizations pay higher average tax rates, thus making the after-tax labor income process less risky than the pre-tax labor income process? The answer that economic theory gives to this question depends on the assumptions about the structure of private insurance markets. If these markets are complete, in that agents can trade a complete set of perfectly enforceable insurance contracts, then complete risk sharing is achieved via private markets and progressive income taxes provide no additional insurance. If, on the other hand, private insurance markets do not implement full risk sharing redistributive taxes might generate welcome additional insurance. As [38], [43] and others have pointed out, this beneficial effect of progressive income taxes has to be traded off against the adverse effect on incentives to supply labor and to accumulate capital, leading to a nontrivial optimal taxation problem.[2]

In this paper we demonstrate that if one models the frictions that lead to incomplete risk sharing in the first place explicitly, then the public provision of insurance may substantially affect the way private insurance markets work.[3] Our main contribution is to show that if private risk sharing is limited because private insurance contracts can only be enforced through exclusion from participating in financial markets in the future, then the provision of public insurance may crowd out the provision of private insurance against idiosyncratic risk, potentially more than one for one. That is, by attempting to better insure households against idiosyncratic risk the government may achieve exactly the opposite, namely a worse allocation of private consumption risk.

Our exact modelling approach follows the work by [24], [25] and [26] and does not impose a priori restrictions on the set of private insurance contracts that can be traded. These contracts, however, can not be fully enforced.[4] The only enforcement mechanism is the threat of exclusion from future credit and insurance markets upon default. Tax liabilities, however, are not subject to this enforcement problem as we assume that the penalty for defaulting on tax payments can be made prohibitively large by the government. If agents default on their private debt, they are banned from future credit and insurance markets, but retain their private (labor) endowment which is still subject to income taxation. A change in the tax system changes the severity of punishment from default by altering the utility an agent can attain without access to insurance markets, and thus changes the extent of enforcement of private contracts. Since enforcement defines the extent through which private contracts are used, a change in the tax system changes the use of private contracts. The allocative and welfare consequences of a change in the tax system then depend on

[2] A second common justification for redistributive taxation is the social desire to attain a more equal income or wealth (and hence consumption and welfare) distribution. Although we believe that this justification is potentially important we will not address this point in this paper. See the seminal paper by [37] for an analysis of the trade-off between the equity and the labor supply incentive effect of redistributive taxation.

[3] In the context of international trade, [16], [17] and [18] discusses how the extent of insurance depends on the underlying friction inhibiting perfect insurance.

[4] Another fraction of the literature derives market incompleteness from *informational frictions* underlying the phenomena of adverse selection and moral hazard (see [12] and their review of the literature). Optimal taxation is this class of models is the main focus of the recent *New Dynamic Public Finance* literature, see [27] and [22] for overviews. [33] is an early study that evaluates the importance of both frictions for economic growth.

the relative magnitudes of the change in public risk sharing implemented by the new tax system and the change in risk sharing through private insurance markets.

We evaluate this trade-off in a series of quantitative examples, motivated by U.S. data, and find that the crowding-out effect from the progressive income tax system characterized in this paper can be quantitatively important. In the examples households face income risk of a magnitude estimated from US household data, and are subject to a simple tax system with a constant marginal tax rate and a constant transfer. To quantify the impact of changes in the tax code on household labor income and consumption risk we construct and compute three measures of risk sharing for the income and consumption distributions: Private risk sharing, defined as the share of *after-tax* income risk *not* reflected in consumption risk (due to the operation of private financial markets); public risk sharing, defined as the reduction in income risk stemming from the progressivity of income taxes and total risk sharing which is (essentially) the sum of the two. When comparing steady state consumption allocations arising under different tax systems we find that making taxes more progressive by construction increases public risk sharing and typically reduces private risk sharing. In some cases the reduction of private risk sharing is bigger than the increase in public risk sharing so that increasing the progressivity of the tax code leads to *less* total risk sharing among households. We also show that this more than one-for-one crowding out result never appears in a standard incomplete markets model in the spirit of [8], Huggett [20] and [3] in which explicit risk sharing is limited for reasons exogenous to the model.

It is important to note that our quantitative analysis only focuses on the risk sharing effect of taxes and therefore abstracts from many elements that are important in the design of optimal taxes such as the presence of distortions of labor supply and savings decisions or a society's preference for redistribution. Thus our findings do not necessarily advocate a particular optimal tax schedule. They simply suggest that, when studying optimal taxation ignoring the effects that the tax system has on the functioning of private financial markets could be a first order omission.

A secondary methodological contribution of this paper is the characterization of the consumption allocation and distribution of a general equilibrium limited commitment model with a *continuum* of agents facing idiosyncratic income risk. This model allows us to analyze insurance mechanisms involving the entire population and not only pairwise relationships. We view this as crucial in our analysis of risk sharing arrangements such as progressive taxation since gains from insurance are particularly sizable among a large pool of agents with mostly idiosyncratic (i.e. largely uncorrelated) income risk. We demonstrate this point by contrasting the consumption allocation in our continuum economy with that arising in a model with only two agents (as studied by [26], [4], [25], among others) and show that the allocation of income risk in these two models is *qualitatively* different.[5] In addition our model endogenously delivers a rich cross-sectional consumption distribution and thus may be of independent interest for the study of other policy reforms where distributional concerns are important. But it is exactly the rich cross-sectional dimension of the model that leads to considerable theoretical and computational complications in solving it. To this end we adapt the

[5] In [28] and [29] we use US household data to evaluate the empirical predictions of the limited commitment model with a continuum of agents for household consumption dynamics and the cross-sectional distribution of consumption.

methodology of [5],[6] who study efficient allocations in an economy with a continuum of agents and private information to our environment with limited commitment. We also show (in the appendix), following [24], how to decentralize efficient allocations as equilibrium allocations in a standard Arrow Debreu equilibrium with individual rationality constraints.

The paper is organized as follows. In the next subsection we briefly relate our work to the existing literature. Section 2 lays out the model environment and defines a competitive equilibrium. In Section 3 we characterize efficient allocations. Section 4 presents qualitative features of these efficient allocations in the continuum economy and compare them to those arising in a simple economy with two agents and perfectly negatively correlated income shocks. Section 5 provides a quantitative thought experiment of changing the progressivity of the income tax code, both within our model and a standard incomplete markets model. Section 6 concludes; and additional definitions, proofs of the main propositions as well as extensions are contained in the appendix.[6]

1.1 Related Literature

The four papers that are most closely related to this work are [29], [42], [7] and [21]. In the first paper we argue that a limited commitment model of the form studied here can rationalize, qualitatively and quantitatively, the empirical fact that in the U.S. consumption inequality in the last 25 years has not increased to the same extent as income inequality. The mechanism in that paper is related to the one we present here: the part of rising income inequality that is due to larger income instability leads to autarky being a more severe punishment, relaxing the enforcement constraints and allowing better consumption insurance, and thus a smaller increase in consumption inequality. But the substantive focus is very different in the two papers: in our previous paper, [29], the emphasis is on capturing the effects on the consumption distribution of an observed change in income inequality and risk. In this work we study the relation between progressive taxation and the functioning of private financial markets in a more general fashion. In particular we study and characterize the relation between publicly provided risk sharing and private financial markets for economies with different structural characteristics and show that this relation change substantially across economies. From a methodological perspective the two papers are also quite distinct: in the current paper we characterize, theoretically and then quantitatively, the stationary solution to a social planner problem of a pure exchange economy. This solution can easily be decentralized as a competitive equilibrium, as we show in the appendix. In contrast, in [29] we solve directly for the competitive equilibrium transition path in a model with physical capital accumulation. [1] show that the equilibrium allocation in the model with capital is typically *not* efficient and thus does *not* coincide with that of the social planner problem. The advantage of analyzing the social planner problem in this paper is that it delivers a fairly sharp characterization of the efficient allocation for iid income shocks, and a simple algorithm for computing efficient (and thus

[6] A separate theoretical appendix contains details of some of the more involved technical arguments that are adaptations of the analysis by [6]. It is available at http://www.econ.upenn.edu/˜dkrueger/research/theoreticalapp.pdf. That appendix also presents the details of the computational algorithm used in the paper.

equilibrium) allocations for serially correlated shocks. On the other hand, the approach pursued here does not easily extend to models with capital accumulation (because with capital the welfare theorems do not apply) and does not lend itself naturally to an analysis of transition paths, something that is needed in [29].

Second, [42] study an optimal unemployment insurance problem with limited commitment and moral hazard. Households can trade private unemployment insurance contracts that are subject to the same underlying limited commitment friction that we model in our work. In addition, the government may provide public unemployment insurance, subject to a moral hazard friction. The public provision of insurance, as in our work, affects the outside option of private insurance contracts and thus determines the extent to which private insurance arises endogenously. In contrast to [42] who restrict attention to iid shocks that can only take two values (which is plausible for their unemployment application), our analysis and algorithm can deal with arbitrary finite state Markov chains. This allows us to more explicitly map our economy to U.S. income data.[7] On the other hand, our method requires the use of dynamic programming techniques whereas their analysis proceeds by fully characterizing the optimal private insurance contract using only the optimality conditions of the sequential problem.[8]

Third, [7] use a limited commitment model to study the effect of mandatory public insurance programs against aggregate risk on private insurance arrangements against idiosyncratic risk. Although their economy is populated by a large number of (potentially heterogeneous) agents, by assumption agents can only enter pairwise insurance arrangements, not involving any other member of the population. Therefore their underlying insurance problem is equivalent to the ones studied by [26] and [4]. Similar to our result they show that the extent to which idiosyncratic shocks can be insured depends negatively on the public provision of insurance against aggregate risk. A similar qualitative result is obtained by [21] in their study of a model with endogenous private insurance markets which are subject to private information (rather than limited enforcement) frictions.

Several other papers study the interaction between private and public insurance. In a model of informal family insurance [15] show that government provided unemployment insurance can crowd out informal insurance provided by the family more than one for one, a result similar to ours. On the empirical side, [13] and [10] measure the degree to which the public provision of health insurance through Medicaid crowds out the private provision of insurance and estimate it to be substantial. [31] and [32] set up a model with a finite, but potentially large number of agents that can engage in mutual insurance schemes. Once they solve for constrained-efficient insurance contracts numerically, however, they need to restrict attention to economies with either two agents (as in [31], in a model with capital accumulation), or they need to assume that agents engage in contracts with the rest of the population, treating the rest of the population as one agent (as in [32]). This again reduces the problem to a bilateral insurance problem as in the other papers discussed

[7]Our work also compares the extent of consumption insurance in the limited commitment model to that arising in a standard Bewley model. On the other hand, [42] can meaningfully study *optimal* policy by explicitly modeling the choice of public insurance subject to a moral hazard friction, whereas our work is silent about where the choice of the government for the tax system comes from.

[8]A similar characterization of the optimal insurance contract with limited commitment in the two-shock iid case can be found in [30]. [45] provides a partial characterization for serially correlated incomes.

previously.[9] [30] study a limited commitment model with a continuum of agents, but focus on endogenizing the value of default through competition. In their paper the interest rate is treated as exogenous, while it is endogenously determined, jointly with the consumption distribution in the current paper. As [30], the recent work by [9] contains an explicit characterization of the stationary consumption distribution under specific assumptions on preferences and the idiosyncratic income process, as well as a complete existence proof of a stationary equilibrium.

2 The Economy

There is a continuum of consumers of measure 1, who have preferences over consumption streams given by

$$U(\{c_t\}_{t=0}^{\infty}) = (1-\beta)E_0\left[\sum_{t=0}^{\infty}\beta^t u(c_t)\right] \tag{1}$$

The period utility function $u : \Re_+ \to D \subseteq \Re$ is assumed to be strictly increasing, strictly concave, twice differentiable and satisfies the Inada conditions. Its inverse is denoted by $C : D \to \Re_+$. Hence $C(u)$ is the amount of the consumption good necessary to yield period utility u. Let $\bar{D} = \sup(D)$; note that we do not assume u to be bounded so that $\bar{D} = \infty$ is possible.

An individual has stochastic endowment process $e \in E$, a finite set with cardinality N, that follows a Markov process with transition probabilities $\pi(e'|e)$. In what follows we use the words endowment and income interchangeably. For each consumer the transition probabilities are assumed to be the same. We assume a law of large numbers,[10] so that the fraction of agents facing shock e' tomorrow with shock e today in the population is equal to $\pi(e'|e)$. We assume that $\pi(e'|e)$ has unique invariant measure $\Pi(.)$. Without loss of generality we normalize average income $\bar{e} = \sum_e e\Pi(e) = 1$.

We denote by e_t the current period endowment and by $e^t = (e_0, .., e_t)$ the history of realizations of endowment shocks; also $\pi(e^t|e_0) = \pi(e_t|e_{t-1}) \cdots \pi(e_1|e_0)$. We use the notation $e^s|e^t$ to mean that e^s is a possible continuation of endowment shock history e^t. We also assume that at date 0 (and hence at every date), the cross-sectional measure over current endowment is given by $\Pi(.)$, so that the aggregate endowment is constant over time. At date 0 agents are distinguished by their initial lifetime utility entitlements[11] w_0 and by their initial shock e_0. Let Φ_0 be the joint measure of utility entitlements and shocks.

The government provides income insurance through a tax policy $\tau(e_t)$ that is constant over time. Since we want to focus on the public and private allocation of risk in this paper we focus on the case in which net

[9]The authors have to do this to avoid the curse of dimensionality. In their set-up of the problem the cumulative Lagrange multipliers on the enforcement constraints for each agent become continuous state variables, in practice ruling out computing allocations for economies with more than a small number of agents. The method of formulating this class of models recursively using cumulative Lagrange multipliers was pioneered by [34].

[10]Note that we do not require pairwise independence of endowment processes across individuals; the assumption of a law of large numbers can then be justified with [19], proposition 2.

[11]As we argue in the appendix, there is a one to one mapping between initial promised utility w_0 in the social planner problem and the associated initial wealth a_0 in the competitive equilibrium required to attain that level of lifetime utility.

revenues generated from the tax system are equal to zero.[12] We take the tax policy $\tau(.)$ as exogenously given, but vary its implied progressivity in our quantitative exercises. For an individual we let $y_t = e_t(1-\tau(e_t))$ denote the after-tax income. Since the function $\tau(.)$ does not depend on time, for a given tax function $\tau(.)$ there is a one-to-one mapping between pre-tax and after-tax endowments. From now on we let $y \in Y \subseteq \Re_{++}$ denote an individual's generic after-tax endowment, following the Markov process π with invariant distribution Π and denote by $y^t = (y_0, \ldots y_t)$ a history of after-tax endowment shocks. Taxes $\tau(.)$ satisfy a period-by-period budget constraint

$$\sum_{e_t} e_t \tau(e_t)\Pi(e_t) = 0 \tag{2}$$

With this assumption resource feasibility for this economy states that the sum of all agents' consumption has to be less or equal than the sum over all individuals' after-tax endowment, which equals 1 in every period.[13] Therefore, once $\tau(.)$ is fixed and hence the after-tax endowment process is specified, we can carry out the subsequent analysis of consumption risk sharing allocations without explicit consideration of government tax policy.

