Sovereign Debt and Structural Reforms*

Andreas Müller† Kjetil Storesletten‡ Fabrizio Zilibotti§

February 15, 2018

Abstract

We construct a dynamic theory of sovereign debt and structural reforms with limited enforcement and moral hazard. A sovereign country in recession would like to smooth consumption and can make costly reforms to speed up recovery. The sovereign can renege on contracts by suffering a stochastic cost. The Constrained Optimum Allocation (COA) prescribes non-monotonic dynamics for consumption and effort and imperfect risk sharing. The COA features debt overhang: reform effort decreases when the recession lasts long. The COA is decentralized by a competitive Markov equilibrium with markets for renegotiable GDP-linked one-period debt. We also consider environments with less complete markets.

JEL Codes: E62, F33, F34, F53, H12, H63

Keywords: Debt Overhang, Default, Dynamic Principal-Agent Model, European Debt Crisis, Markov Equilibrium, Moral hazard, Renegotiation, Risk premia, Risk Sharing, Sovereign Debt, Structural Reforms.

---

*We would like to thank the editor, three referees, Manuel Amador, George-Marios Angeletos, Cristina Arellano, Marco Bassetto, Tobias Broer, Fernando Broner, Alessandro Dovis, Jonathan Eaton, Patrick Kehoe, Adrien Leclerc, Enrique Mendoza, Juan-Pablo Nicolini, Ugo Panizza, Victor Rios-Rull, Ctirad Slavik, Aleh Tsyvinski, Jaume Ventura, Christopher Winter, Tim Worrall, and seminar participants at Annual Meeting of the Swiss Society of Economics and Statistics, Barcelona GSE Summer Forum, Brown University, CEMFI, Columbia University, CERGE-EI, CREi, EIEF Political Economy Workshop, ESSIM 2016, European University Institute, Goethe University Frankfurt, Graduate Institute of Geneva, Humboldt University, Istanbul School of Central Banking, LMU Munich, NORMAC, Oxford, Royal Holloway, Swiss National Bank, Università Cà Foscari, Universitat Autonoma Barcelona, University College London, University of Cambridge, University of Konstanz, University of Mannheim, University of Oslo, University of Oxford, University of Pennsylvania, University of Padua, University of Toronto, University of Zurich, and Yale University. We acknowledge support from the European Research Council (ERC Advanced Grant IPCDP-324085).

†University of Essex, Department of Economics, andreas.mueller@essex.ac.uk.

‡University of Oslo, Department of Economics, kjetil.storesletten@econ.uio.no.

§Yale University, Department of Economics, fabrizio.zilibotti@yale.edu.

In this paper, we propose a normative and positive dynamic theory of sovereign debt in an environment characterized by informational frictions. The theory rests on two building blocks. The first is that sovereign debt is subject to limited enforcement, and that countries can renege on their obligations subject to real costs as in, e.g., Aguiar and Gopinath (2006), Arellano (2008) and Yue (2010). The second is that countries can undertake *structural* policy reforms to speed up recovery from an existing recession.[1] The reform effort is assumed to be unobservable and subject to moral hazard.

The theory is motivated by recent episodes during the Great Recession, especially in Europe, where sovereign debt crises and economic reforms have been salient intertwined policy issues. Greece, for instance, saw its debt-GDP ratio soar from 103% in 2007 to 172% in 2011 in spite of a 53% haircut agreed with foreign creditors in 2011. While creditors and international organizations pushed hard the Greek government to introduce structural reforms that would help the economy recover and meet its international financial obligations, such reforms were forcefully opposed internally. Opposers maintained that the reforms would imply major sacrifice for domestic residents while a large share of the benefits would accrue to foreign lenders. Meanwhile, international organizations stepped in to provide financial assistance and access to new loans, asking in exchange fiscal restraint and a commitment to economic reforms. Our theory rationalizes these dynamics.

The model economy is a dynamic endowment economy subject to income shocks following a two-state Markov process. The economy (henceforth, the *sovereign*) starts in a unique recession whose duration is stochastic. Costly structural reforms increase the probability that the recession ends. Consumers' preferences induce a desire for consumption and effort smoothing. We first characterize the solution of two planning problems: the first best and the constrained optimum allocation (COA) subject to limited enforcement and moral hazard. In the first best, the planner provides the sovereign country with full insurance by transfering resources to it during recession and reversing the transfer as soon as the recession ends. The sovereign exerts the efficient level of costly effort for as long as the recession lasts.

The first best is not implementable in the presence of informational frictions for two reasons. First, the sovereign has access to a stochastic outside option whose realization is publicly observable. This creates scope for opportunistic deviations involving cashing in transfers for some time, and then unilaterally quit (i.e., *default on*) the insurance contract as soon as the realization of the outside option is sufficiently favorable. Second, the sovereign has an incentive to shirk and rely on the transfers rather than exerting the required reform effort to increase output.

The COA is characterized by means of a promised utility approach in the vein of Spear and Srivastava (1987), Thomas and Worrall (1988 and 1990), and Kocherlakota (1996). The optimal contract is subject to an incentive compatibility constraint (IC) that pins down the effort choice and a participation constraint (PC) that captures the limited commitment. The COA has the following features: throughout recession, within spells of slack PC (i.e., when the realized cost of default is high), the planner front-loads the sovereign's consumption and decreases it over time in order to provide dynamic incentives for reform effort (as in Hopenhayn and Nicolini 1997). In this case, the solution is dictated by the IC and is history-dependent: consumption and promised utility fall over time, while effort follows non-monotonic dynamics for reasons to which we return below. Whenever the PC binds (i.e., the sovereign faces an attractive outside option), the planner increases discretely

[1] Examples of such reforms include labor and product market deregulation, and the establishment of fiscal capacity that allows the government to raise tax revenue efficiently (see, e.g., Ilzkovitz and Dierx 2011). While these reforms are beneficial in the long run, they entail short-run costs for citizens at large, governments or special-interest groups (see, e.g., Blanchard and Giavazzi 2003).

consumption and promised utility in order to prevent the sovereign from leaving the contract.

Next, we move to the decentralization of the COA that is the main contribution of the paper. We show that the COA can be implemented by a competitive Markov equilibrium where the sovereign issues two one-period securities paying returns contingent on the aggregate state of the economy (*GDP-linked bonds*). The bonds are sold to profit-maximizing international creditors who hold a well-diversified portfolio. They are defaultable and renegotiable. When she faces a low realization of the default cost, the sovereign could, in principle, default, pay a cost, and restart afresh with zero debt. However, costly default can be averted by renegotiation: when a credible default threat is on the table, a syndicate of creditors makes a take-it-or-leave-it debt haircut offer, as in Bulow and Rogoff (1989). In equilibrium, there is no outright default, but recurrent debt renegotiations.[2] The key for this market arrangement to attain constrained efficiency is that creditors have all the bargaining power, ex-post. Importantly, this market arrangement is Markovian, and does not rely on complicated mechanisms to coordinate future punishment.