A consumption allocation $c = \{c_t(w_0, y^t)\}$ specifies how much an agent of type (w_0, y_0) consumes who experienced a history of endowment shocks y^t. Individuals, at any point in time, have the option to renege on existing allocations. The punishment for doing so is that agents that default on their allocations are banned from participating in future insurance arrangements and have to consume their *after-tax* endowments.[14] The expected continuation utility for an agent who defaults after history y^t is given by $U^{Aut}(y_t)$ where U is the solution to the functional equation

$$U(y) = (1-\beta)u(y) + \beta \sum_{y'} \pi(y'|y)U(y') \tag{3}$$

Note that $U^{Aut}(y_t)$ is strictly increasing in y_t, as long as the income shocks are uncorrelated or positively correlated over time.

Individuals have no incentive to default on a consumption allocation c, at any point in time and any

[12] Our theoretical analysis fully extends to the case of constant positive, but exogenous, government spending that needs to be financed through taxes.

[13] This immediately follows from $\bar{e} = 1$ and equation (2).

[14] Alternatively we could assume that households, after default, are allowed to save at an exogenous return r. The theoretical analysis of the dynamic programming problem below would not be affected by this extension. In section 5 we comment on how our quantitative results are affected by the choice of ruling out saving after default.

The general assumption of the consequences of bankruptcy is motivated by current US bankruptcy laws. Agents filing for bankruptcy under Chapter 7 must surrender all their assets above certain exemption levels, and the receipts from selling these assets are used to repay the consumer's debt. Remaining debt is discharged. In most cases of Chapter 7 bankruptcy debtors have no non-exempt assets (see [44]), so the consequences of filing for bankruptcy only entail restrictions on future credit. Individuals that declared personal bankruptcy are usually denied credit for seven years from major banks and credit card agencies. We view our assumption of being banned forever as an easily tractable approximation, keeping in mind that it may overstate the punishment from default somewhat.

For an explicit model of bankruptcy within the context of the standard incomplete markets model see [11]. The attractive feature of their model and the large literature it has spawned is that default indeed occurs with positive probability in equilibrium.

contingency, if and only if an allocation satisfies following enforcement constraints

$$U_t(w_0, y^t, c) \equiv (1-\beta)\left(u\left(c(w_0, y^t)\right) + \sum_{s>t}\sum_{y^s|y^t} \beta^{s-t}\pi(y^s|y^t)u\left(c(w_0, y^s)\right)\right) \geq U^{Aut}(y_t) \qquad \forall y^t \tag{4}$$

i.e. if the continuation utility from c is at least as big as the continuation utility from defaulting on c, for all histories y^t. Since there is no private information and markets are complete, exclusion will not happen in equilibrium as nobody would offer a contract to an individual for a contingency at which this individual would later default with certainty.

We now define a competitive equilibrium for the economy described above, and argue in the appendix that the solution to the social planner problem in the next section can be decentralized as such a competitive equilibrium. We will follow the approach of [24]. Consider an agent with period zero endowment of y_0 and initial wealth of a_0. Wealth is measured as entitlement to the period 0 consumption good. Let Θ_0 be the joint distribution over (a_0, y_0) and denote by $p_t(y^t)$ the date zero price[15] of a contract that specifies delivery of one unit of the consumption good at period t to/from a person who has experienced endowment shock history y^t. For each contingency $c_t(a_0, y^t) - y_t$ is the net trade of individual (a_0, y_0) for that contingency. In period 0 there is no uncertainty, so normalize the price of the consumption good at period 0 to 1.

A household of type (a_0, y_0) chooses an allocation $\{c_t(a_0, y^t)\}$ to solve

$$\begin{aligned} & & \max U_0(a_0, y_0, c) \qquad (5)\\ \text{s.t. } c_0(a_0, y_0) + \sum_{t=1}^{\infty}\sum_{y^t|y_0} p_t(y^t)c_t(a_0, y^t) & \leq & a_0 + y_0 + \sum_{t=1}^{\infty}\sum_{y^t|y_0} p_t(y^t)y_t \qquad (6)\\ U_t(a_0, y^t, c) & \geq & U^{Aut}(y_t) \qquad (7) \end{aligned}$$

Note that, as in [24], the continuing participation constraints enter the individual consumption sets directly.

Definition 1 *An equilibrium consists of prices $\{p_t(y^t)\}_{t=0}^{\infty}$ and allocations $\{c_t(a_0, y^t)\}_{t=0}^{\infty}$ such that*

1. *given prices, the allocation solves household's problem for almost all (a_0, y_0)*
2. *markets clear, i.e. for all t,*

$$\int \sum_{y^t} c_t(a_0, y^t)\pi(y^t|y_0)d\Theta_0 = \int \sum_{y^t} y_t\pi(y^t|y_0)d\Theta_0. \tag{8}$$

As is clear from the equilibrium definition our economy does not include physical capital accumulation or government debt, so assets are in zero net supply and the aggregate asset to income ratio is identically

[15] Note that in standard Arrow Debreu equilibrium theory with finitely many consumers, a complete description of the state of the economy would be *everybody's* endowment shock history, and all prices would be contingent on this complete state. With atomistic individuals, the assumed law of large number and no aggregate risk, attention can be restricted to equilibria in which prices (and quantities) depend only on idiosyncratic income histories y^t.

equal to zero. While this may seem unrealistic, we deliberately chose to abstract from both types of assets. In a closed economy with incomplete markets and precautionary savings motives an increase in income risk leads to higher precautionary saving, hence higher investment, a higher steady state capital stock and thus higher steady state production (see [3]). In our economy relaxed borrowing constraints drive the interest rate up and thus, in a version of the model with capital, the aggregate capital stock down. Since in this paper we want to focus on the risk sharing properties of different taxation schemes rather than the effects of taxation and income risk on capital accumulation, we compromise on realism to more clearly isolate the *potential* importance of the crowding-out mechanism introduced in this paper.

With respect to government debt, the government budget constraint would mandate that, for the same amount of outstanding government debt, the amount of taxes levied to finance the interest payments on the debt would vary across steady states, due to changes in the interest rate. Since the comparison of private households' welfare across economies with different tax burdens seems problematic, we also abstract from government debt in this paper.

3 Efficient Allocations

We now proceed directly to the study of efficient allocations. The presence of the infinite number of dynamic constraints (4) restricting the choice of consumption allocations does not easily allow for a direct characterization of the competitive equilibrium. Therefore in this section we follow [5] and [6] to first characterize efficient allocations and then argue in the appendix that they can be decentralized as competitive equilibrium allocations. As shown by [5] and [6] solving for efficient allocations does reduce to solving a standard dynamic programming problem. As they, however, we also have to restrict our analysis to stationary allocations, i.e. to allocations for which the cross-sectional consumption and lifetime utility distribution is constant over time.

The key insight of [5] is to analyze the problem of finding efficient allocations in terms of state contingent *utility* promises rather than state contingent consumption. An allocation is then a sequence $\{h_t(w_0, y^t)\}_{t=0}^{\infty}$ that maps initial utility entitlements w_0 and sequences of shocks y^t into levels of current utility in period t. Here $h_t(w_0, y^t)$ is the current period utility that an agent of type (w_0, y_0) receives if she experienced a history of endowment shocks y^t. Note that $c_t(w_0, y^t) = C(h_t(w_0, y^t))$ where C is the inverse of the period utility function defined in Section 2. For any allocation $h = \{h_t(w_0, y^t)\}_{t=0}^{\infty}$ define

$$U_t(w_0, y^t, h) = (1-\beta)\left(h_t(w_0, y^t) + \sum_{s>t}^{\infty} \sum_{y^s|y^t} \beta^{s-t}\pi(y^s|y^t) h_t(w_0, y^s)\right) \quad (9)$$

Equation (9) defines the continuation utility from an allocation h of agent of type (w_0, y_0) from date t and shock history y^t onwards.

In the appendix we formally define an efficient allocation. That definition says that a utility allocation is efficient if it attains the utility promises made by the initial distribution Φ_0 in a way that is individually rational, i.e. satisfies

$$\begin{aligned} w_0 &= U_0(w_0, y_0, h) && (10) \\ U_t(w_0, y^t, h) &\geq U^{Aut}(y_t) \qquad \forall y^t && (11) \end{aligned}$$

is resource feasible

$$\sum_{y^t} \int \left(C(h_t(w_0, y^t)) - y_t\right) \pi(y^t|y_0) d\Phi_0 \leq 0 \tag{12}$$

and there is no other allocation that satisfies the promise keeping constraint (10), the limited enforcement constraints (11) and requires fewer aggregate resources.[16]

In order to use recursive techniques, however, we have to restrict ourselves to stationary allocations. Define Φ_t to be the joint measure over endowment shocks y_t and continuation utilities $U_t(w_0, y^t, h)$ for a given allocation. An allocation is stationary if $\Phi_t = \Phi_0 = \Phi$. In the next subsections we will characterize it and demonstrate how to compute it.

3.1 Recursive Formulation

In order to solve for stationary efficient allocations we consider the problem of a social planner (whom one can also interpret as a financial intermediary) that is responsible of allocating resources to a given individual and who can trade resources at a fixed intertemporal price $\frac{1}{R}$. In this subsection we set up such a planners' recursive problem and in the next subsection its solution. We then show that the planners' policy functions induce a Markov process over utility promises and income shocks which has a unique invariant distribution, and finally we argue how to compute an R^* at which the resources needed to deliver utility promises dictated by the stationary distribution equal the aggregate endowment in the economy.

For constant $R \in (1, \frac{1}{\beta}]$, consider the following functional equation. Individual state variables are the promise to expected discounted utility that an agent enters the period with, w, and the current income shock y. The planner chooses how much current period utility to give to the individual, h, and how much to promise her for the future, $g_{y'}$, conditional on her next periods endowment realization y'. We now make the following assumptions on the individual endowment process[17]

Assumption 2 $\pi(y'|y) = \pi(y')$ *for every* $y', y \in Y$

Assumption 3 $\pi(y) > 0$, *for all* $y \in Y$

[16]Note that a Φ_0 that puts positive mass on (w_0, y_0) and satisfies $w_0 < U^{Aut}(y_0)$ does not permit a constrained feasible allocation, as promise keeping and enforcement constraints are mutually exclusive. We restrict attention to Φ_0 with the property that only initial utility entitlements at least as big as the utility from autarky have positive mass.

[17]For the quantitative analysis we will relax these assumptions, and there is no problem in setting up the recursive problem with serially correlated shocks. However, we were unable to prove that the state space is compact with persistent shocks.

The operator T_R defining the functional equation of the planner's problem is:

$$T_R V(w) = \min_{h,\{g_{y'}\}_{y'\in Y}\in D} \left\{ \left(1-\frac{1}{R}\right) C(h) + \frac{1}{R} \sum_{y'\in Y} \pi(y') V(g_{y'}) \right\} \tag{13}$$

$$\text{s.t. } w = (1-\beta)h + \beta \sum_{y'\in Y} \pi(y') g_{y'} \tag{14}$$

$$g_{y'} \geq U^{Aut}(y') \qquad \forall y' \in Y \tag{15}$$

where $V(w)$ is the resource cost for the planner to provide an individual with expected utility w when the intertemporal shadow price of resources for the planner is $\frac{1}{R}$. The cost consists of the cost for utility delivered today, $(1-\frac{1}{R})C(h)$, and expected cost from tomorrow on, $\sum_{y'} \pi(y')V(g_{y'})$, discounted to today. [5] and [6] show that a stationary allocation $\{h_t(w_0, y^t)\}_{t=0}^{\infty}$ is efficient if it is induced by an optimal policy from the functional equation above with $R > 1$ and satisfies the resource constraint with equality.[18]

Equation (14) is the promise-keeping constraint: an individual that is entitled to w in fact receives utility w through the allocation rules $\{h(.), g_{y'}(.)\}_{y'\in Y}$. The enforcement constraints in equation (15) state that the social planner for each state tomorrow has to guarantee individuals an expected utility promise at least as high as obtained with the autarkic allocation. The utility in autarky is given as the solution to the functional equation in (3).

3.2 Existence and Characterization of Policy Functions for Fixed R

We first prove the existence of optimal allocation rules in the problem with the additional constraints $g_{y'} \leq \bar{w}$ in (15), where $\bar{w}$ is an upper bound on future utility promises. We then characterize the solution of this problem and show that the additional constraints are not binding so that the solution to the problem with additional constraints is also solution to the original problem.[19] The modified Bellman equation is defined on $C(W)$, that is, the space of continuous and bounded functions on W, where $W = \{w \in \Re | \underline{w} \leq w \leq \bar{w}\} \subseteq D$ is a compact subset of $\Re$ and $\underline{w} := \min_y U^{Aut}(y)$. This gives us a standard *bounded* dynamic programming problem. From now on we will denote by T_R the operator defined above, but *including* the additional constraints.

Note that with the additional constraints on future utility promises, (14) and (15) imply that for every w in W possible choices h for current utility satisfy

$$\underline{h}(w) := \frac{w - \beta\bar{w}}{(1-\beta)} \leq h \leq \frac{w - \beta \sum \pi(y')U^{Aut}(y')}{(1-\beta)} =: \bar{h}(w) \tag{16}$$

[18]A policy $(h, \{g_{y'}\})$ induces an allocation, for all (w_0, y_0), in the following way: $h_0(w_0, y_0) = h(w_0, y_0), w_1(w_0, y^1) = g_{y_1}(w_0, y_0)$ and recursively $w_t(w_0, y^t) = g_{y_t}(w_{t-1}(w_0, y^{t-1}), y_t)$ and $h_t(w_0, y^t) = h(w_t(w_0, y^t), y_t)$. Adaptations of their proofs to our environment are contained in a separate theoretical appendix, available at http://www.econ.upenn.edu/˜dkrueger/research/theoreticalapp.pdf

[19]Note that if we had assumed that u and hence C are bounded functions this complication is avoided as the upper bound on u serves as upper bound $\bar{w}$. The results to follow do *not* require boundedness of u.

Accordingly define $\underline{h} := \underline{h}(\underline{w})$ and $\bar{h} := \bar{h}(\bar{w})$. We will show below that we can choose $\bar{w} = \max_y U^{Aut}(y) + \varepsilon$, for $\varepsilon > 0$ arbitrarily small, without the constraints $g_{y'}(w) \leq \bar{w}$ binding at the optimal solution, for all $w \in W$. In order to assure that the constraint set of our dynamic programming problem is compact, for all $w \in W$ we need (since D need not be compact) the following

Assumption 4 $[\underline{h}, \bar{h}] \subseteq D$

Assumption 3 is an assumption purely on the fundamentals (u, π, Y) of the economy and hence is straightforward to check. In particular, note that $\bar{h}(\bar{w}) = u(y_{\max}) \in D$ and $\underline{h}(\underline{w}) = u(y_{min}) - \beta[u(y_{\max}) - Eu(y)] \in D$ as long as $\frac{y_{\max}}{y_{\min}}$ is sufficiently small and/or β is sufficiently small.[20]

Using the standard theory of dynamic programming with bounded returns it is easy to show that the operator T_R has a unique fixed point $V_R \in C(W)$ and that for all $v_0 \in C(W)$, $||T_R^n v_0 - V_R|| \leq \frac{1}{R^n}||v_0 - V_R||$, with the norm being the sup-norm. Also V_R is strictly increasing, strictly convex and continuously differentiable and the optimal policies $h(w), g_{y'}(w)$ are continuous, single-valued functions.[21]

We will now use the first order conditions to characterize optimal policies.

$$\begin{aligned} C'(h) &\leq \frac{1-\beta}{\beta(R-1)} V'(g_{y'}) \\ &= \frac{1-\beta}{\beta(R-1)} V'(g_{y'}) \qquad \text{if} \qquad g_{y'} > U^{Aut}(y') \\ w &= (1-\beta)h + \beta \sum_{y' \in Y} \pi(y') g_{y'} \end{aligned} \tag{17}$$

The envelope condition is:

$$V'(w) = \frac{(R-1)}{R(1-\beta)} C'(h) \tag{18}$$

First we characterize the behavior of h and $g_{y'}$ with respect to w. The planner reacts to a higher utility promise w today by increasing current and expected future utility, i.e. by smoothing the cost over time and across states. The enforcement constraints, though, prevent complete cost smoothing across different states: some agents have to be promised more than otherwise optimal in certain states to be prevented from defaulting in that state. This is exactly the reason why complete risk sharing may violate the enforcement constraints.

Lemma 5 *Let assumptions 1-3 be satisfied. The optimal policy h, associated with the minimization problem in (13) is strictly increasing in w. The optimal policies $g_{y'}$, are constant in w and equal to $U^{Aut}(y')$ or strictly increasing in w, for all $y' \in Y$. Furthermore*

$$\begin{aligned} g_{y'}(w) &> U^{Aut}(y') \textit{ and } g_{\bar{y}'}(w) > U^{Aut}(\bar{y}') \textit{ imply } g_{y'}(w) = g_{\bar{y}'}(w) \\ g_{y'}(w) &> U^{Aut}(y') \textit{ and } g_{\bar{y}'}(w) = U^{Aut}(\bar{y}') \textit{ imply } g_{y'}(w) \leq g_{\bar{y}'}(w) \textit{ and } y' < \bar{y}' \end{aligned}$$

[20] For CRRA utility with coefficient of relative risk aversion $\sigma \geq 1$ assumption 3 is always satisfied.

[21] The proofs of these results are adaptations of proofs by [6] and contained in the separate theoretical appendix to this paper.

Proof. See Appendix ■

The last part of the lemma states that future promises are equalized across states whenever the continuing participation constraints permit it. Promises are increased in those states in which the constraints bind.

Now we state a result that is central for the existence of an upper bound $\bar{w}$ of utility promises. For promises that are sufficiently high it is optimal to deliver most of it in terms of current period utility, and promise less for the future than the current promises. This puts an upper bound on optimal promises in the long run, the main result in this section, stated in Theorem 7

Lemma 6 *Let assumptions 1-3 be satisfied. For every $(w, y') \in W \times Y$, if $g_{y'}(w) > U^{Aut}(y')$, then $g_{y'}(w) < w$. Furthermore, for each y', there exists a unique $w_{y'}$ such that $g_{y'}(w_{y'}) = w_{y'} = U^{Aut}(y')$.*

Proof. See Appendix ■

Theorem 7 *Let assumptions 1-3 be satisfied. There exists a $\bar{w}$ such that $g_{y'}(w) < w$ for every $w \geq \bar{w}$ and every $y' \in Y$.*

Proof. See Appendix ■

Note that the preceding theorem implies that whenever $w \in [\underline{w}, \bar{w}] = W$, then for all $y' \in Y$, the constraint $g_{y'}(w) \leq \bar{w}$ is never binding; since the constraint set in the original dynamic programming problem without the additional constraints is convex, the policy functions characterized in this section are also the optimal policies for the original problem for all $w \in W$. For any $(w_0, y_0) \in W \times Y$ these policies then induce efficient sequential allocations.

The policy functions $g_{y'}$ together with the transition matrix π induce a Markov process on $W \times Y$. In the next subsection we will show that this Markov process has a unique invariant measure, the long-run cross sectional distribution of utility promises (and hence welfare) and income, for any given fixed $R \in (1, \frac{1}{\beta})$.

3.3 Existence and Uniqueness of an Invariant Probability Measure

Let $\mathcal{B}(W)$ and $\mathcal{P}(Y)$ the set of Borel sets of W and the power set of Y. The function $g_{y'}(w)$, together with the transition function π for the endowment process, defines a Markov transition function on income shock realizations and utility promises $Q : (W \times Y) \times (\mathcal{B}(W) \times \mathcal{P}(Y)) \rightarrow [0, 1]$ as follows:

$$Q(w, y, \mathcal{W}, \mathcal{Y}) = \sum_{y' \in \mathcal{Y}} \begin{cases} \pi(y') & \text{if } g_{y'}(w) \in \mathcal{W} \\ 0 & \text{else} \end{cases} \tag{19}$$

Given this transition function, we define the operator T^* on the space of probability measures $\Lambda((W \times Y), (\mathcal{B}(W) \times \mathcal{P}(Y))$ as

$$(T^*\lambda)(\mathcal{W}, \mathcal{Y}) = \int Q(w, y, \mathcal{W}, \mathcal{Y}) d\lambda = \sum_{y' \in \mathcal{Y}} \pi(y') \int_{\{w \in W | g_{y'}(w) \in \mathcal{W}\}} d\lambda \tag{20}$$

for all $(\mathcal{W}, \mathcal{Y}) \in \mathcal{B}(W) \times \mathcal{P}(Y)$. Note that T^* maps Λ into itself (see [40], Theorem 8.2). An invariant probability measure associated with Q is defined to be a fixed point of T^*. We now show that such a probability measure exists and is unique.

Theorem 8 *Let assumptions 1-3 be satisfied. Then there exists a unique invariant probability measure* Φ *associated with the transition function* Q *defined above. For all* $\Phi_0 \in \Lambda((W \times Y), (\mathcal{B}(W) \times \mathcal{P}(Y))$, $(T^*\Phi_0)^n$ *converges to* Φ *in total variation norm.*

Proof. See Appendix ■

Note that Lemma 6 and Theorem 7 above imply that any ergodic set of the Markov process associated with Q must lie within $[U^{Aut}(y_{\min}), U^{Aut}(y_{\max})] \times Y$ and that the support of the unique invariant probability measure is a subset of this set.

So far we proved that, for a fixed intertemporal price R, policy functions $(h, g_{y'})$, cost functions V and invariant probability measures Φ exist and are unique. From now on we will index $(h, g_{y'}), V$ and Φ by R to make clear that these functions and measures were derived for a fixed R. In the next section we will discuss how to find the intertemporal price R^* associated with an *efficient* stationary allocation. This requires the allocation to satisfy the aggregate resource constraint with equality, a constraint that we have not yet imposed. We will do so in the next subsection, in order to solve for R^*.

3.4 Determination of the "Market Clearing" R

In the previous section we showed that for a fixed $R \in (1, \frac{1}{\beta})$ there exists a unique stationary joint distribution Φ_R over (w, y). We now describe how to find the intertemporal price $1/R^*$ such that the stationary consumption distribution associated with this R^* exactly exhausts the aggregate endowment.

Define the "excess demand function" $d : \left(1, \frac{1}{\beta}\right] \to \Re$ as

$$d(R) = \int V_R(w) d\Phi_R - \int y d\Phi_R \tag{21}$$

Since by assumption $\bar{y} := \int y d\Phi_R$ does not vary with R, the behavior of d depends on how V_R and Φ_R vary with R. The behavior of Φ_R with respect to R in turn depends on the behavior of $g_{y'}^R$ with respect to R as $g_{y'}^R$ determines the Markov process to which Φ_R is the invariant probability measure.

Computationally it is straightforward to determine the R^* such that $d(R^*) = 0$ by simply solving the associated dynamic programming problem for a fixed R, determine the associated unique invariant distribution, compute $d(R)$ from equation (21) and increase R if $d(R) < 0$ and lower it if $d(R) > 0$.

For this algorithm to be justified it is useful to establish further theoretical properties of $d(R)$. In the appendix we argue, again straightforwardly adapting arguments by [6] to the current environment that $d : d : \left(1, \frac{1}{\beta}\right] \to \Re$ is continuous and increasing.

Furthermore we demonstrate that under assumptions 2-4 and

Assumption 9 $U^{Aut}(y_{\max}) > u(\bar{y})$

we have $d(R = 1/\beta) > 0$. That is, for the interest rate equal to its complete market level the efficient consumption allocation provides perfect consumption insurance at $c = \bar{y}$, which, under assumption 9 violates the enforcement constraint for the household with the high income realization. Thus, under the maintained assumption the excess demand function d is weakly increasing and positive at $R = 1/\beta$.

As R approaches 1 from above, two things can happen: either the efficient consumption allocation is autarkic (and evidently satisfies $d = 0$) or the efficient allocation features some, but not complete (under assumption 9) risk sharing. In the appendix we also show that under

Assumption 10 $\beta \frac{u'(y_{\min})}{u'(y_{\max})} \geq 1$

autarky is not an efficient allocation. Under assumptions 2-4 and 10 in our quantitative work it is always true that

$$\lim_{R \searrow 1} d(R) < 0$$

but we were not able to show this result theoretically. Furthermore, in our numerical work the market clearing interest rate R^* is always unique, but since we can only show theoretically that $d(.)$ is increasing, but not that it is strictly increasing, we could not establish uniqueness of a market clearing interest rate either.

4 Qualitative Features of the Efficient Allocation

In this section we illustrate some of the qualitative features of the efficient allocation characterized in the section above. To do so we consider a simple numerical example of our economy in which the after-tax income process is *iid* can take only two values, $0 \leq y_l < y_h \leq 2$, which are equally likely. Note that since the average after-tax endowments are normalized to 1 we have $y_l = 2 - y_h$.

In order to highlight the qualitative differences of efficient allocations in our model with a continuum of agents and in the model with a small number of agents (the case typically studied in the literature, see e.g. [4], [7] or [25] we also present results for a limited commitment model with two agents $i = 1, 2$, each of which has endowment $y^i \in \{y_l, y_h\}$. We assume that in the two agent economy incomes are perfectly negatively correlated, so that if agent 1 has income y_l, agent 2 has income y_h and vice versa.[22] Consequently, as in the continuum economy, average income in the economy is nonstochastic and equal to 1. In accordance with the continuum economy we also assume that the income process in the two-agent economy is iid over time, with equal probability of each agent being rich in every period.

[22] [24], [26] and [4] all analyze limited commitment models with a small number of (types of) agents.

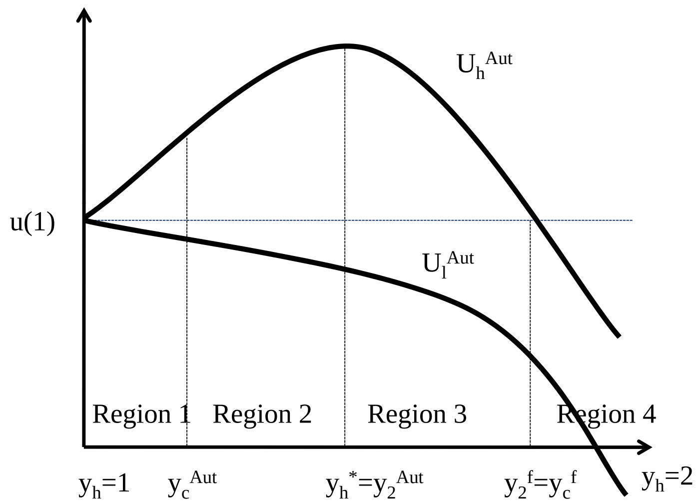

Figure 1: Value functions in Autarky

For both economies the utilities from autarky are given by

$$\begin{aligned} U^{Aut}(y_l) &= \left(1-\frac{1}{2}\beta\right)u(2-y_h)+\frac{1}{2}\beta u(y_h) \\ U^{Aut}(y_h) &= \frac{1}{2}\beta u(2-y_h)+\left(1-\frac{1}{2}\beta\right)u(y_h) \end{aligned}$$

We note that the size of y_h is a measure of income risk, with higher y_h associated with more income risk. Note that $U^{Aut}(y_l)$ is strictly decreasing in y_h, whereas $U^{Aut}(y_h)$ is strictly increasing in y_h at $y_h = 1$, and strictly concave with unique maximum $y_h^* \in (1, 2)$ satisfying

$$\frac{1}{2}\beta u'(2-y_h) = \left(1-\frac{1}{2}\beta\right)u'(y_h)$$

The values of autarky are plotted as a function of y_h in figure 1.

To put this example into the context of the tax system we will use in our quantitative examples, let the pre-tax endowment take two values $e_l < e_h$ with equal probability and recall we have assumed that mean income equals 1. The tax system is characterized by a constant marginal tax rate τ and a constant transfer d. Budget balance of the government (see equation (2)) implies that $\tau = d$ and therefore after-tax incomes are given by

$$\begin{aligned} y_l &= (1-\tau)e_l + d = e_l + \tau(1-e_l) \\ y_h &= (1-\tau)e_h + d = e_h - \tau(e_h - 1) \end{aligned}$$

and average taxes (net of transfers) are given by

$$t(e) = \frac{\tau e - \tau}{e} = \tau\left(1 - \frac{1}{e}\right).$$

Thus as long as $\tau > 0$ the tax system is progressive ($t'(e) > 0$) and the progressivity of the tax system increases with τ (since$\frac{dt'(e)}{\delta\tau} > 0$) and y_h decreases with τ. Similarly a $\tau < 0$ stands for a regressive tax system. The set of admissible tax rates ranges from $\underline{\tau} = \frac{e_l}{e_l - 1} < 0$, resulting in $y_l = 0$, to $\tau = 1$, resulting in $y_l = y_h = 1$ and thus perfect income insurance.[23]

It follows that a more progressive tax system is equivalent to a reduction in y_h, holding the pre-tax endowment process constant. Thus all comparative statics results with respect to y_h to follow can be interpreted as a change in the progressivity of the tax code; with *more progressivity representing a lower value of* y_h. From figure 1 we see that, starting from a very progressive tax system (a y_h below y_h^*) making the tax system even more progressive (lowering y_h) reduces the value of autarky for high income households and thus relaxes the limited enforcement constraint, potentially enabling *more* private risk sharing. On the other hand, if taxes become more progressive starting from a low level (a decline in y_h starting from a high level $y_h > y_h^*$) then the enforcement constraint tightens, private insurance is crowded out, potentially more than one for one. How strongly private insurance is affected by such a change in public insurance depends crucially on the *level* of private insurance, the determinants of which we discuss next.