That a Markov equilibrium with only two assets decentralizes the COA is at first view surprising. In our environment, there is a continuum of states associated with the possible realization of the stochastic process, and only two securities suggesting that financial markets cannot span the full set of states. Moreover, there is moral hazard, and creditors cannot commit to punish opportunistic behavior, in contrast with the planner, who can design dynamic incentives under full commitment. The decentralization result hinges on two features. First, the risk of renegotiation (where renegotiations follow a particular protocol) introduces a sufficient degree of state contingency to attain constrained efficiency. Second, the equilibrium debt dynamics and its endogenously evolving price provide efficient dynamic incentives for the sovereign to exert the second-best reform effort.

In the competitive equilibrium, debt accumulates and consumption falls over time as long as the recession continues and debt is not renegotiated. Interestingly, the reform effort is a non-monotonic function of debt. This result stems from the interaction between limited enforcement and moral hazard. Under full enforcement, effort would increase monotonically over the recession as debt accumulates. Absent moral hazard, effort would be constant when the PC is slack and decrease every time debt is renegotiated. When both informational constraints are present, effort increases with debt at low level. However, for sufficiently high debt levels the relationship is flipped: there, issuing more debt deters reforms because, due to the high probability of renegotiation, most of the gains from an economic recovery would accrue to foreign lenders in the form of capital gains on the outstanding debt. In this region, the reform effort falls over time as debt accumulates. This *debt overhang* curtails consumption smoothing: when sovereign debt is high, investors expect low reform effort, are pessimistic about the economic outlook, and request even higher risk premia. Interestingly, this debt overhang is constrained efficient under the informational constraints postulated in our theory.

After deriving the main decentralization result, we consider economies with more market incompleteness. In particular, we consider an economy in the spirit of Eaton and Gersovitz (1981) where the sovereign can issue only one bond whose return is not contingent on the level of GDP. This economy fails to attain the COA: the sovereign has access to less consumption smoothing and ends up providing an inefficient effort level.[3] This extension is interesting because in reality markets for GDP-linked bonds are often missing. In the (arguably realistic) one-asset environment, there is scope for policy in-

[2] This approach conforms with the empirical observations that unordered defaults are rare events, and that there is great heterogeneity in the terms at which debt is renegotiated, as documented by Tomz and Wright (2007) and Sturzenegger and Zettelmeyer (2008).

[3] We also study a case in which we do not allow renegotiation. This reduces further reduces consumption and effort smoothing and welfare.

tervention. In particular, an international institution such as the IMF can improve welfare by means of an assistance program. During the recession, the optimal program entails a persistent budget support through extending loans on favorable terms. When the recession ends, the sovereign is settled with a (large) debt on market terms. We also discuss the possibility that the international institution takes control over the reform process overcoming the friction associated with the non-contractible nature of the reform effort.

Our analysis is related to a large international and public finance literature. In a seminal contribution, Atkeson (1991) studies the optimal contract in an environment in which an infinitely-lived borrower faces a sequence of two-period lived lenders. The borrower can use funds to invest in productive future capacity or to consume the funds. Contracting is subject to moral hazard as lenders cannot observe the allocation to investment or consumption. Our paper differs from Atkeson's in various aspects. First, the environment is different: in our theory, all agents have an infinite horizon and investments in structural reforms affect the future stochastic process of income, while in Atkeson's model investments only affect next period's income. Second, we provide a novel COA decentralization result through a Markov equilibrium with renegotiable one-period bonds. Third, Atkeson (1991) emphasizes the result that the optimal contract involves capital outflow from the borrower during the worst aggregate state. Our model predicts instead that in a recession the borrower keeps accumulating debt and renegotiates it periodically.

A number of recent papers deal with the dynamics of sovereign debt under a variety of informational and contractual frictions. Dovis (2017) studies the efficient risk-sharing arrangement between international lenders and a sovereign borrower with limited commitment and private information about domestic productivity. In his model the constrained efficient allocation can be implemented as a competitive equilibrium with non-contingent defaultable bonds of short and long maturity. He does not consider the interaction between structural reforms and limited commitment. Aguiar *et al.* (2017) study a model à la Eaton and Gersovitz (1981) with limited commitment assuming, as we do, that the borrower has a stochastic default cost. Their research is complementary to ours insofar as they focus on debt maturity in rollover crises, from which instead we abstract. Jeanne (2009) also studies a rollover crisis in an economy where the government takes a policy action that affects the return to foreign investors (e.g., the enforcement of creditor's right) but this can be reversed within a time horizon that is shorter than that at which investors must commit their resources.

Our work is also related to the literature on debt overhang initiated by Krugman (1988). He constructs a static model with exogenous debt showing that a large debt can deter the borrower from undertaking productive investments. In this regime, it may be optimal for the creditor to forgive debt. A number of papers consider distortions associated with high indebtedness in the presence of informational imperfections. Aguiar and Amador (2014) show that high debt increases the volatility of consumption by reducing risk sharing. Aguiar, Amador, and Gopinath (2009) consider the effect of debt on investment volatility. When an economy is indebted, productivity shocks gives rise to larger dispersion in investment rates. Aguiar and Amador (2011) consider a politico-economic model where capital income is subject to risk of ex-post expropriation and the government can default on external debt. A country with a large sovereign debt position has a greater temptation to default, and therefore investments are low. Conesa and Kehoe (2015) construct a theory predicting that the governments of highly indebted countries may opt to gamble for redemption.

Our research is related also to the literature on endogenous incomplete markets due to limited enforcement or limited commitment. This includes Alvarez and Jermann (2000) and Kehoe and Perri (2002). The analysis of constrained efficiency is related to the literature on competitive risk sharing contracts with limited commitment, including, among others, Thomas and Worrall (1988), Marcet and Marimon (1992), Phelan (1995), Kocherlakota (1996), and Krueger and Uhlig (2006).

Finally, our work is related, more generally, to recent quantitative models of sovereign default such as Aguiar and Gopinath (2006), Arellano (2008), and Chatterjee and Eyigungor (2012).[4] Of particular interest is Abraham *et al.* (2017) who study an Arellano (2008) economy, extended to allow shocks to government expenditures, and compare quantitative outcomes of the market allocation to the optimal design of a Financial Stability Fund, interpreted as the solution to a planning problem. Broner *et al.* (2010) study the implications for the incentives to default of having part of the government debt held by domestic residents. Song *et al.* (2012) and Müller *et al.* (2016) focus on Markov equilibria to study the politico-economic determination of debt in open economies where governments are committed to honor their debt.

The rest of the paper is organized as follows. Section 1 describes the model environment. Section 2 solves for the first best and the COA under limited commitment and moral hazard. Section 3 provides the main decentralization result. Section 4 considers a one-asset economy, and discuss policy interventions to restore efficiency. Section 5 concludes. Two online appendixes contain, respectively, the proofs of the main propositions and lemmas (Appendix A) and additional technical material referred in the text (Appendix B).

1 The model environment

The model economy is a small open endowment economy populated by an infinitely-lived representative agent. A benevolent *sovereign* makes decisions on behalf of the representative agent. The stochastic endowment follows a two-state Markov switching process, with realizations $\underline{w}$ and $\bar{w}$, where $0 < \underline{w} < \bar{w}$. We label the two endowment states *recession* and *normal times*, respectively. An economy starting in recession remains in the recession with probability $1-p$ and switches to normal times with probability p. Normal times is assumed to be an absorbing state.[5] This assumption aids tractability and enables us to obtain sharp analytical results. During recession, the sovereign can implement a costly reform policy to increase p. In our notation p denotes both the *reform effort* and the probability that the recession ends. The sovereign can smooth consumption by contracting with a financial intermediary that has access to an international market offering a gross return R.

The sovereign's preferences are described by the following expected utility function:

$$E_0 \sum \beta^t \left[u\left(c_t\right) - \phi_t I_{\{\text{default in } t\}} - X\left(p_t\right) \right].$$

The utility function u is twice continuously differentiable and satisfies $\lim_{c \to 0} u(c) = -\infty$, $u'(c) > 0$, and $u''(c) < 0$. $I \in \{0, 1\}$ is an indicator switching on when the economy is in a default state and ϕ is an associated utility loss. In the planning allocation, the cost ϕ accrues when the sovereign opts out of the dynamic contract offered by the planner. In the market allocation it accrues when sovereign reneges on a debt contract with international lenders.[6] In recession, ϕ follows a stochastic process assumed to

[4] Other papers studying restructuring of sovereign debt include Asonuma and Trebesch (2016), Bolton and Jeanne (2007), Hatchondo *et al.* (2014), Mendoza and Yue (2012), and Yue (2010).

[5] In a previous version of this paper, we considered the possibility that the economy could fall recurrently into recession.

[6] The cost ϕ is exogenous and publicly observed, and captures in a reduced form a variety of shocks including both taste shocks (e.g., the sentiments of the public opinion about defaulting on foreign debt) and institutional shocks (e.g., the election of a new prime minister, a new central bank governor taking office, the attitude of foreign governments, etc.). Alternatively, ϕ could be given a politico-economic interpretation, as reflecting special interests of the groups in power. For instance, the government may care about the cost of default to its constituency rather than to the population at large. In the welfare analysis, we stick to the interpretation of a benevolent government and abstract from politico-economic

be *i.i.d.* over time and to be drawn from the p.d.f. $f(\phi)$ with an associated c.d.f. $F(\phi)$. We assume that $F(\phi)$ is continuously differentiable everywhere, and denote its support by $\aleph \equiv [\phi_{\min}, \phi_{\max}] \subseteq \mathbb{R}^{+}$, where $0 \leq \phi_{\min} < \phi_{\max} < \infty$. The assumption that shocks are independent is for simplicity. In order to focus on debt dynamics in recessions we assume that there is full enforcement in normal times (i.e., in normal time ϕ is arbitrarily large). This is again for simplicity. In an earlier version of the paper (Müller *et al.* 2015), we show that our main results are robust to assuming that the distribution of ϕ is the same irrespective of whether the economy is in normal times or recession.

The function $X(p)$ represents the cost of reform, assumed to be increasing and convex in the probability of exiting recession, $p \in [\underline{p}, \bar{p}] \subset [0,1]$. X is assumed to be twice continuously differentiable, with the following properties: $X(\underline{p}) \geq 0$, $X'(\underline{p}) = 0$, $X'(p) > 0\ \forall p > \underline{p}$, $X''(p) > 0$, and $\lim_{p \to \bar{p}} X'(p) = \infty$. In normal times, $X = 0$.

The time line of events is as follows: at the beginning of each period, the aggregate state (recession or normal time) is observed; then, ϕ is realized and immediately becomes common knowledge; finally, effort is exercised. Throughout the paper, we focus on an economy starting in a recession.

2 Planning allocation

We first characterize a planning allocation. The planning problem is formulated as a one-sided commitment program following a promised-utility approach. The planner has access to a storage technology with a return $R = 1/\beta$ and maximizes profits subject to a promise-keeping constraint.[7]

Let ν denote the promised utility, i.e., the expected utility promised to the sovereign in the beginning of the period, before the realization of ϕ is observed. ν is the key state variable of the problem. We denote by ω_ϕ and $\bar{\omega}_\phi$ the promised continuation utilities conditional on the realization ϕ and on the economy staying in recession or switching to normal time, respectively. We denote by $P(\nu)$ the expected present value of profits accruing to the planner conditional on delivering the promised utility ν in the most cost-effective way.

The optimal value $P(\nu)$ satisfies the following functional equation:

$$P(\nu) = \max_{\{c_\phi, p_\phi, \omega_\phi, \bar{\omega}_\phi\}_{\phi \in \aleph}} \int_\aleph \left[\underline{w} - c_\phi + \beta\left((1 - p_\phi) P(\omega_\phi) + p_\phi \bar{P}(\bar{\omega}_\phi)\right)\right] dF(\phi), \tag{1}$$

where the maximization is subject to the promise-keeping constraint

$$\int_\aleph \left[u(c_\phi) - X(p_\phi) + \beta\left((1 - p_\phi)\omega_\phi + p_\phi \bar{\omega}_\phi\right)\right] dF(\phi) \geq \nu; \tag{2}$$

the planner's profit function in normal times,

$$\bar{P}(\bar{\omega}_\phi) = \max_{\bar{c} \in [0, \bar{w}]} \bar{w} - \bar{c} + \beta \bar{P}\left(\beta^{-1}\left[\bar{\omega}_\phi - u(\bar{c})\right]\right); \tag{3}$$

and

$$c_\phi \geq 0,\ p_\phi \in [\underline{p}, \bar{p}],\ \nu, \omega_\phi \geq \alpha - E[\phi]. \tag{4}$$

factors, although the model could be extended in this direction.

[7] The problem would be identical under two-sided lack of commitment under some restrictions on the state space. In particular, the problem would require some upper bound on the sovereign's initial promised utility to ensure that the principal does not find it optimal to exit the contract in some states of the world. We return to this point below.

Our main insights carry over to the case in which $\beta R < 1$.

For now we treat α simply as a parameter capturing the value for the sovereign of not being in the contract. Its interpretation is discussed in more detail below. The constraint $\omega_\phi \geq \alpha - E[\phi]$ reflects the fact that the sovereign could attain the expected utility $\alpha - E[\phi]$ by defaulting in every state of the world.

We introduce two informational frictions. The first is limited enforcement: the sovereign can quit the contract when the realization of ϕ is sufficiently low. This is captured by a participation constraint (PC). The second is moral hazard: the reform effort is chosen by the sovereign and is not observed by the planner. This is captured by an incentive-compatibility constraint (IC).

When the planning problem is subject to a limited enforcement problem, the allocation is subject to the following PC:

$$u(c_\phi) - X(p_\phi) + \beta((1-p_\phi)\omega_\phi + p_\phi\bar{\omega}_\phi) \geq \alpha - \phi, \quad \phi \in \aleph. \tag{5}$$

Here, $\alpha - \phi$ is the sovereign's stochastic outside option if she quits the contract, where ϕ is the shock discussed above.