4.1 Three Risk Sharing Regimes

In both economies efficient consumption allocations are either characterized by autarky (everybody consumes its after-tax income in all states), perfect risk sharing (everybody consumes average income of 1 in all states) or partial, but not perfect risk sharing. For both models, define critical income values for y_h (or equivalently, critical levels of tax progression τ)

$$\begin{aligned} 1 &\leq y_j^{Aut} \leq y_j^f \leq 2 \\ 1 &\leq y_j^{Aut} \leq y_j^f \leq 2 \end{aligned}$$

where $j = 2$ stands for the two agent economic and $j = c$ for the continuum economy. If $y_h \in [1, y_j^{Aut}]$, the constraint efficient consumption allocation is autarkic, if $y_h \in [y_j^f, 2]$, it is characterized by perfect risk

[23] Although regressive taxes $\tau < 0$ may seem empirically implausible and perhaps undesirable for a benevolent government that values equity among ex ante heterogeneous households, there is no conceptual reason to rule them out in the analysis of our model.

sharing, and if $y_h \in (y_j^{Aut}, y_j^f)$ it is characterized by partial risk sharing.

Figure 1 shows these regions. In region 1 the allocation is autarkic for both the two agent and the continuum economy, and in region 4 there is full risk sharing in *both* economies. Region 2, in which $y_h \in (y_c^{Aut}, y_2^{Aut})$, features autarky in the 2 household model but partial risk sharing in the continuum economy. Finally, in region 3 there is partial risk sharing in both economies.

4.1.1 Full Risk Sharing

Perfect risk sharing entails consuming average income $\bar{y} = 1$ for all agents, in each state. For this allocation to be constrained efficient it must satisfy the enforcement constraints. Since $U^{Aut}(y_h) \geq U^{Aut}(y_l)$ this requires

$$u(1) \geq \frac{\beta}{2} u(2 - y_h) + \left(1 - \frac{\beta}{2}\right) u(y_h) \tag{22}$$

in both economies $j = 2, c$.

The critical value $y_2^f = y_c^f > 1$ satisfies the above equations with equality. Perfect risk sharing occurs for exactly the same set of y_h values (and thus the same range of tax progressivity) in the continuum economy and the two-agent economy. As long as there is perfect risk sharing, a marginal change in tax progressivity y_h has no effect on the consumption allocation in either economy. In this range there is exactly a one-for-one crowding out of private risk sharing from public risk sharing in both economies.

4.1.2 Autarky

Autarky may be the only feasible allocation, and thus the (constrained-) efficient allocation. For the continuum economy autarky is efficient if and only if (see Lemma 16 in the appendix)

$$u'(y_h) \geq \beta u'(2 - y_h)$$

and for the two agent economy it is efficient if and only if (see [29], section 3)

$$u'(y_h) \geq \frac{\beta}{2 - \beta} u'(2 - y_h)$$

Thus $y_c^{Aut}, y_2^{Aut} \in (1, 2)$ and that $y_c^{Aut} < y_2^{Aut}$.Therefore the set of values of income (risk) y_h for which the constrained efficient allocation is autarkic is strictly bigger in the two agent economy than in the continuum economy. In this sense, there is more risk sharing possible in a continuum economy than in the two-agent economy.

In this region of the parameter space, a small change in y_h (equivalently, in τ) changes the consumption distribution one-for one with the income distribution. There is no crowding-out effect induced by a change in the tax system. Again, the absence of a crowding out effect occurs for a wider set of parameter values (tax rates) in the two agent economy, relative to the continuum economy.

Note that $y_2^{Aut} = y_2^*$, the level of y_h that maximizes the value of autarky for the currently rich household. Therefore in the region where the value of autarky increases in y_h and thus decreases with the progressivity of taxes the allocation in the two-agent economy is *always* autarkic. In other words, in the two agent economy (with only two possible income realizations that are iid and perfectly negatively correlated across agents) it is a theorem that public insurance can *never crowd in* private insurance. The same is not true in the model with a continuum of households where region 2 of figure 1 features both partial risk sharing and $U^{Aut}(y_h)$ increasing in y_h and thus declining in tax progressivity τ.

4.1.3 Partial Risk Sharing

For all $y_h \in (y_c^{Aut}, y_c^f)$ (respectively, for all $y_h \in (y_2^{Aut}, y_2^f)$ in the two agent economy) the stationary constrained efficient consumption distribution is characterized by partial risk sharing. In the next subsection we will characterize this distribution further in both economies, with particular focus on how it changes with the measure of inequality y_h,and thus with the degree of tax progressivity.

Two Agent Economy [25] show that the constrained efficient consumption allocation, conditional on $y_h \in (y_2^{Aut}, y_2^f)$ is fully characterized by the number c_h, the consumption level of households with currently high income.[24] This number is determined as the smallest solution of the equation

$$U^{Aut}(y_h) = \frac{1}{2}\beta u(2 - c_h) + \left(1 - \frac{1}{2}\beta\right) u(c_h)$$

and satisfies $c_h \in (1, y_h)$. Recall that $y_h^* = y_2^{Aut}$ is the unique maximum of $U^{Aut}(y_h)$ and thus within the range $y_h \in (y_2^{Aut}, y_2^f)$, an increase in y_h unambiguously reduces the value of $U^{Aut}(y_h)$. This relaxes the enforcement constraint of the income-rich household and thus, from the previous equation, reduces c_h and increases $c_l = 2 - c_h$. That is, consumption dispersion declines with an increase in y_h.Thus if there is partial insurance to start with, then an increase in public risk sharing through the tax system (a reduction of y_h) unambiguously increases consumption risk. In summary, in the two-agent economy with only two possible income realizations that are iid and perfectly negatively correlated across agents we obtain a theoretically sharp and unambiguous result linking the level of publicly provided insurance and the extent of crowding-out of private insurance: for small τ private markets provide full insurance which is crowded out by public insurance one for one. For intermediate values of τ there is partial private insurance and more than 100% crowding-out. Finally, for large τ private markets are inactive and there is no crowding-out of these private markets. Finally, crowding-in is theoretically impossible in this economy.

Continuum Economy For the continuum economy, under partial risk sharing (that is, for all $y_h \in (y_c^{Aut}, y_c^f)$), the consumption dynamics and distribution is more complex. Lemma 5 and 6 show that the

[24] Currently poor households consume $c_l = 2 - c_h$. The efficient consumption allocation in the two agent model is history-independent and only depends on the current state.

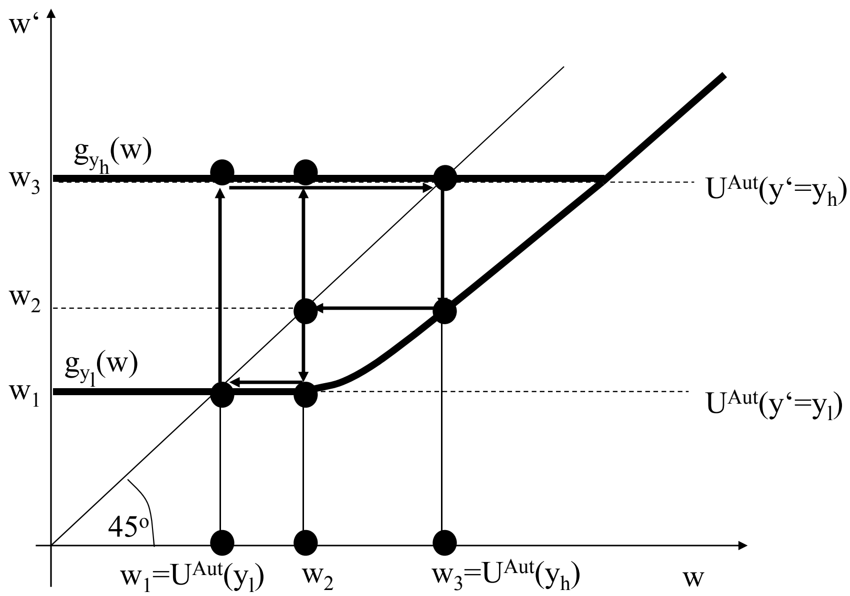

Figure 2: Policy Function for Promised Utility w

optimal policy function $g_{y'}(w)$, as a function of utility promises w, is constant and equal to the value of autarky $U^{Aut}(y')$, intersects the 45^0 line and at some point $w > U^{Aut}(y')$ starts to monotonically increase. If $g_{y_l}(w) > U^{Aut}(y_l)$ and $g_{y_h}(w) > U^{Aut}(y_h)$, then $g_{y_l}(w) = g_{y_h}(w)$. Figure 2 plots a typical policy function for utility promises tomorrow, $g_{y'}(w)$, against utility promises today, w, conditional on tomorrow's shock being either $y' = y_l$ or $y' = y_h$.

Figure 2 can be used to deduce the dynamics of utility promises w (and hence consumption, which is a strictly monotone function of w), as well as the invariant distribution over utility promises and hence consumption. First, the support of the stationary distribution of utility w is equal to $[U^{Aut}(y_l), U^{Aut}(y_h)]$, as shown in the theoretical analysis. For *all* $w \in [U^{Aut}(y_l), U^{Aut}(y_h)]$ an agent with high income $y' = y_h$ receives continuation utility $w' = U^{Aut}(y_h)$. History is forgotten in this event, as with $y' = y_h$ future utility does not depend on present utility entitlements w, which summarize the history of past endowment shocks. For agents with $y' = y_l$ history *does* matter. An agent starting with $w = w_3 = U^{Aut}(y_h)$ that receives $y' = y_l$ drops to $w_2 = g_{y'}(U^{Aut}(y_h)) < U^{Aut}(y_h)$, and, upon a further bad shocks, works herself downwards through the promised utility distribution. In a finite number of steps an agent with a string of bad shocks arrives at $w_1 = U^{Aut}(y_l)$, with any good shock putting her immediately back to $w_3 = U^{Aut}(y_h)$. Consumption obeys the same dynamics as utility entitlements since it is a strictly monotonic function of utility entitlements. The stationary utility entitlement (and thus consumption) distribution associated with the policy functions is depicted in figure 3.

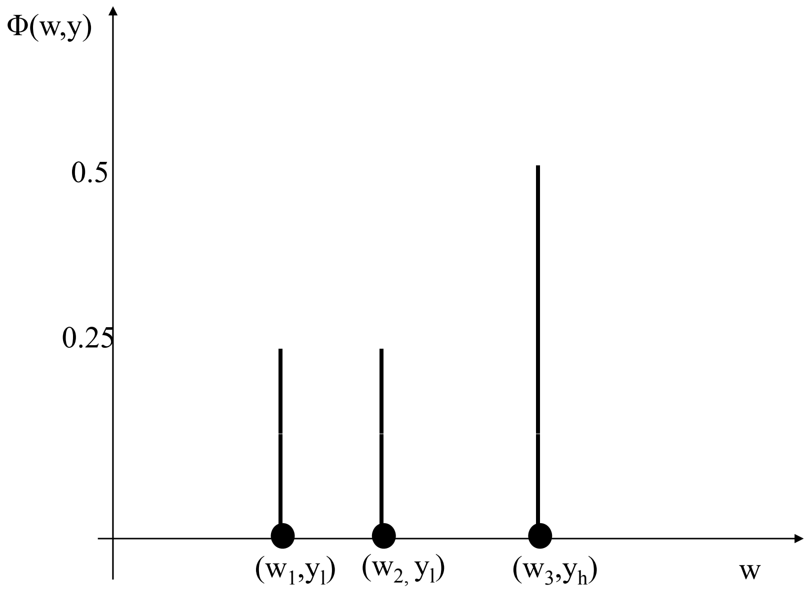

Figure 3: Invariant Consumption Distribution Φ

The efficient stationary consumption distribution is formally characterized as follows:[25]

Proposition 11 *For a given interest rate R the constrained efficient stationary consumption allocation is characterized by a number $n > 2$, and ordered consumption levels $c_1, c_2, \ldots c_n$, ordered lifetime utility levels $w_{1,} w_2, \ldots, w_n$ and associated probabilities $\pi_1, \pi_2, \ldots, \pi_n$ such that:*

1. *The stationary consumption and utility distribution is given by*

$$\pi_1 = 0.5^{n-1} \text{ and } \pi_j = 0.5^{n-j+1} \text{ for } j = 2, \ldots, n$$

2. *The consumption and utility levels satisfy*

$$\begin{aligned} w_1 &= U^{Aut}(y_l), \; w_n = U^{Aut}(y_h) \\ w_j &= (1-\beta)u(c_j) + 0.5\beta \left(w_{\max\{j-1,1\}} + w_n \right) \text{ for } j = 1, \ldots, n \end{aligned} \tag{23}$$

[25] [42], [30] and [9] prove this result, the latter two in a model with *exogenous* interest rates. [42] provide this characterization in an economy with endogenous interest rates, and without appealing to recursive techniques. Their approach to obtain an explicit characterization, to the best of our judgment, cannot fully be extended to cases with more than two income shocks (or to non-iid income processes).

As this paper [30] use recursive techniques. In contrast to the current paper the consequence of default in their model is not financial autarky, but the best insurance contract a competing financial intermediary offers, but this does not affect the characterization of the efficient consumption and lifetime utility allocation. Therefore the proof of the proposition in this paper is identical to the one in [30] and hence omitted.

and

$$\begin{aligned} u'(c_j) &= \beta R u'(c_{j-1}) \text{ for } j = 3, \ldots, n && (24) \\ u'(c_2) &\geq \beta R u'(c_1) && (25) \end{aligned}$$

The interest rate R itself is determined from the resource constraint

$$\sum_{j=1}^{n} \pi_j c_j = 1 \tag{26}$$

Even though the stationary consumption distribution is fairly sharply characterized by the previous proposition,[26] due to the endogenous interest rate R we found it impossible to establish sharp comparative statics with respect to the degree of tax progressivity τ. However, the following result immediately follows from the previous proposition.

Corollary 12 *In the continuum economy we have $c_1 = y_l$ and $c_n < y_h$.*

In summary, in the case of partial risk sharing the continuum economy insures households against bad income shocks by allowing consumption to decline slowly over time, relative to the two agent economy. This comes at the cost that consumption eventually falls to a lower level than in the two agent economy, in the event of a sequence of bad income shocks. Corollary 12 implies that, in stark contrast to the two-agent economy, an increase the progressivity of taxes (and thus an increase in y_l) raises the lower end of the of the support of the consumption distribution. Thus more publicly provided insurance, while leading to more consumption dispersion and lower minimum consumption in the two-agent model, leads to an *increase* of *minimum* consumption in the continuum model.

Furthermore, in the continuum economy, for a $y_h \in (y_c^{Aut}, y_2^{Aut})$ in region 2 of figure 1 the value of autarky of high income households is declining in the progressivity of the tax code (increasing in y_h) and the resulting allocation still displays partial private insurance. In this region of the tax policy space more public insurance can potentially *crowd in* private insurance, something that cannot possibly happen in the two household version of the model. These results demonstrate that the limited commitment model with only two (types of) households has qualitatively different implications for the impact of a change in the progressivity of taxation on the efficient consumption allocation than the continuum model with its richer consumption distribution. In the remainder of the paper we now quantitatively evaluate whether there is crowding-out or crowding-in of private insurance by a progressive tax system in a realistically parameterized version of the continuum economy, as well as to document its size.

[26] For a given n, equations (23), (24) and (26) form a system of $2n+1$ equations in the unknowns $c_1, c_2, \ldots c_n$, $w_1, w_2, \ldots, w_n$ and R. If for a given n, equation (25) is satisfied an there in no larger n such that this is true, we have found the optimal step size. This simple algorithm is only applicable in the *iid* case with two shocks, however. For a more general endowment process the computational method based on the recursive social planner problem discussed above needs to be used to compute stationary constrained-efficient allocations.