When the planning problem is also subject to moral hazard, the allocation is subject to the following IC:

$$p_\phi = \tilde{p}_\phi = \arg\max_{p\in[\underline{p},\bar{p}]} -X(p) + \beta((1-p)\omega_\phi + p\bar{\omega}_\phi). \tag{6}$$

Note that the sovereign chooses effort *after* the planner has set the future promised utilities.

2.1 First best

We start by characterizing the first-best allocation given by the program (1)–(4). The proof, which follows standard methods, can be found in Appendix B.

Proposition 1 *Given a promised utility ν, the first best allocation satisfies the following properties. The sequences for consumption and promised utilities are constant at the level $c^{FB}(\nu)$, $\omega^{FB} = \nu$, and $\bar{\omega}^{FB} = \nu + X\left(p^{FB}(\nu)\right)/\left(1-\beta\left(1-p^{FB}(\nu)\right)\right)$, implying full consumption insurance, $c^{FB}(\nu) = \bar{c}^{FB}\left(\bar{\omega}^{FB}\right)$. Moreover, effort is constant during recession at the level $p^{FB}(\nu)$ and the functions $c^{FB}(\nu)$ and $p^{FB}(\nu)$ are strictly increasing and strictly decreasing functions, respectively, satisfying:*

$$\frac{\beta}{1-\beta(1-p^{FB}(\nu))}\left(\underbrace{(\bar{w}-\underline{w})\times u'\left(c^{FB}(\nu)\right)}_{\text{output increase if recovery}} + \underbrace{X\left(p^{FB}(\nu)\right)}_{\text{saved effort cost if recovery}}\right) = X'\left(p^{FB}(\nu)\right) \tag{7}$$

$$\frac{u\left(c^{FB}(\nu)\right)}{1-\beta} - \frac{X\left(p^{FB}(\nu)\right)}{1-\beta(1-p^{FB}(\nu))} = \nu. \tag{8}$$

The solution to the functional equation (3) in normal times is given by

$$\bar{P}(\bar{\omega}_\phi) = \frac{\bar{w} - \bar{c}^{FB}(\bar{\omega}_\phi)}{1-\beta}, \quad \bar{c}^{FB}(\bar{\omega}_\phi) = u^{-1}((1-\beta)\bar{\omega}_\phi). \tag{9}$$

The first-best allocation features perfect insurance: the sovereign enjoys a constant consumption irrespective of the aggregate state and exerts a constant reform effort during recession. Moreover, consumption c^{FB} is strictly increasing in ν while effort p^{FB} is strictly decreasing in ν: a higher promised utility is associated with higher consumption and lower effort.

2.2 Constrained Optimum Allocation (COA)

Next, we characterize the COA. The planning problem (1)-(4) is subject to the PC (5) and to the IC (6). Note that the planning problem is evaluated after the uncertainty about the endowment state has been resolved, but before the realization of ϕ. For didactic reasons, we first study limited enforcement separately, assuming that effort is controlled directly by the planner. Then, we generalize the analysis to the more interesting case in which there is moral hazard. We start by establishing a property of the COA that holds true in both environments.

Lemma 1 *Assume that the profit functions P and $\bar{P}$ are strictly concave. Define the sovereign's discounted utility conditional on the promised utility ν and the realization ϕ as $\mu_\phi(\nu) \equiv u(c_\phi(\nu)) - X(p_\phi(\nu)) + \beta\left[(1-p_\phi(\nu))\omega_\phi(\nu) + p_\phi(\nu)\bar{\omega}_\phi(\nu)\right]$. Then, the COA features a unique threshold function $\tilde{\phi}(\nu)$ such that the PC binds if $\phi \in \left[\phi_{\min}, \tilde{\phi}(\nu)\right]$ and is slack if $\phi \in \left[\tilde{\phi}(\nu), \phi_{\max}\right]$. Moreover,*

$$\mu_\phi(\nu) = \begin{cases} \alpha - \phi & \text{if} \quad \phi \in \left[\phi_{\min}, \tilde{\phi}(\nu)\right], \\ \alpha - \tilde{\phi}(\nu) & \text{if} \quad \phi \in \left[\tilde{\phi}(\nu), \phi_{\max}\right]. \end{cases}$$

The functions P and $\bar{P}$ are value functions of the planning problem. Proving that $\bar{P}$ is strictly concave is straightforward. The strict concavity of P is more difficult to establish analytically. We will prove that P is strictly concave in the case without moral hazard, while in the case with moral hazard we will guess and verify it numerically. The lemma formalizes the intuitive property that (i) if in a state ϕ_a the planner promises the agent more than her reservation utility, it is then optimal for her to do so for all $\phi > \phi_a$; moreover, promised utility is equalized across such states; (ii) if in a state ϕ_b the planner promises the agent the reservation utility, it is then optimal for her to do so for all $\phi < \phi_a$.

2.2.1 Limited Enforcement without Moral Hazard

In this environment, there is no IC and the planner chooses directly the constrained efficient effort level. We prove in Appendix B (Proposition 8) that the planner's profit function $P(\nu)$ is decreasing, strictly concave and continuously differentiable at interior $\nu > \alpha - E[\phi]$. Moreover, the first-order conditions (FOCs) of the planning problem are necessary and sufficient. The proof follows the strategy in Thomas and Worrall (1990). Since in the normal state there is full enforcement, then, $\bar{c}(\bar{\omega}_\phi)$ and $\bar{P}(\bar{\omega}_\phi)$ are as in the first best (cf. Equation 9).

Combining the FOCs with respect to c_ϕ, ω_ϕ, and $\bar{\omega}_\phi$ with the envelope condition (see proof of Proposition 2 in Appendix B) yields:[8]

$$\frac{1}{u'(\bar{c}(\bar{\omega}_\phi))} - \frac{1}{u'(c_\phi)} = 0 \tag{10}$$

$$u'(c_\phi) = -\frac{1}{P'(\omega_\phi)}, \quad \forall \omega_\phi > \alpha - E[\phi]. \tag{11}$$

Combining these with the FOC with respect to p_ϕ yields

$$X'(p_\phi) = \beta\left((\bar{\omega}_\phi - \omega_\phi) - u'(c_\phi) \times \left(\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi)\right)\right). \tag{12}$$

[8]Recall that the first-best allocation implies that $\bar{P}'(\bar{\omega}_\phi) = -1/u'(\bar{c}(\bar{\omega}_\phi))$. Thus, equations (10) and (11) imply that the marginal profit loss associated with promised utilities is equalized across the two aggregate states when $\omega_\phi > \alpha - E[\phi]$, i.e., $\bar{P}'(\bar{\omega}_\phi) = P'(\omega_\phi)$.

Equation (10) establishes that the planner provides the sovereign with full insurance against the realization of the stochastic process for GDP, i.e., she sets $\bar{c}\left(\bar{\omega}_{\phi}\right)=c_{\phi}$ and equates the marginal profit loss associated with promised utilities in the two aggregate states. Equation (12) establishes that effort is set at the constrained efficient level. The marginal cost (left hand-side) equals two sum of the marginal benefits accruing to the sovereign and to the planner, respectively (right hand-side).

Proposition 2 below provides a formal characterization of the COA.[9]

Proposition 2 *The constrained optimal allocation (COA) is characterized as follows. Let ν denote the initial promised utility. Then, the threshold function $\tilde{\phi}(\nu)$ is decreasing and implicitly defined by the condition*

$$\nu=\alpha-\left[\int_{\phi_{\min}}^{\tilde{\phi}(\nu)}\phi dF(\phi)+\tilde{\phi}(\nu)\left[1-F(\tilde{\phi}(\nu))\right]\right], \tag{13}$$

such that:

1. *If $\phi<\tilde{\phi}(\nu)$, the PC is binding, and the solution for $\left(c_{\phi},p_{\phi},\omega_{\phi},\bar{\omega}_{\phi}\right)$ is determined by (10), (11), (12), and by (5) holding with equality. Moreover, $\omega_{\phi}>\nu$.*

2. *If $\phi\geq\tilde{\phi}(\nu)$, the PC is slack, and the solution for $\left(c_{\phi},p_{\phi},\omega_{\phi},\bar{\omega}_{\phi}\right)$ is given by $\omega_{\phi}=\nu$, $c_{\phi}=c\left(\nu\right)$, $\bar{\omega}_{\phi}=\bar{\omega}\left(\nu\right)$, and $p_{\phi}=p\left(\nu\right)$, where the functions $c\left(\nu\right)$, $\bar{\omega}\left(\nu\right)$ and $p\left(\nu\right)$ are determined by*

   $$u(c(\nu))-X(p(\nu))+\beta\left[p(\nu)\bar{\omega}(\nu)+(1-p(\nu))\nu\right]=\alpha-\tilde{\phi}(\nu),$$

   (10), and (12), respectively. The solution is history-dependent. The reform effort is strictly decreasing and consumption and future promised utility are strictly increasing in ν.

The COA under limited enforcement has standard back-loading properties (cf. Ljungqvist and Sargent 2012, ch. 20). Whenever the sovereign's PC is not binding, consumption, effort and promised utility remain constant over time. Consumption remains constant even as the recession ends. Thus, the COA yields full consumption insurance across all states in which the PC is slack. Whenever the PC binds, the planner increases the sovereign's consumption and promised utilities while reducing her effort in order to meet her PC.[10]

The upper panels of Figure 1 describe the dynamics of consumption and effort (left panel), and promised utilities (right panel) under a particular sequence of realizations of ϕ.[11] In this numerical example, the recessions last for 13 periods. Thereafter, the economy attains full insurance. Consumption and effort are initially low and high, respectively. In periods 9 and 11, the PC binds and the planner must increase consumption and promised utility, and reduce effort. Consumption and effort remain constant when the PC is slack. Note also that consumption remains constant when the recession ends.

[9] The proof strategy builds on Thomas and Worrall (1990) and can be found in Appendix B. The proof of Proposition 3 below includes the proof of Proposition 2 as a special case.

[10] Although we have assumed that the planner controls effort directly, the same allocation would obtain if the planner did not control effort ex ante but could observe it ex post and punish deviations. In particular, if the sovereign did not exert the effort level prescribed by the optimal contract, the planner would terminate the contract in the following period and force the sovereign to pay the ensuing realization of ϕ. Since ϕ is assumed to be arbitrarily large in normal times (and $p\geq\underline{p}>0$), this threat would be sufficient to discipline the sovereign and sustain the allocation of Proposition 2.

[11] In all numerical examples, the parameters of the model are calibrated as described in Appendix B.

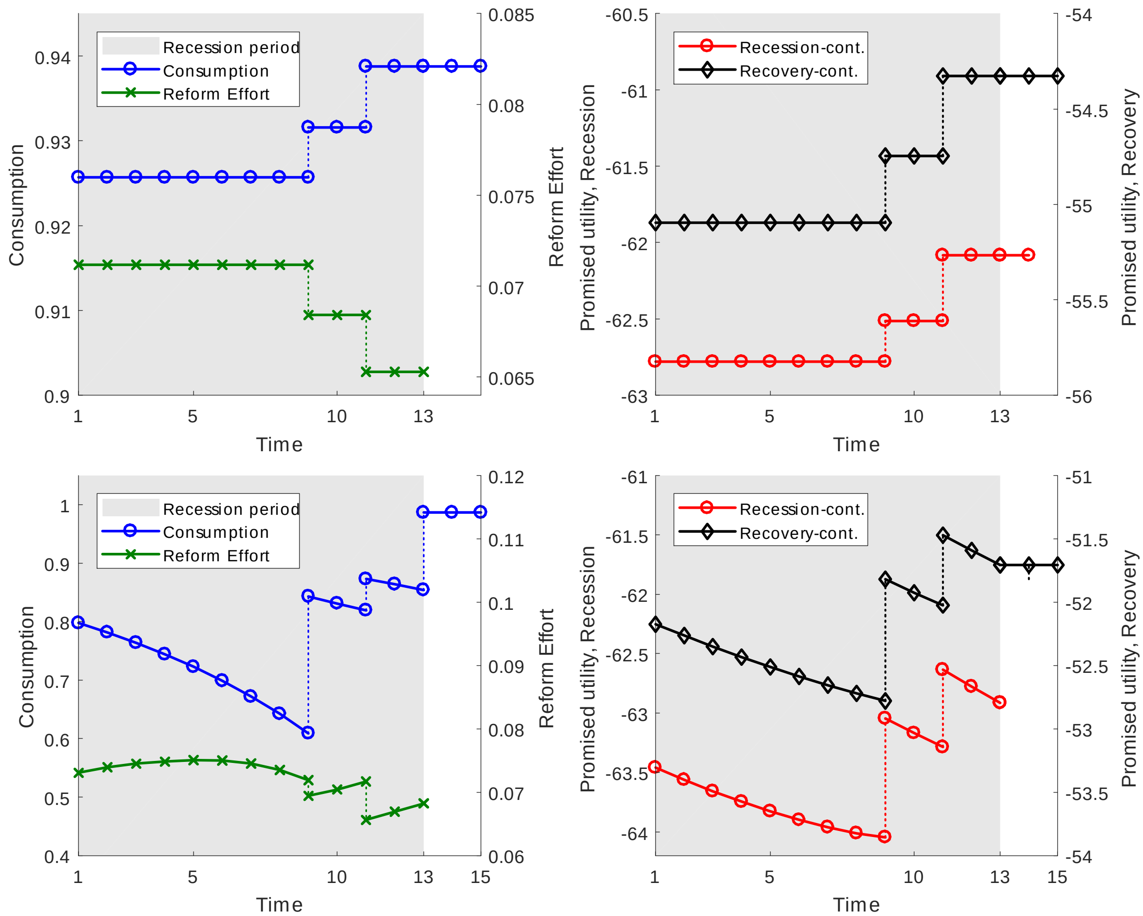

Figure 1: Simulation of consumption, effort, and promised utilities for a particular sequence of ϕ's. In this particular simulation the recession ends in period 13. The top panels show the planner allocation without moral hazard, the bottom panels with moral hazard.

2.2.2 Limited Enforcement with Moral Hazard

Next, we consider the more interesting case in which the planner cannot observe effort and is subject to both a PC and an IC. The COA features an important qualitative difference from the earlier special case: within each spell in which the PC is slack, the planner front-loads consumption and promised utility to incentivize the sovereign to provide effort. Therefore, moral hazard prevents full insurance even across the states of nature in which the PC is slack.