5 Quantitative Evaluation

In this section it is our goal to study the quantitative impact of changes in the progressivity of the tax system on the amount of risk sharing in equilibrium. In particular we use the model to measure the extent to which public risk sharing mechanisms (i.e. progressive taxes) and private risk sharing mechanisms (i.e. financial markets) interact in insulating private consumption from random income fluctuations. In order to do so we specify and estimate a simple statistical process for pre-tax labor income risk on US household data, and then study how a change in the tax system affects the extent of consumption risk sharing in the steady state.

We would like to stress that we restrict attention to the *long-run* consequences of different tax codes on private financial markets and overall risk sharing, rather than characterizing the entire transition path induced by a tax reform. Therefore the analysis in this section will focus on the positive (as opposed to the welfare) implications of the model and will not use the model to study *optimal* policy, since optimal tax policy is likely to depend on a variety of factors omitted here, including the explicit consideration of transitional dynamics.[27]

5.1 Functional Forms and Parameterization

We now describe the estimation of the pre-tax labor income process, the class of tax functions we consider in our experiments and the parameterization of preferences and the consequences of default.

5.1.1 Labor Income Risk

We specify the process for log pre-tax labor income of household i as a simple $AR(1)$ process

$$\log e_{it} = \rho \log e_{it-1} + \varepsilon_{it} \tag{27}$$

This process is meant to capture idiosyncratic labor income shocks (risk) of US households, and is fully characterized by the two parameters ρ and σ_ε. In order to separately identify the two parameters in (27) we use micro data with a panel dimension provided the US Consumer Expenditure Survey (CEX). We select the set of all households in the CEX over the period 2000-2007 whose head is between the age of 25 and 60 and who have positive labor income for two consecutive periods.[28] Consistent with the model we measure income as real labor earnings before taxes from all members of the household. Since in the model all households have

[27] Using our methodology to study transitions is not immediate. An unexpected change in government policies alters the set of feasible distributions of lifetime expected discounted utilities this economy can attain with given aggregate resources (which remain unchanged). Thus, for a particular agent the promised utility w she entered the period with is not necessarily a valid description of her state after the change in fiscal policy (a probability zero event) anymore. Consequently a method that employs promised expected utility as a state variable cannot be employed to compute transitional dynamics induced by unexpected policy innovations. Any transition analysis in this economy has to tackle the (sequential) competitive equilibrium directly, as we do in [29].

[28] A significant fraction of households in the CEX sample report their labor income in the past year at two consecutive points in time, on average 10 month apart. We use CEX as opposed to PSID as CEX has a larger sample size (although the panel dimension is much smaller). We conjecture that similar estimates would be obtained from PSID data since [23] document that the CEX and the PSID income data align rather well along a number of cross-sectional dimensions.

the same size we divide real total labor income by the number of adult equivalents in the household. Then, in order to exclude from our data permanent differences across households and aggregate risk,[29] the income measures for each year are regressed on a set of individual controls (which include quarter and education dummies, a quartic in age and age-education interactions). The residuals from these regressions are the data equivalents of $\log e_{it}$ in the process specified in (27) . Since for each household we have exactly two observations we can estimate time varying parameters ρ_t and $\sigma_{\varepsilon t}$ using the following simple cross sectional moment conditions:

$$\begin{aligned}\rho_t &= \frac{cov(\log e_{it}, \log e_{it-1})}{var(\log e_{it-1})} \\ \sigma_{\varepsilon t}^2 &= var(\log e_{it}) - \rho_t var(\log e_{it-1}).\end{aligned}$$

Finally, we obtain estimates of ρ and σ_ε as the simple time averages of ρ_t and $\sigma_{\varepsilon t}$ from the first quarter of 2000 to the first quarter of 2007. This results in estimates of $\rho = 0.8014$ (with a standard error of 0.03) and $\sigma_\varepsilon^2 = 0.1849$ (with a standard error of 0.021). These estimates reveal that labor income risk is quite persistent, but also contains a sizeable transitory component (possibly due to measurement error). These two general findings are consistent with a number of studies (see for example [35], and the large literature that followed) that estimate statistical processes for household earnings or income.

In order to map the estimated process into our theory in which pre-tax labor income follows a finite state Markov chain we discretize the continuous AR(1) process into a finite state Markov chain with 5 states using the procedure in [41]. Finally we re-normalize the value of all income states (after translating these states from logs into levels) such that mean pre-tax labor income equals to 1.

5.1.2 Fiscal Policy

Since the purpose of our quantitative exercise is to document the potential quantitative importance of the crowding-out mechanism, rather than to argue that the crowding-out effect is larger than one in the US economy under the actual (and quite complex) tax system we restrict ourselves to the same simple one-parameter family of tax functions as in section 4. For this tax system the degree of public risk sharing can be varied in a transparent way. Therefore, as above we assume that the tax code is given by a constant marginal tax rate τ and a fixed deduction (or transfer) d.

Recall that, given our normalization of average pre-tax income in the economy to $\bar{e} = 1$, the government budget constraint implies $d = \tau$, and therefore after-tax income y is given by

$$y = (1 - \tau)e + \tau.$$

The policy parameter $\tau \in (\frac{-e_{\min}}{1-e_{\min}}, 1]$ here measures the constant marginal tax rate but also, given a balanced budget, the size of lump sum transfers to households. Since marginal taxes are proportional and transfers are

[29] These are not explicitly modeled in our theoretical analysis that focuses on *idiosyncratic risk*.

lump sum, the higher is τ the larger is the degree of redistribution from the lucky to the unlucky households, i.e. the extent of public risk sharing. Notice that as τ approaches $\frac{-e_{\min}}{1-e_{\min}}$ from above, the tax system actually magnifies income risk faced by households. At the other extreme when $\tau = 1$ the government tax and transfer system eliminates income risk faced by households altogether: after-tax income y is constant and equal to 1 regardless of a household's pre-tax income realization e.

5.1.3 Preferences

We assume that households have log-utility, $u(c) = \log(c)$ and document results for various combinations of time discount factors $\beta \in (0, 1)$. The essential trade-off determining the extent of private risk sharing in equilibrium involves a comparison between the value of staying in the risk sharing agreement, relative to the value of being excluded from financial markets. The impact on both of these values of varying risk aversion σ are qualitatively similar to the impact of varying β. A higher β as well as a higher risk aversion σ increases the value of having access to risk sharing arrangements, relative to autarky, and hence relaxes the debt constraints, resulting in increase in private risk sharing.

5.2 Three Measures of Risk Sharing

Before we present our numerical results we define three measures of risk sharing which we will use to quantify the change in after-tax income and consumption risk faced by households induced by changes in the tax code. We define *Total Intermediation* (TI) of risk as one minus the ratio between the standard deviations of consumption to the standard deviation of pre-tax income:

$$TI = 1 - \frac{std(c)}{std(e)}.$$

Note that when $std(c) = 0$, $TI = 1$, consumption does not vary at all across individuals and the economy exhibits complete risk sharing. If $std(c) = std(e)$, $TI = 0$ and consumption varies one for one with pre-tax endowments For $0 < TI < 1$ there is some, but not complete risk sharing, with higher TI indicating higher risk sharing.

We can decompose TI into two components reflecting risk intermediation enforced by the government (GI) via the tax system and the additional risk intermediation achieved by private insurance contracts, (PI). Similar to TI we define as

$$GI = 1 - \frac{std(y)}{std(e)} \qquad PI = 1 - \frac{std(c)}{std(y)} \tag{28}$$

Note that given our tax function it follows that $std(y) = (1-\tau)std(e)$ so that $GI = \tau$. Thus GI measures nothing else but the progressivity of the tax code.

To interpret PI note that if $std(c) = 0$, $PI = 1$ and private markets completely intermediate income risk. If, on the other hand $std(c) = std(y) \neq 0$, $PI = 0$ and private markets do not achieve any risk sharing

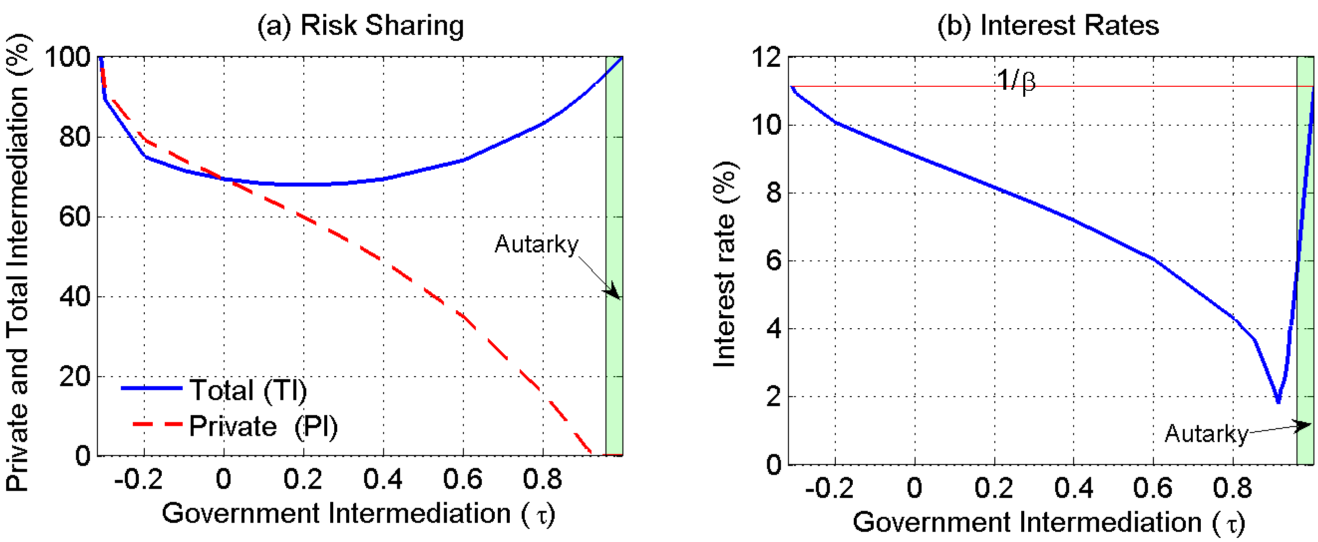

Figure 4: The effects of public risk sharing

beyond that implemented by the tax system. A simple calculation shows that

$$TI = \tau + (1 - \tau)PI \tag{29}$$

showing that total intermediation equals government intermediation plus private intermediation of that part of risk that is not already removed by the tax system.

5.3 Quantitative Results: Limited Commitment Model

Figure 4 plots the measures TI and PI (panel a) as well as the interest rate (panel b), as a function of τ, for $\beta = 0.9$. Panels (a) and (b) show that both total intermediation TI and the real interest rate are U-shaped functions of government intermediation τ. When τ is sufficiently close to $\frac{-e_{\min}}{1-e_{\min}}$ (which in our example is around -0.32) the tax system is regressive enough to drive the value of after-tax income in the lowest state close to 0. Therefore the value of autarky approaches $-\infty$, and the first best, full risk sharing allocation is enforceable. In this case PI is equal to 100%, and so is total intermediation (as is clear from equation 29); the gross real interest rate is equal to its complete markets value of $1/\beta$.

On the other end of the spectrum, if $\tau = 1$ full insurance is achieved through government intermediation alone: total intermediation is again 100% and the interest rate is $1/\beta$. Notice that when government intermediation is sufficiently high individual income risk is sufficiently low and the efficient allocation is the autarkic. The region of τ where autarky is efficient is given by the shaded area[30] in figure 4, and corresponds to region 1 in figure1.

[30]Technically, this region is defined by the condition of lemma 16 in the appendix.

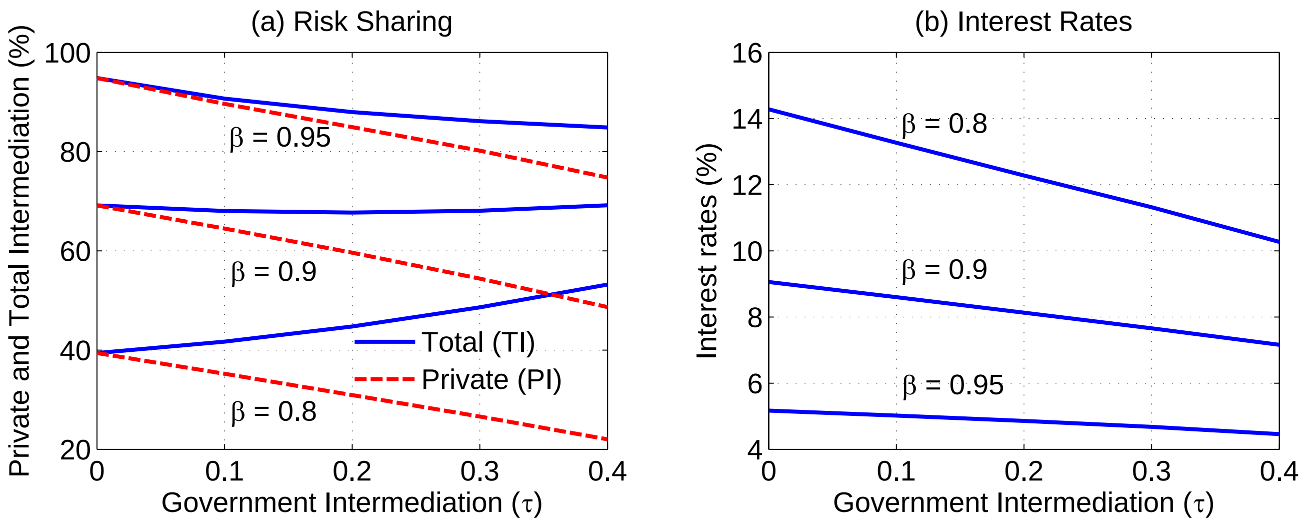

Figure 5: The effects of public risk sharing in three economies

In the middle range of government intermediation τ prefect risk sharing is not achievable and therefore total intermediation is less than 100%. The corresponding interest rate falls below its complete markets level of $1/\beta$. The fact that the first best allocation can be achieved with extremely regressive or extremely progressive taxes is a strong prediction of this model, but not one that we think is very relevant for the design of optimal policy, as obviously in real economies there are many factors which we abstract from in our setup (e.g. disincentive effects on labor supply, equity considerations) that will make such extreme policies undesirable. The more relevant conclusion from our model is that government tax policy, regardless of what motivates it, has a potential effect on the incentives that sustain of private risk sharing arrangements.

To evaluate the magnitude of this impact in an empirically plausible range of fiscal policy, figure 5 below reports private intermediation and total risk intermediation (panel a) and interest rates (panel b) for values of τ ranging from 0 (a flat tax) to 40% (which approximates the degree of public redistribution observed in some European countries) in three economies, characterized by $\beta = 0.95, 0.9, 0.8$. The values for βs are chosen to show three possible patterns of interaction between public and private risk sharing we discuss now.