Let us start from the IC (6). The FOC yields $X'(\tilde{p}_\phi) = \beta(\bar{\omega}_\phi - \omega_\phi)$, or equivalently

$$\tilde{p}_\phi = \Upsilon(\bar{\omega}_\phi - \omega_\phi). \tag{14}$$

where $\Upsilon(x) \equiv (X')^{-1}(\beta x)$. The properties of X and its inverse imply that $\tilde{p}_\phi$ is increasing in the promised utility difference $\bar{\omega}_\phi - \omega_\phi$. Equation (14) is the analogue of (12). Effort is distorted because the sovereign does not internalize the benefits accruing to the planner.

The FOCs with respect to ω_ϕ and $\bar{\omega}_\phi$, together with the envelope condition yield (see proof of

Proposition 3 in Appendix A):

$$\begin{aligned}
\frac{1}{u'(\bar{c}(\bar{\omega}_\phi))} - \frac{1}{u'(c_\phi)} &= \frac{\Upsilon'(\bar{\omega}_\phi - \omega_\phi)}{\Upsilon(\bar{\omega}_\phi - \omega_\phi)} \left[\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi)\right] &(15)\\
\frac{1}{u'(c_\phi)} &= -P'(\omega_\phi) + \frac{\Upsilon'(\bar{\omega}_\phi - \omega_\phi)}{1-\Upsilon(\bar{\omega}_\phi - \omega_\phi)} \left[\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi)\right], \quad \forall \omega_\phi > \alpha - E[\phi] &(16)\\
0 &= \theta_\phi \times [\omega_\phi - (\alpha - E[\phi])], &(17)
\end{aligned}$$

where $\theta_\phi \geq 0$ is the Lagrange multiplier on the constraint $\omega_\phi \geq \alpha - E[\phi]$.[12] This constraint binds when ν is sufficiently low. Equation (16) is then replaced by the complementary slackness condition (17) implying that $\omega_\phi = \alpha - E[\phi]$ and $\theta_\phi > 0$.

The FOC (15) is the analogue of (10). With moral hazard, the planner does no longer provide the sovereign with full insurance against the realization of the aggregate state: by promising a higher consumption if the economy recovers (thereby curtailing insurance), she incentivizes effort provision. Note that the effort wedge is proportional to the elasticity of effort Υ'/Υ.

The FOC (16) is the analogue of (11). There, the marginal utility of consumption simply equaled the profit loss associated with an increase in promised utility. Here, the planner finds it optimal to open a wedge that is proportional to the elasticity of effort: she front-loads consumption in order to make the sovereign more eager to leave a recession.

We can now proceed to a full characterization of the COA with moral hazard. As is common for problems with both limited enforcement and moral hazard, it is difficult to prove that the program is globally concave and to establish analytically that the profit function is concave. We therefore assume that $P(\nu)$ is strictly concave in ν, and verify this property numerically.[13]

Proposition 3 *Suppose that effort is not observable. Assume that P is strictly concave. Then, P is differentiable at the interior of its support, and the optimal contract satisfies the following properties: (i) $p_\phi = \Upsilon(\bar{\omega}_\phi - \omega_\phi)$ as in (14), and (ii) the threshold function $\tilde{\phi}(\nu)$ is decreasing and implicitly defined by equation (13), and such that:*

1. *If $\phi < \tilde{\phi}(\nu)$, the PC is binding, and the solution for $(c_\phi, \omega_\phi, \bar{\omega}_\phi, \theta_\phi)$ is determined by (15), (16), (17), and by (5) holding with equality.*

2. *If $\phi \geq \tilde{\phi}(\nu)$, the PC is slack, and the solution for $(c_\phi, \omega_\phi, \bar{\omega}_\phi, \theta_\phi)$ is determined by (15), (16),(17), and*
$$u(c_\phi) - X(p_\phi) + \beta\left[p_\phi \bar{\omega}_\phi + (1-p_\phi)\omega_\phi\right] = \alpha - \tilde{\phi}(\nu). \tag{18}$$
The solution is history-dependent, i.e., $c_\phi = c(\nu)$, $\omega_\phi = \omega(\nu)$, and $\bar{\omega}_\phi = \bar{\omega}(\nu)$. Promised utility falls over time, i.e., $\omega_\phi = \omega(\nu) \leq \nu$, with $\omega(\nu) < \nu$ for $\theta_\phi = 0$ (i.e., if ν is sufficiently large). The function $c(\nu)$ is strictly increasing. Effort is non-monotone, being strictly increasing in ν at low levels of ν and strictly decreasing in ν at high levels of ν.

[12] Recall that the planner cannot promise a utility below $\alpha - E[\phi]$. This constraint adds the Lagrange multiplier θ_ϕ to the problem (see Appendix A for details). The issue did not arise in the case without moral hazard because the optimal choice of ω_ϕ was non-decreasing, making the lower constraint on ω_ϕ never binding. Nor does the problem arise for the choice of $\bar{\omega}_\phi$ since there is no PC in the normal state.

[13] We prove that under the assumption that P is strictly concave, $P(\nu)$ must be differentiable at interior $\nu > \alpha - E[\phi]$. Moreover, the FOCs are necessary. Although we cannot establish in general that they are also sufficient, this turns out to be the case in all parametric examples we considered.

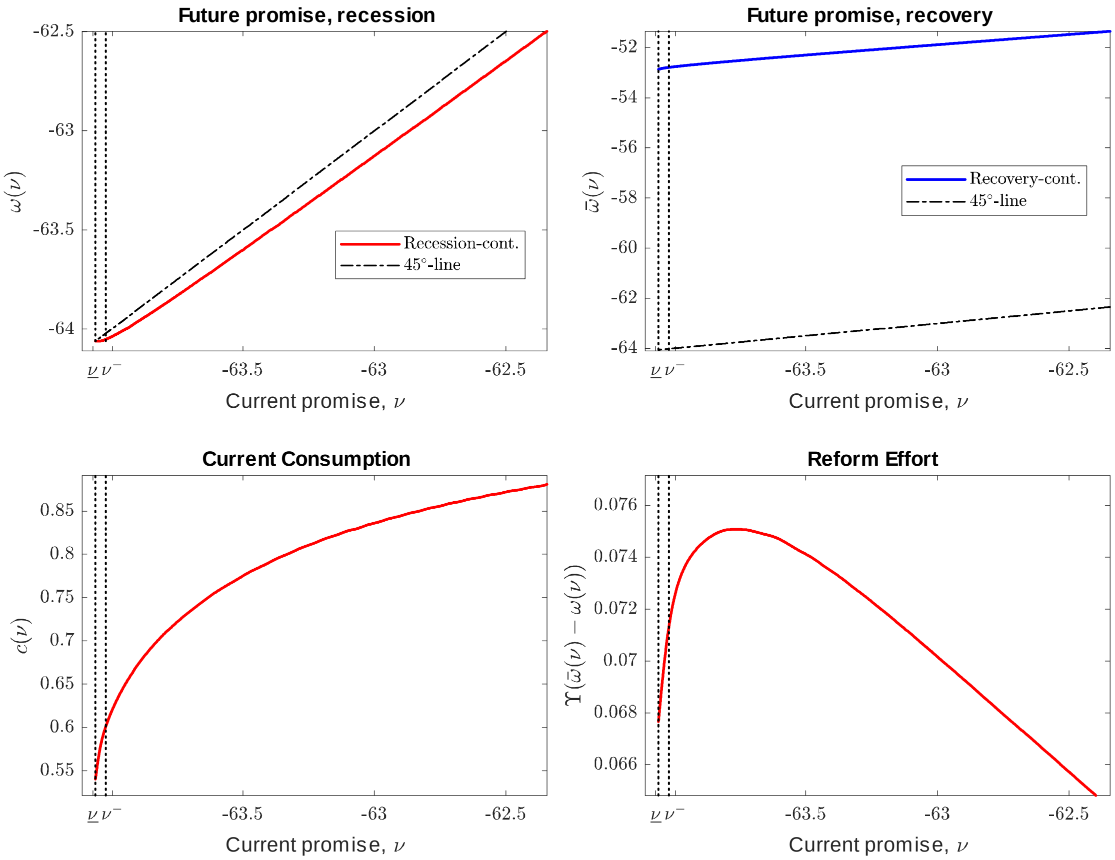

Figure 2: Policy functions for state-contingent promised utility, consumption, and effort conditional on the maximum cost realization $\phi_{\max}$.

Figure 2 summarizes the results of the proposition with the aid of a numerical example. All panels show policy functions for an economy in recession, conditional on a slack PC. The figure illustrates that promised utilities $\omega_\phi(\nu)$ and $\bar{\omega}_\phi(\nu)$ and consumption $c(\nu)$ are weakly increasing in ν. The upper left panel shows the law of motion of ν. The fact that the function $\omega_\phi(\nu)$ is below the 45-degree line implies that promised utility falls over time and converges to the lower bound $\underline{\nu} \equiv \alpha - E[\phi]$.[14] In the range $[\underline{\nu}, \nu^-]$ the planner is constrained by the inability to abase promised utility below $\underline{\nu}$ (more formally, ν^- is the lowest level of promised utility for which the constraint $\omega(\nu) \geq \underline{\nu}$ is slack). The upper right and lower left panels imply that, along a continuing recession in which the PC is slack, the sovereign's consumption declines in both aggregate states.

Combining (16) with the envelope condition $P'(\omega_\phi) = -1/u(c(\omega_\phi))$ for $\nu > \nu^-$ and denoting $\nu' = \omega_\phi$ and $\bar{\nu}' = \bar{\omega}_\phi$, yields:

$$\frac{1}{u'(c_\phi)} - \frac{1}{u'(c(\nu'))} = \frac{\Upsilon'(\bar{\nu}' - \nu')}{1 - \Upsilon(\bar{\nu}' - \nu')}\left(\bar{P}(\bar{\nu}') - P(\nu')\right). \tag{19}$$

We label equation (19) a Conditional Euler Equation (CEE). The CEE describes the optimal consumption dynamics for states where the PC does not bind next period and $\nu > \nu^-$. As noted above,

[14] Note that, as ν falls, the probability that the PC binds increases. At $\underline{\nu}$ the PC binds almost surely in the next period, and the sovereign receives the realized reservation utility $(\alpha - \phi')$ if the economy remains in recession.

the elasticity $\Upsilon'/(1-\Upsilon)$ captures moral hazard. In its absence, $\Upsilon' = 0$ and the planner delivers perfect consumption smoothing (see Proposition 2). Under moral hazard, the right hand-side of (19) is positive, implying that consumption decreases over time as long as the economy remains in recession and the PC is slack.[15]

Combining the CEE (19) with the Euler equation describing the consumption dynamics upon recovery (15), yields a conditional version of the so-called *Inverse Euler Equation* (CIEE):

$$\frac{1}{u'(c_\phi)} = \left(1 - \Upsilon\left(\bar{\nu}' - \nu'\right)\right)\frac{1}{u'(c(\nu'))} + \Upsilon\left(\bar{\nu}' - \nu'\right)\frac{1}{u'(\bar{c}(\bar{\nu}'))}. \tag{20}$$

The CIEE equates the inverse marginal utility in the current period with next period's expected inverse marginal utility conditional on the PC being slack and $\nu > \nu^-$. The key difference relative to the standard inverse Euler Equation in the dynamic contract literature (cf. Rogerson 1985) is that our CIEE only holds in states where the PC is slack. If there were no limited enforcement, then our CIEE would boil down to the standard result.

Next, consider the effort dynamics. Recall that, in the absence of moral hazard, effort and promised utility stay constant over time when the PC is slack. One might expect that, with moral hazard, the planner would back-load effort to incentivize its provision. However, this conjecture is wrong; effort is hump-shaped in promised utility (cf. Figure 2). In a range of low ν, effort is increasing in promised utility, implying that, as ν falls, effort decreases over time along the equilibrium path even when the PC is slack.

The reason for this non-monotonicity is subtle. Recall that, when the PC is slack, the optimum effort choice (14) yields $\tilde{p}_\phi(\nu) = \Upsilon(\bar{\omega}_\phi(\nu) - \omega_\phi(\nu))$. Since both $\bar{\omega}_\phi$ and ω_ϕ are increasing in ν, the effect of ν on $\tilde{p}_\phi$ is in general ambiguous. In the range $\nu \le \nu^-$, the constraint $\omega_\phi \ge \alpha - E[\phi]$ is binding, and the planner must set $\omega_\phi = \underline{\nu}$. Thus, $p_\phi(\nu) = \Upsilon(\bar{\omega}_\phi(\nu) - \underline{\nu})$, where both Υ and $\bar{\omega}_\phi(\nu)$ are increasing functions. Therefore, $\tilde{p}_\phi$ must be increasing in ν (hence, declining over time).[16] Conversely, effort is decreasing in promised utility when ν is high, in which case the planner uses the gap in promised utility across aggregate states to further incentivize effort provision.

The lower panels of Figure 1 illustrate simulated dynamics of consumption, promised utilities, and effort in the case of moral hazard under the same sequence of realizations of ϕ as in the upper panels. Moral hazard affects qualitatively the behavior of all variables when the PC is slack. Now, consumption and promised utilities fall over time, while effort falls or rises over time depending on ν. When ν becomes sufficiently low (i.e., in periods 7, 8, and 9), the reform effort starts falling. As the recession ends, consumption increases, implying that aggregate risk is not fully insured.

In conclusion, the combination of limited enforcement and moral hazard delivers effort dynamics qualitatively different from models with only one friction. In our model effort is hump-shaped over time (even when the PC remains slack), while effort is monotone increasing in pure moral hazard models (see for example Hopenhayn and Nicolini 1997) and weakly decreasing in pure limited enforcement models (see Proposition 2). The dynamics of consumption echo the typical properties of models with dynamic moral hazard as long as the PC is slack, namely the planner curtails consumption in order to extract higher effort over time. However, in our case the planner periodically increases consumption and promised utility whenever the PC is binding. This averts the immiseration that would arise in the absence of limited enforcement.