In order to better understand the results in the figure it is useful to differentiate (29) with respect to τ yielding

$$\frac{\partial TI}{\partial \tau} = \underbrace{1 - PI}_{Direct} + \underbrace{(1-\tau)\frac{\partial PI}{\partial \tau}}_{Indirect} \tag{30}$$

The first term in the right hand side of (30), labeled "direct", represents the direct effect of public insurance on total intermediation. Note that this effect is decreasing in PI, as government intermediation τ partially displaces private intermediation, but is always non negative. The second term in the right hand side of (30) captures the indirect effect of τ. Changes in tax progressivity τ, by changing the relative attractiveness of autarky, alter the tightness of the enforcement constraints and hence the extent of private intermediation. In theory (as discussed in section 4) $\frac{\partial PI}{\partial \tau}$ could be either positive or negative: an increase in τ reduces dispersion in after tax income and whether $\frac{\partial PI}{\partial \tau}$ is positive or negative depends on whether the dispersion in the consumption distribution is reduced more or less than the dispersion in income distribution. Figure 5 shows that for our parameterizations $\frac{\partial PI}{\partial \tau} < 0$ for all three values of β and the range of $\tau \in [0, 40\%]$ we consider. This result is manifested in the negatively sloped dashed lines that plot private intermediation (PI) as a function of public intermediation τ. Notice that this effect is the distinguishing feature of our model as it captures the response of private financial markets to changes in publicly provided insurance. The figure also shows total intermediation (the solid lines). For $\beta = 0.95$ TI is decreasing in τ, for $\beta = 0.9$ is mildly U-shaped[31] and for $\beta = 0.8$ it is increasing in τ. In other words, for high β government intermediation crowds out private intermediation more than 1 to 1 while for low β this does not happen.

Why does the magnitude of the crowding-out effect crucially depend on the time discount factor? Different values for β simply capture the extent of private intermediation and thus the magnitude of the positive direct effect in equation (30). For high β private intermediation is high (i.e. private financial markets work well), the direct effect $(1 - PI)$ of increasing public intermediation is small and thus $\frac{\partial TI}{\partial \tau}$ is more likely to be negative. When β is low private intermediation is small (financial markets do not work as well) and hence the positive direct effect of increasing government intermediation is larger (in absolute value) than the negative indirect effect. Thus more government intermediation increases total intermediation.

More generally parameters configurations that reduce equilibrium private intermediation (for example low risk aversion, low variance or high persistence of income shocks, possibility of saving after default) increase $1-PI$, the direct effect in $\frac{\partial TI}{\partial \tau}$, and make the crowding-out effect smaller. Finally notice (panel b) that in the range of τ we display in figure 5 interest rates are always a declining function of public intermediation. This provides further direct evidence of the tightening effect induced by higher public risk sharing. Since lower income risk (due to higher τ) tightens borrowing constraints (by raising the value of autarky) it reduces the aggregate demand for credit, thus lowering the required equilibrium interest rate.

5.3.1 Crowding-out or Crowding-in?[32]

The discussion in section 4.1.3 suggests the possibility that in the economy with a continuum of agents public insurance might crowd-in private insurance i.e. $\frac{\partial PI}{\partial \tau} > 0$. The theory also suggests that crowding-in should happen when the individual income process exhibit low volatility so that reducing the volatility of income through taxes should reduce the value of autarky in the high states.[33] The results in figure 5 show that for the baseline parameters considered here $\frac{\partial PI}{\partial \tau}$ is negative for all values of τ in the range from 0% to 40%. Crowding-in never happens for these magnitudes of public insurance. This is because for these values of

[31] The U-shape is more transparent if TI is plotted for a larger range of τ , as we did in figure 4.
[32] We thank an anonymous referee for very useful comments leading to this section.
[33] In terms of figure 1 crowding-in is more likely to happen in region 2 where there is little private risk sharing to start with.

taxes τ individual after-tax income risk is high, and the public provision of more risk sharing increases the value of autarky even for high income states and hence tightens the binding enforcement constraints. In the context of figure 1 we are located in region 3.

How about that region of the parameter space in which the tightness of the enforcement constraints for high income states is declining in the progressivity of the income tax code (corresponding to region 2 of figure 1)? Roughly speaking, the set of tax rates τ for which this is the case is given by $\tau \in [80\%, 96\%]$. As figure 4 demonstrates in this region total intermediation increases with τ, but not private intermediation. In other words, the relaxation of the enforcement constraints permits the social planner to compress the consumption distribution (the standard deviation of consumption falls as τ increases). However, since the after-tax income distributions gets compressed at a faster rate, $\frac{std(c)}{std(y)}$ is increasing in τ and therefore $\frac{\partial PI}{\partial \tau} < 0$.

Qualitatively, alternative specifications of the model that raise the value of autarky (such as a larger persistence of the income process or the option of households to save at an exogenous storage return after default) and bring efficient risk sharing allocations closer to autarky for a larger set of taxes[34] tend to enlarge the region of τ for which the relevant enforcement constraints are relaxed as tax progressivity τ increases and for crowding-in might occur. In our many experiments with alternative parameter values, including different specification for the income process and different redistributive tax schemes, we did not find instances in which significant crowding-in arose for a robust set of taxes τ, although we cannot rule out the possibility of it occurring.[35] Obviously, in models in which public insurance can affect the functioning of credit markets through alternative channels[36] the crowding-in result might emerge more frequently and/or be more significant quantitatively.

5.4 Quantitative Results: Standard Incomplete Markets Model

In this section we contrast our findings on the effects of changes in public risk sharing (government intermediation) in a limited commitment economy to the effects of the same changes in a standard incomplete markets model. In this economy agents are only permitted to trade a single uncontingent bond and they face an exogenously specified constant borrowing limit $\underline{b}$.[37] By assumption enforcement frictions are absent in this model. The specific model we consider is most similar to the one studied by [20] and shares the same market structure and the same continuum of households with the models of [8] and [3]. The household problem in recursive formulation reads as

$$v(a, y) = \max_{-\underline{b} \leq a' \leq y + Ra} (1 - \beta)u(y + Ra - a') + \beta \sum_{y'} v(a', y')\pi(y'|y)$$

[34] In the context of figure 1, elements that enlarge region 2.

[35] For the instances in which we detected crowding-in, the range of τ for which it occurred was so narrow and the magnitude of the crowding-in so small that we could not comfortably dismiss the possibility that these findings were driven by numerical approximation errors associated with our algorithm and/or the coarseness of the 5-state income process we employ.

[36] See for example [2] who consider a model where taxes can have an effect on private financial markets through capital accumulation.

[37] Since average income is normalized to $\bar{y} = 1$, $\underline{b}$ has the interpretation of the fraction of average income that a household can borrow.

where a are holdings of the one-period bond at the beginning of the period and R is the gross real interest rate on these bonds. As with the previous model we compare stationary equilibria under different tax systems. To enable an exact comparison with the limited commitment economy we also use the same preferences and multiple discount factors, while we set the maximum amount that can be borrowed by households to an amount equivalent to 100% of average income.[38] Figure 6 reports how total and private intermediation (panel a) and interest rates (panel b) respond to changes in government intermediation in the standard incomplete markets economy.

First note that, similar to the previous model, as government intermediation increases private intermediation (the dashed lines in figure 6) falls, suggesting the presence of a crowding-out effect under this market structure as well. The intuition behind this crowding-out effect is quite different here, though. When larger government intermediation reduces income risk of households, it also reduces the incentive of consumers of engaging in precautionary saving. With a weaker precautionary motive households behave more like "Permanent Income" consumers, which leads to a more dispersed long-run wealth distribution as τ increases.[39] Such a more dispersed wealth distribution in turn is associated with a consumption distribution with larger variance and thus a lower extent of private intermediation. Notice for example that in figure 6, for the low value of $\beta = 0.8$ and significant public intermediation τ private intermediation PI turns negative: as the definition of PI in equation (28) makes clear for this constellation of parameters and government policy the dispersion of the consumption distribution is larger than the dispersion of the after-tax income distribution. This can only happen if the distribution of capital income displays a large variance, which in turn requires a large cross-sectional dispersion in asset holdings. Consistent with this argument in experiments with economies that feature much tighter borrowing constraints we found, not surprisingly, that the crowding-out effect is significantly smaller. With less generous borrowing constraints the long-run asset distribution is more narrowly bounded (at least from below) and therefore the corresponding consumption distribution is significantly less dispersed.

We conclude this section by highlighting two additional crucial differences between the responses to changes in government intermediation in the two models. First, although the crowding-out effect of private insurance from public insurance can be substantial even in the standard incomplete markets model we never found it to be larger than a 100% in any of the many quantitative examples we considered. Therefore in this model in which the structure of financial markets is unaffected by government policy (both the set of assets that are being traded as well as the borrowing constraints are policy-invariant) more public intermediation always leads to better overall consumption insurance (and consequently to higher ex-ante steady state welfare).[40]

[38] We obtain qualitatively similar findings for even tighter levels of the household borrowing constraint.

[39] In this model the desire to engage in precautionary saving is driven both by strictly convex marginal utility as well as potentially binding borrowing constraints.

[40] We want to stress that although we experimented with many possible parameters configurations and never have encountered the crowding-out effect to exceed 100% in the standard incomplete markets model we were not able to obtain a formal theoretical proof of this result. We therefore think it is conceivable (albeit not very likely, given our numerical results) that, even in the standard incomplete market model the long-run crowding out of private intermediation from public risk sharing could potentially

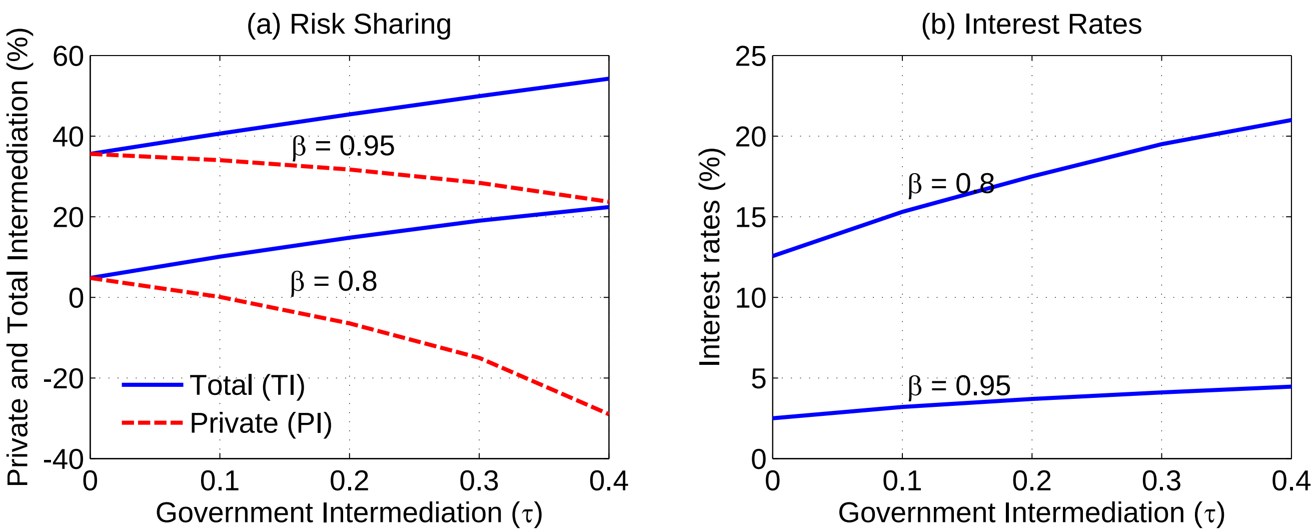

Figure 6: The effects of public risk sharing in standard incomplete markets

Finally the effect of government policy on real interest rates is qualitatively different in the limited commitment and the standard incomplete markets model. In the former more publicly provided risk sharing caused, in the relevant range of τ, a reduction in the equilibrium interest rate (see again figure 5, panel b), because larger τ lead to tightened enforcement constraints and thus reduced borrowing. In the standard incomplete markets economy in contrast the equilibrium interest rate is increasing in government intermediation (see figure 5, panel b). Higher government intermediation mitigates labor income risk and thus reduces the precautionary demand for assets (the supply of loans) which in turn drives up the equilibrium interest rate. This effect is largely absent in the limited commitment economy, due to the availability of a full set of state-contingent assets.[41]

6 Conclusion

In this paper we presented a model that highlights a mechanism through which the provision of public income insurance through progressive income taxation endogenously impacts the operation of private financial markets. By changing the incentives to default on private financial contracts government policy alters the extent to which private financial can provide consumption insurance against after-tax income risk. We demonstrate that when private labor income insurance markets are active, public risk sharing provided via taxes affects the extent of private risk sharing. In order to gain some insights into the potential quantitative magnitude of this effect we measured the extent of household labor income risk from US household data and

exceed 100%.

[41]Due to the presence of (state-contingent) borrowing constraints in the limited commitment model the precautionary motive to save is not entirely absent from this model either.

confronted consumers in our model with this risk. In our quantitative examples we found that the magnitude of the crowding-out effect can be substantial, potentially more than 100%. By attempting to provide better consumption insurance the government induces more consumption risk in equilibrium.

In contrast, if private insurance markets are assumed to be missing for model-exogenous reasons (and thus there is no interaction between the extent of public insurance and the structure of private markets), as in the standard incomplete markets model developed by [8], a tax reform that reduces the variance of after-tax income serves as am effective partial substitute for private insurance markets and always increases the amount of consumption risk sharing in the economy. This finding indicates that the assumption about the exact structure of private capital markets is crucial when analyzing social insurance policies.

On the methodological side we showed how to characterize and compute stationary efficient (and thus equilibrium) allocations in a limited commitment model with a continuum of households. We argued that the properties of the consumption risk sharing allocations in the continuum economy are qualitatively different than in an economy with a small number of agents, and so is the impact of public risk sharing on private consumption insurance, justifying the additional theoretical and computational burden required to deal with the continuum economy.

In order to isolate the effect of the tax system on private insurance markets and on risk sharing as clearly as possible we focused on comparisons of steady state equilibria and abstracted from several features of actual economies that are potentially important in the analysis of tax policy, most notably its potential distortions of labor-leisure and capital accumulation decisions as well as its redistributive consequences. A comprehensive quantitative positive and normative analysis of progressive taxation that incorporate the effects we highlight in this work into a model featuring these omitted distortions and equity concerns and considers transitional dynamics is called for, in our view. We defer such analysis to future research.

References

[1] Abraham, A. and E. Carceles-Poveda (2006): "Complete Markets, Enforcement Constraints and Intermediation," *Computing in Economics and Finance* 320, Society for Computational Economics.

[2] Abraham, A. and E. Carceles-Poveda (2010): "Endogenous Trading Constraints with Incomplete Asset Markets,", *Journal of Economic Theory* 145, 974-1004

[3] Aiyagari, R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659-684.

[4] Alvarez, F. and U. Jermann (2000): "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68, 775-798.

[5] Atkeson, A. and R. Lucas (1992): "On Efficient Distribution with Private Information," *Review of Economic Studies*, 59, 427-453

[6] Atkeson, A. and R. Lucas (1995): "Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance," *Journal of Economic Theory*, 66, 64-88.

[7] Attanasio, O. and V. Rios-Rull (2000): "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?," *European Economic Review,* 44, 1225-1258.

[8] Bewley, T. (1986), "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in Werner Hildenbrand and Andreu Mas-Colell (eds.) *Contributions to Mathematical Economics in Honor of Gerard Debreu*, North-Holland, 79-102

[9] Broer, T. (2009): "Heterogeneous Individuals in the International Economy," Ph.D. Thesis. European University Institute.

[10] Brown, J., N. Coe and A, Finkelstein (2007): "Medicaid Crowd-Out of Private Long Term Care Insurance Demand: Evidence from the Health and Retirement Survey," in James Poterba (ed.) *Tax Policy and the Economy*, The MIT Press, 1-34.

[11] Chatterjee, S., D. Corbae, M. Nakajima and V. Rios-Rull (2007): "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default," *Econometrica,* 75, 1525-1589.

[12] Cole, H. and N. Kocherlakota (2001): "Efficient Allocations With Hidden Income and Hidden Storage," *Review of Economic Studies,* 68, 523-542.

[13] Cutler, D. and J. Gruber (1996): "Does Public Insurance Crowd Out Private Insurance?," *Quarterly Journal of Economics*, 111, May 1996, 391-430.

[14] Díaz-Giménez, J., V.Quadrini and J.V. Ríos-Rull (1997): "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth", *Federal Reserve Bank of Minneapolis Quarterly Review*, 21, 3-21.

[15] Di Tella, R. and R. MacCulloch (2002): "Informal Family Insurance and the Design of the Welfare State," *Economic Journal*, 112, 481-503.

[16] Dixit, A. (1987): "Trade and Insurance with Moral Hazard," *Journal of International Economics*, 23, 201-220.

[17] Dixit, A. (1989a): "Trade and Insurance with Imperfectly Observed Outcomes," *Quarterly Journal of Economics*, 104, 195-203.

[18] Dixit, A. (1989b): "Trade and Insurance with Adverse Selection," *Review of Economic Studies*, 56, 235-247.

[19] Feldman, M. and C. Gilles (1985): "An Expository Note on Individual Risk without Aggregate Uncertainty," *Journal of Economic Theory,* 35, 26-32.

[20] Huggett, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953-969.

[21] Golosov, M. and A. Tsyvinski (2007): "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122, 487-534.

[22] Golosov, M., A. Tsyvinski and I. Werning (2007): "New Dynamic Public Finance: A User's Guide," *NBER Macroeconomics Annual* 2006, 317-364.

[23] Heathcote J., F. Perri and G. Violante (2009): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, forthcoming

[24] Kehoe, T. and D. Levine (1993): "Debt Constrained Asset Markets," *Review of Economic Studies*, 60, 865-888.

[25] Kehoe, T. and D. Levine (2001): "Incomplete Markets versus Debt Constrained Markets," *Econometrica,* 69, 575-598.

[26] Kocherlakota, N. (1996): "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies*, 63, 595-609.

[27] Kocherlakota, N. (2006): "Advances in Dynamic Optimal Taxation," Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Volume I.

[28] Krueger, D. and F. Perri (2005), "Understanding Consumption Smoothing: Evidence from the US Consumer Expenditure Survey," *Journal of the European Economic Association Papers and Proceedings*, 3, 340-349.

[29] Krueger, D. and F. Perri (2006), "Does Income Inequality lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73, 163-193.

[30] Krueger, D. and H. Uhlig (2006), "Competitive Risk Sharing Contracts with One-Sided Limited Commitment," *Journal of Monetary Economics*, 53, 1661-1691

[31] Ligon, E., J. Thomas and T. Worrall (2000), "Mutual Insurance, Individual Savings, and Limited Commitment," *Review of Economic Dynamics*, 3, 1-47.

[32] Ligon, E., J. Thomas and T. Worrall (2002), "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence in Village Economies," *Review of Economic Studies*, 69, 209-244.

[33] Marcet, A. and R. Marimon, (1992): "Communication, Commitment, and Growth," *Journal of Economic Theory*, 58, 219-249.

[34] Marcet, A. and R. Marimon, (2010): "Recursive Contracts," *Working paper*, European University Institute.

[35] MaCurdy, Thomas (1982): "The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics*, 18, 83-114.

[36] Miao, J. (2002): "Stationary Equilibria of Economies with a Continuum of Heterogeneous Consumers," *Working paper*, Boston University.

[37] Mirrless, J. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175-208.

[38] Mirrlees, J. (1974): "Notes on Welfare Economics, Information and Uncertainty, in *Essays on Economic Behavior under Uncertainty,* edited by M. Balch, D. McFadden and S. Wu, Amsterdam: North Holland.

[39] Nelson, J. (1994): "On Testing for Full Insurance Using Consumer Expenditure Survey Data," *Journal of Political Economy*, 102, 384-394.

[40] Stokey, N. and R. Lucas , with E. Prescott (1989): "Recursive Methods in Economic Dynamics," Harvard University Press, Cambridge.

[41] Tauchen, G. and R. Hussey (1991) "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59(2), 371-96

[42] Thomas, J. and T. Worrall (2007): "Unemployment Insurance under Moral Hazard and Limited Commitment: Public versus Private Provision," *Journal of Public Economic Theory*, 9, 151-181.

[43] Varian, H. (1980): "Redistributive Taxation as Social Insurance," *Journal of Public Economics*, 14, 49-68.

[44] White, M. (1998): "Why it Pays to File for Bankruptcy: A Critical Look at Incentives under U.S. Bankruptcy Laws and A Proposal for Change," *University of Chicago Law Review,* 65, 685-732.

[45] Zhang, Y. (2009): "A Characterization of Optimal Contracts with One-Sided Limited Commitment," *Working paper*, University of Iowa.

APPENDIX

A Sequential Definition of an Efficient Allocation

Definition 13 *An allocation* $\{h_t(w_0, y^t)\}_{t=0}^{\infty}$ *is constrained feasible with respect to a joint distribution over utility entitlements and initial endowments,* Φ_0*, if for almost all* $(w_0, y_0) \in supp(\Phi_0)$

$$\begin{aligned} w_0 &= U_0(w_0, y_0, h) && (31) \\ U_t(w_0, y^t, h) &\geq U^{Aut}(y_t) \qquad \forall y^t && (32) \\ \lim_{t\to\infty} \beta^t \sup_{y^t} U_t(w_0, y^t, h) &= 0 && (33) \\ \sum_{y^t} \int \left(C(h_t(w_0, y^t)) - y_t\right) \pi(y^t|y_0) d\Phi_0 &\leq 0. \quad \forall t && (34) \end{aligned}$$

An allocation $\{h_t(w_0, y^t)\}_{t=0}^{\infty}$ *is efficient with respect to* Φ_0 *if it is constrained feasible with respect to* Φ_0 *and there does not exist another allocation* $\{\hat{h}_t(w_0, y^t)\}_{t=0}^{\infty}$ *that is constrained-feasible with respect to* Φ_0 *and such that*

$$\sum_{y^t} \int C(\hat{h}_t(w_0, y^t))\pi(y^t|y_0) d\Phi_0 < \sum_{y^t} \int C(h_t(w_0, y^t))\pi(y^t|y_0) d\Phi_0 \textit{ for some } t \quad (35)$$

The separate theoretical appendix[42] shows that optimal policies from the recursive formulation of the social planner problem in the main text induce efficient allocations in the sense defined above.

B Determination of the "Market Clearing" R: Theoretical Results

B.1 Continuity of $d(R)$

Showing that d is continuous follows arguments that are identical to the ones in [6]. These are given in the separate extended theoretical appendix.

B.2 The Case $R = \frac{1}{\beta}$

In this subsection we characterize optimal policies of the planner for $R = \frac{1}{\beta}$ and provide a sufficient condition for the result $\lim_{R \nearrow \frac{1}{\beta}} d(R) > 0$ to hold. Note that for $R = \frac{1}{\beta}$

$$g_{y'}(w) = \begin{cases} w & \text{if } w \geq U^{Aut}(y') \\ U^{Aut}(y') & \text{if } w < U^{Aut}(y') \end{cases} \quad (36)$$

from the first order conditions of the recursive planners' problem (which still has a unique solution as all the results of Section 3.2 go through). Now there is a continuum of invariant measures associated with the Markov chain induced by the optimal policies. From (36) it is clear that any such measure $\Phi_{\frac{1}{\beta}}$ satisfies $w \notin \text{supp}\left(\Phi_{\frac{1}{\beta}}\right)$ for all $w < U^{Aut}(y_{\max})$ as the probability of leaving such a w is at least $\pi(y_{\max})$ and the probability of coming back (into a small enough neighborhood) is 0. Therefore all w in the support of any possible invariant measure satisfy $g_{y'}(w) = w$. From the promise-keeping constraint $h(w) = w$ follows, and each individuals' consumption is constant over time: for $R = \frac{1}{\beta}$ there is complete risk sharing in the long

[42]Available at http://www.econ.upenn.edu/˜dkrueger/research/theoreticalapp.pdf

run. For complete risk sharing to be efficient it has to satisfy the resource constraint. Since the cost function V_R is strictly increasing in w, the one of the continuum of invariant measures with lowest cost is

$$\Phi_{\frac{1}{\beta}}(w,y) = \begin{cases} \pi(y) & \text{if } w = U^{Aut}(y_{\max}) \\ 0 & \text{if } w \neq U^{Aut}(y_{\max}) \end{cases} \tag{37}$$

All individuals receive utility promises $w = U^{Aut}(y_{\max})$ and hence the same current utility $h(U^{Aut}(y_{\max})) = U^{Aut}(y_{\max})$. This allocation has per-period resource cost $C(U^{Aut}(y_{\max}))$ and is resource feasible if and only if $C(U^{Aut}(y_{\max})) \leq \bar{y}$, or applying the strictly increasing period utility function u to both sides, if and only if $U^{Aut}(y_{\max}) \leq u(\bar{y})$. Let the net resource cost of this allocation be denoted by

$$d\left(\frac{1}{\beta}\right) = C(U^{Aut}(y_{\max})) - \bar{y} \tag{38}$$

We summarize the discussion in the following

Lemma 14 *Let assumptions 2-4 be satisfied. For $R = \frac{1}{\beta}$ any solution to the recursive social planners' problem exhibits complete risk sharing. There exists an efficient stationary allocation with complete risk sharing if and only if $U^{Aut}(y_{\max}) \leq u(\bar{y})$.*

Intuitively, the lemma states that it is constrained efficient to share resources equally among the population in this economy if the agents with the highest incentive to renege on this sharing rule, namely the agents with currently high income, find it in their interest to accept constant consumption at $c = \bar{y}$ and lifetime utility $u(\bar{y})$, rather than to leave and obtain lifetime utility $U^{Aut}(y_{\max})$.Using arguments similar to showing continuity of $d(R)$ on $(1, \frac{1}{\beta})$ one can show that $\lim_{R \nearrow \frac{1}{\beta}} d(R) = d\left(\frac{1}{\beta}\right)$, where $d\left(\frac{1}{\beta}\right)$ is defined in (38). In order to rule out complete risk sharing[43] it is sufficient to make assumption 9 in the main text.

Lemma 15 *Let assumptions 2-9 be satisfied. Then $\lim_{R \nearrow \frac{1}{\beta}} d(R) > 0$.*

Proof. Applying the strictly increasing cost function C to the inequality of assumption 4 gives

$$d\left(\frac{1}{\beta}\right) = C\left(U^{Aut}(y_{\max})\right) - \bar{y} > 0$$

■

Note that assumption 9 is satisfied if the time discount factor β is sufficiently small, agents are not too risk-averse or the largest endowment shock is sufficiently large.

B.3 The Case of R Approaching 1

In this subsection we provide necessary and sufficient conditions for autarky (all agents consume their endowment in each period) to be an efficient allocation and discuss the characteristics of consumption allocations for R approaching 1 from above.

If agents are very impatient and/or the risk of future low endowments is low, then it is not efficient for the planner to persuade currently rich agents to give up resources today in exchange for insurance tomorrow. We now formally state and prove this result.

Lemma 16 *Let assumptions 2-4 be satisfied. Autarky is efficient if and only if*

$$\beta \frac{u'(y_{\min})}{u'(y_{\max})} < 1 \tag{39}$$

[43] If there is complete risk sharing under a particular tax system (remember that the tax system maps a given pre-tax income process into a particular after-tax income process), then a small tax reform has no effect on the extent of risk sharing since the resulting allocation is the complete risk sharing allocation: our crowding-out effect is absent.

Proof. For the if-part we note that the autarkic allocation satisfies the first order conditions for some $R > 1$ if (39) holds. Since autarky is constrained feasible, it is efficient.

For the only-if part we first show that there is an allocation attaining a distribution of utility that stochastically dominates the utility distribution in autarky and requires no more resources. It is then immediate that autarky is not efficient. In autarky the measure over utility entitlements and endowment shocks is given by

$$\Phi^{Aut}(\{U^{Aut}(y), y\}) = \pi(y) \tag{40}$$

We show that there exist allocations that attain the joint measure $\hat{\Phi}$ defined as

$$\begin{aligned} \hat{\Phi}(\{U^{Aut}(y), y\}) &= \pi(y) \qquad \text{all } y \neq y_{\min} \\ \hat{\Phi}(\{U^{Aut}(y_{\min}), y_{\min}\}) &= \pi(y_{\min})(1 - \pi(y_{\max})) \\ \hat{\Phi}(\{\tilde{w}, y_{\min}\}) &= \pi(y_{\min})\pi(y_{\max}) \end{aligned} \tag{41}$$

where $\tilde{w} = U^{Aut}(y_{\min}) + \varepsilon$ for small $\varepsilon > 0$. Define $\delta_{\max}$ and $\delta_{\min}$ implicitly by

$$\begin{aligned} \tilde{w} &= (1-\beta)(u(y_{\min}) + \delta_{\min}) + \beta \sum_y \pi(y) U^{Aut}(y) \\ U^{Aut}(y_{\max}) &= (1-\beta)(u(y_{\max}) - \delta_{\max}) + \beta \sum_{y \neq y_{\min}} \pi(y) U^{Aut}(y) + \beta \pi(y_{\min}) \tilde{w} \end{aligned} \tag{42}$$

Since $\tilde{w} = U^{Aut}(y_{\min}) + \varepsilon$, we have

$$\begin{aligned} \delta_{\max} &= \frac{\beta \pi(y_{\min})}{(1-\beta)} \varepsilon \\ \delta_{\min} &= \frac{\varepsilon}{(1-\beta)} \end{aligned} \tag{43}$$

The autarkic allocation exhausts all resources. The new allocation reduces consumption for the $\pi(y_{\max})$ agents with $y_{\max}$ by $\delta_{\max}$ and increases consumption for $\pi(y_{\max})\pi(y_{\min})$ agents by $\delta_{\min}$. Hence, compared to the autarkic allocation the change in resource requirements is given by

$$\begin{aligned} \Delta &= -\pi(y_{\max}) C'(u(y_{\max})) \delta_{\max} + \pi(y_{\max}) \pi(y_{\min}) C'(u(y_{\min})) \delta_{\min} \\ &= \frac{\pi(y_{\min})\pi(y_{\min})\varepsilon}{(1-\beta)} \left(\frac{-\beta}{u'(y_{\max})} + \frac{1}{u'(y_{\min})} \right) \end{aligned} \tag{44}$$

Therefore $\Delta \leq 0$ if and only if

$$\beta \frac{u'(y_{\min})}{u'(y_{\max})} \geq 1 \tag{45}$$

Under this condition the new allocation is resource feasible, incentive feasible and attains $\hat{\Phi}$, a distribution that dominates Φ^{Aut}. It is straightforward to construct the sequential allocation h induced by the recursive policies supporting $\hat{\Phi}$. By reducing $h_0(U^{Aut}(y_{\min}), y_{\min})$ so that the agents receiving discounted utility $\hat{w}$ under $\hat{\Phi}$ receive $U^{Aut}(y_{\min})$ the new allocation attains Φ^{Aut} but with less resources, a contradiction to the assumption that autarky is constrained efficient. ■

The previous lemma provides a condition under which $d(R) = 0$ as R approaches 1, with autarky as the (efficient) allocation. In order to assure that autarky is not efficient assumption 10 is thus sufficient (and also necessary).

B.4 Existence of Market Clearing R

The arguments in this section constitute elements of an existence proof of a market clearing interest rate R^*, which remains incomplete because (even under assumption 10) we cannot show[44] (but verify easily in numerical work) that

$$\lim_{R \searrow 1} d(R) < 0$$

As indicated above, some of our results and proof strategies resemble [6] closely. The basic strategy to compute and prove existence of a stationary general equilibrium also exhibits similarities to the approach taken in the standard incomplete markets literature, see [8], [20] and [3]. The main difference is that these authors, due to the simple asset structure in their models, can tackle the equilibrium directly. As we do, they first, for a fixed and constant interest rate, solve a simple dynamic programming problem[45] (in their case for the single household, with assets as state variable, in our case for the social planner, with utility promises as state variables). Then they let the optimal policies induce a Markov process to which a unique invariant distribution is shown to exist.[46] Finally the market clearing interest rate is determined from the goods or asset market clearing condition.[47]

C Decentralization

In this section we describe how to decentralize a stationary efficient allocation $h = \{h_t(w_0, y^t)\}_{t=0}^{\infty}$ induced by the optimal policies from the recursive planners' problem as a competitive equilibrium as defined in the main text. Let $\beta^t \pi(y^t|y_0)\mu(a_0, y^t) \geq 0$ be the Lagrange multiplier associated with the enforcement constraint at history y^t and $P(y^t) = \{y^\tau | \pi(y^t|y^\tau) > 0\}$ be the set of all endowment shock histories that can have y^t as its continuation. Using the first order necessary conditions of the household's maximization problem (5) one obtains

$$\beta \frac{u'\left(c_t(a_0, y^{t+1})\right)\pi(y^{t+1}|y_0)}{u'\left(c_t(a_0, y^t)\right)\pi(y^t|y_0)} = \frac{p_{t+1}(y^{t+1})}{p_t(y^t)} \frac{1 + \sum_{y^\tau \in P(y^t)} \mu(a_0, y^\tau)}{1 + \sum_{y^\tau \in P(y^{t+1})} \mu(a_0, y^\tau)} \tag{46}$$

Obviously, an agent whose participation constraint does not bind at contingency y^{t+1}, following history y^t, faces the standard complete markets Euler equation (as $\mu(a_0, y^{t+1}) = 0$).

Now consider the efficient allocation of utilities $\{h_t(w_0, y^t)\}$ determined in the main text. Combining the first order condition and the envelope condition from the planners problem we have for an agent that is unconstrained[48] (see (17) and (18)):

$$\frac{1}{R} = \beta \frac{C'(h_t(w_0, y^t))}{C'(h_{t+1}(w_0, y^{t+1}))} \equiv \beta \frac{u'(c_{t+1}(w_0, y^{t+1}))}{u'(c_t(w_0, y^t))} \tag{47}$$

[44] This result can be shown for CRRA utility and iid income with state space $Y = \{y_l, y_h\}$ but not (at least not by us) for more general economies. Also see [9] for a proof in a small open economy version of this model.

[45] As in our model, boundedness of the state space for assets from above has to be assured. [20] assumes that income can only take two values, but doesn't need the stochastic process to be *iid* over time nor any assumption on the period utility function. [3] assumes *iid* income and u to be bounded and to have bounded relative risk aversion -see his working paper. We do not require any boundedness assumption on u, but need the *iid* assumption.

[46] The theorems invoked for the existence of a unique invariant measure are similar in spirit; in particular they all require a "mixing condition" that assures that there is a unique ergodic set. In their setting agents with bad income shocks run down their assets, and good income shocks induce upward jumps in the asset position; in our setting agents with bad shocks move down in the entitlement distribution towards $U^{Aut}(y_{\min})$, with good shocks inducing jumps towards higher w, due to binding participation constraints.

[47] [20] provides no theoretical properties of the excess asset demand function, in [3] the presence of physical capital, which makes the supply of assets interest-elastic, assures (together with continuity of the asset demand function) the existence of a market-clearing interest rate.

[48] If no agent is unconstrained the allocation is autarkic and we can take $\frac{1}{R} = \beta \frac{u'(y_{\min})}{u'(y_{\max})}$.

This suggests that the equilibrium prices satisfy (with normalization of $p_0 = 1$)

$$p_t(y^t) = \frac{\pi(y^t|y_0)}{R^t} = p_t \pi(y^t|y_0). \tag{48}$$

with $p_t = R^{-t}$. That is, the price for a commodity delivered contingent on personal histories is composed of two components, an aggregate intertemporal price $p_t = R^{-t}$ and an individual specific, history dependent component, equal to the probability that the personal history occurs.

Given prices, the initial wealth level that makes the efficient consumption allocation affordable for an agent of type (w_0, y_0) is given by

$$a_0 = c_0(w_0, y_0) - y_0 + \sum_{t=1}^{\infty} \sum_{y^t|y_0} \frac{\pi(y^t|y_0)}{R^t} \left(c_t(w_0, y^t) - y_t\right) = a_0(w_0, y_0) < \infty \tag{49}$$

where the last inequality follows from the fact that the efficient consumption allocation is bounded from above.[49] Finally, the equilibrium consumption allocation corresponding to the efficient allocation is given by[50]

$$c_t(a_0, y^t) = c_t(a_0^{-1}(w_0, y_0), y^t) = C\left(h_t(w_0, y^t)\right). \tag{50}$$

The preceding discussion can be summarized in the following

Theorem 17 *Suppose that $\{h_t(w_0, y^t)\}_{t=0}^{\infty}$ is a stationary efficient allocation (with associated shadow interest rate $R > \frac{1}{\beta}$). Then prices $\{p_t(y^t)\}$ and allocations $\{c_t(a_0, y^t)\}$, as defined in (48) and (50) are an equilibrium for initial wealth distribution Θ_0 derived from Φ_0 and (49).*

Proof. The allocation satisfies the resource constraint (8) since the efficient allocation does and Θ_0 is derived from Φ_0. Also the allocation satisfies the continuing participation constraints, and, by construction of $a_0(w_0, y_0)$, the budget constraint. It remains to be shown that $\{c_t(a_0, y^t)\}$ is utility maximizing among the allocations satisfying the budget and the enforcement constraints. The first order conditions

$$(1-\beta)\beta^t \pi(y^t|y_0) u'\left(c_t(a_0, y^t)\right) \left(1 + \sum_{y^\tau \in P(y^t)} \mu(a_0, y^\tau)\right) = \lambda(a_0, y_0) p(y^t) \tag{51}$$

are sufficient for consumer optimality.[51] Define Lagrange multipliers $\mu(a_0, y_0) = 0$, $\lambda(a_0, y_0) = (1 - \beta) u'\left(c_0(a_0, y_0)\right)$ and recursively

$$1 + \sum_{y^\tau | y^t} \mu(a_0, y^\tau) = \frac{u'\left(c_0(a_0, y_0)\right)}{(\beta R)^t u'\left(c_t(a_0, y^t)\right)} \tag{52}$$

Note that the allocation by construction (see 47) satisfies $\frac{u'(c_t(a_0, y_0))}{\beta R u'(c_{t+1}(a_0, y^{t+1}))} \geq 1$, with equality if the limited enforcement constraint is not binding. Hence $\mu(a_0, y^{t+1}) \geq 0$ and $\mu(a_0, y^{t+1}) = 0$ if the constraint is not binding. By construction the allocation and multipliers satisfy the first order conditions. ■

[49] Therefore, to decentralize a particular stationary efficient consumption allocation we require a very particular initial distribution over initial assets. In this sense one of the primitives of our model, Θ_0, can't be chosen arbitrarily, which is true in all steady state analyses.

[50] Given that the optimal recursive policy function $h(., y)$ is a strictly increasing function in w, the $h_t(., y^t)$ and hence the $c_t(., y^t)$ are strictly increasing in w_0. Therefore $a_0(., y_0)$ is strictly increasing and thus invertible. We denote its inverse by a_0^{-1}.

[51] The consumer maximization problem is a strictly convex programming problem (the constraint set with the debt constraints remains convex). Note that since the efficient consumption allocation is bounded from above, the expected continuation utility from any history y^T onward, discounted at market prices R^{-T} goes to zero as $T \to \infty$ (i.e. the relevant transversality condition is satisfied). For details see the separate theoretical appendix, available at http://www.econ.upenn.edu/˜dkrueger/research/theoreticalapp.pdf.

D Proofs of Results in Main Text

Proof of Lemma 5:

We want to show that for all $\underline{w} \leq w < \hat{w} \leq \bar{w}$, $h(w) < h(\hat{w})$. Suppose not. Then from (17) $V_R'(g_{y'}(w)) \geq V_R'(g_{y'}(\hat{w}))$ for all y' such that $g_{y'}(\hat{w}) > U^{Aut}(y')$, and hence $U^{Aut}(y') < g_{y'}(\hat{w}) \leq g_{y'}(w)$ for all those y' by strict convexity of V_R. From promise keeping there must exist $\bar{y}'$ such that $g_{\bar{y}'}(w) < g_{\bar{y}'}(\hat{w}) = U^{Aut}(\bar{y}')$, a violation of the enforcement constraint.

Now, since h is strictly increasing in w, $C'(h(w)) < C'(h(\hat{w}))$. Suppose that $g_{y'}(w) > U^{Aut}(y')$. Then from (17) we have $V_R'(g_{y'}(w)) < V_R'(g_{y'}(\hat{w}))$ and from the strict convexity of V_R it follows that $g_{y'}(\hat{w}) > g_{y'}(w)$. Obviously, if $g_{y'}(w) = U^{Aut}(y')$ then $g_{y'}(\hat{w}) \geq g_{y'}(w)$.

Thus we conclude that either $g_{y'}(\hat{w}) > g_{y'}(w)$ or $g_{y'}(w) = g_{y'}(\hat{w}) = U^{Aut}(y')$ ■

Proof of Lemma 6:

V_R is strictly convex and differentiable. By assumption $g_{y'}(w) > U^{Aut}(y')$. Combining (17) and (18) we obtain $\beta R V_R'(w) = V_R'(g_{y'}(w))$. Since $R < \frac{1}{\beta}$ we have $V_R'(w) > V_R'(g_{y'}(w))$. By strict convexity of V_R the first result follows. Hence $g_{y'}(.)$ are always strictly below the 45^0 line in its strictly increasing part. But $g_{y'}(w) \geq U^{Aut}(y')$ for all w. Hence for $w < U^{Aut}(y')$ it follows that $g_{y'}(w) = U^{Aut}(y') > w$. By continuity of $g_{y'}(.)$ we obtain that $g_{y'}(U^{Aut}(y')) = U^{Aut}(y')$, and from the first result it follows that $g_{y'}(w) < w$ for all $w > U^{Aut}(y')$ ■

Proof of Theorem 7:

Take $\bar{w} = \max_y U^{Aut}(y) + \varepsilon$, for $\varepsilon > 0$. If $g_{y'}(w) > U^{Aut}(y')$, then the previous lemma yields the result. If $g_{y'}(w) = U^{Aut}(y')$, then $g_{y'}(w) = U^{Aut}(y') \leq \max_y U^{Aut}(y) < \bar{w}$

Proof of Theorem 8:

We first prove that there exists $w^* \in W$ such that $w^* > U^{Aut}(y_{\max})$ and $g_{y_{\max}}(w^*) = U^{Aut}(y_{\max})$, from which it follows that $g_{y_{\max}}(w^*) = U^{Aut}(y)$ for all $w \leq w^*$.

Suppose, to obtain a contradiction, that $g_{y_{\max}}(w) > U^{Aut}(y_{\max})$ for all $w \in W, w > U^{Aut}(y_{\max})$. Then by Lemma 5 we have $g_{y'}(w) = g_{y_{\max}}(w)$, for all $y' \in Y$ and all $w > U^{Aut}(y_{\max})$. By continuity of $g_{y'}$ and Lemma 6 we conclude that $g_{y'}(U^{Aut}(y_{\max})) = U^{Aut}(y_{\max})$, for all $y' \in Y$. But since $U^{Aut}(y_{\max}) > U^{Aut}(y')$ for all $y' \neq y_{\max}$, by Lemma 6 it follows that $g_{y'}(U^{Aut}(y_{\max})) < U^{Aut}(y_{\max})$ for all $y' \neq y_{\max}$, a contradiction.

We now can apply Stokey et al., Theorem 11.12. For this it is sufficient to prove there exists an $\varepsilon > 0$ and an $N \geq 1$ such that for all $(w, y) \in (W, Y)$ we have $Q^N((w, y, U^{Aut}(y_{\max}), y_{\max}) \geq \varepsilon$.

If $w^* \geq \bar{w}$ this is immediate, as then for all $(w, y) \in (W, Y)$, $Q((w, y, U^{Aut}(y_{\max}), y_{\max}) \geq \pi(y_{\max})$, since $g_{y_{\max}}(w) = U^{Aut}(y_{\max})$ for all $w \in W$. So suppose $w^* < \bar{w}$. Define

$$d = \min_{w \in [w^*, \bar{w}]} \{w - g_{y_{\max}}(w)\} \tag{53}$$

Note that d is well-defined as $g_{y_{\max}}$ is a continuous function and that $d > 0$ from Lemma 6 Define

$$N = \min\{n \in N | \bar{w} - nd \leq w^*\} \tag{54}$$

and $\varepsilon = \pi(y_{\max})^N$. Suppose an individual receives $y_{\max}$ for N times in a row, an event that occurs with probability ε. For (w, y) such that $w \leq w^*$ the result is immediate as for those $w, g_{y_{\max}}(w) = U^{Aut}(y_{\max})$ and $g_{y_{\max}}(U^{Aut}(y_{\max})) = U^{Aut}(y_{\max})$. For any $w \in (w^*, \bar{w}]$ we have $g_{y_{\max}}(w) \leq w - d$, $g_{y_{\max}}(g_{y_{\max}}(w)) \leq w - 2d$, etc. The result then follows by construction of (N, ε) ■