[15] The right-hand side of (19) is positive since $\Upsilon' > 0$ and $\bar{P}(\bar{\nu}') > P(\nu')$ (for the latter property, see proof of Proposition 3 in Appendix A) This implies that the marginal utility of consumption must be rising (and consumption falling) over time as ν falls.

[16] One can prove that, by continuity, this property carries over to a range of low ν's in a right neighborhood of $\underline{\nu}$.

2.3 Primal Formulation

The COA was characterized by solving a *dual* planning problem, i.e., maximizing profits subject to a promised-utility constraint. We can alternatively characterize the COA as the solution of a *primal* problem where the planner maximizes the sovereign's discounted utility subject to a promised-expected-profit constraint. The reason for laying out the primal program is that its formulation is directly comparable with the market equilibrium. Therefore, it will be useful in the derivation of the main decentralization result.

Let $\mu^{pl}(\pi)$ and $\mu_\phi^{pl}(\pi)$ denote the sovereign's discounted utility, respectively, before and after the realization of ϕ. We can write the primal problem as:

$$\begin{aligned}\mu^{pl}(\pi) &= \max_{\left\{c_\phi,p_\phi,\pi'_\phi,\bar{\pi}'_\phi\right\}_{\phi\in\aleph}} \int_\aleph \mu_\phi^{pl}(\pi)\, dF(\phi) \qquad (21)\\ &= \max_{\left\{c_\phi,p_\phi,\pi'_\phi,\bar{\pi}'_\phi\right\}_{\phi\in\aleph}} \int_\aleph \left[u(c_\phi) - X(p_\phi) + \beta\left((1-p_\phi)\,\mu^{pl}\left(\pi'_\phi\right) + p_\phi\bar{\mu}^{pl}\left(\bar{\pi}'_\phi\right)\right)\right] dF(\phi),\end{aligned}$$

where $\bar{\mu}^{pl}(\bar{\pi}) = u(\bar{w} - (1-\beta)\bar{\pi})/(1-\beta)$, subject to the promised-profit constraint

$$\pi = \int_\aleph b_\phi dF(\phi), \qquad (22)$$

having defined $b_\phi = \underline{w} - c_\phi + \beta\left((1-p_\phi)\,\pi'_\phi + p_\phi\bar{\pi}'_\phi\right)$ as the planner's discounted profit after the realization of ϕ. The problem is subject to a set of PCs and ICs

$$u(c_\phi) - X(p_\phi) + \beta\left((1-p_\phi)\,\mu^{pl}\left(\pi'_\phi\right) + p_\phi\bar{\mu}^{pl}\left(\bar{\pi}'_\phi\right)\right) \geq \alpha - \phi, \quad \phi \in \aleph, \qquad (23)$$

$$p_\phi = \arg\max_{p\in[\underline{p},\bar{p}]} -X(p) + \beta\left((1-p)\,\mu^{pl}\left(\pi'_\phi\right) + p\bar{\mu}^{pl}\left(\bar{\pi}'_\phi\right)\right), \qquad (24)$$

and to the boundary conditions $c_\phi \geq 0$ and $\pi'_\phi \leq \left(\Phi^{pl}\right)^{-1}(\phi_{\max})$ (where Φ^{pl} is defined below).

The primal allocation is identical to the dual one as long as $\pi = P(\nu)$. Conversely, the dual allocation is identical to the primal one as long as $\nu = \mu^{pl}(\pi)$. Thus, $\mu^{pl} = P^{-1}$. The analysis of the dual problem established that the COA features threshold properties, with a threshold function $\tilde{\phi}(\nu)$ decreasing in ν. The same property is inherited by the primal problem, where the threshold $\Phi^{pl}(\pi) \equiv \tilde{\phi}(P^{-1}(\pi))$ is an increasing function of π.

It is useful to precede the characterization result by some definitions.

Definition 1 *Let*

1. $\Psi^{pl}(\pi',\bar{\pi}') = \Upsilon(\bar{\mu}^{pl}(\bar{\pi}') - \mu^{pl}(\pi'))$*;*
2. $\left\langle \Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi), C^{pl}(\pi)\right\rangle = \arg\max_{c\geq 0,\pi',\bar{\pi}'} \left\{ \begin{array}{c} u(c) - X\left(\Psi^{pl}(\pi',\bar{\pi}')\right) \\ +\beta\left(\begin{array}{c}\Psi^{pl}(\pi',\bar{\pi}') \times \bar{\mu}^{pl}(\bar{\pi}') + \\ \left(1-\Psi^{pl}(\pi',\bar{\pi}')\right)\times\mu^{pl}(\pi')\end{array}\right)\end{array}\right\}$*, subject to* $\pi' \leq \left(\Phi^{pl}\right)^{-1}(\phi_{\max}) \equiv \pi_{\max}$ *and* $c = \underline{w} - b^{pl}(\pi) + \beta\left(\left(1-\Psi^{pl}(\pi',\bar{\pi}')\right)\pi' + \Psi^{pl}(\pi',\bar{\pi}')\,\bar{\pi}'\right)$*, where the function* $b^{pl} : [0,\pi_{\max}] \to R$ *is recursively defined as follows:*

$$b^{pl}(\pi) = \left(\frac{\pi - \int_{\phi_{\min}}^{\Phi^{pl}(\pi)} b^{pl}\left(\left(\Phi^{pl}\right)^{-1}(\phi)\right) dF(\phi)}{1 - F\left(\Phi^{pl}(\pi)\right)}\right). \qquad (25)$$

3. $W^{pl}(\pi) = u\left(C^{pl}(\pi)\right) - X\left(\Psi^{pl}\left(\Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi)\right)\right) + \beta \begin{pmatrix} \left(1-\Psi^{pl}\left(\Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi)\right)\right)\mu^{pl}\left(\Pi^{pl}(\pi)\right) \\ +\Psi^{pl}\left(\Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi)\right)\bar{\mu}^{pl}\left(\bar{\Pi}^{pl}(\pi)\right) \end{pmatrix}$.

Ψ^{pl} is the optimal incentive-compatible effort function. The functions in parts 2 and 3 of Definition 1 are all conditional on realizations of ϕ such that the PC is slack: Π^{pl} and $\bar{\Pi}^{pl}$ are the discounted profits under recession and normal state; b^{pl} denotes the transfer from the sovereign to the principal; finally, W^{pl} is the discounted utility conditional on realizations of ϕ such that the PC does not bind. Finally, note that $W^{pl}(\pi) = \alpha - \Phi^{pl}(\pi)$. We can now move to the primal characterization of the COA.

Proposition 4 *The planning problem has the following primal characterization. The planner sets* $c_\phi = C^{pl}\left(\wp^{pl}(\pi,\phi)\right)$, $\pi'_\phi = \Pi^{pl}\left(\wp^{pl}(\pi,\phi)\right)$, $\bar{\pi}'_\phi = \bar{\Pi}^{pl}\left(\wp^{pl}(\pi,\phi)\right)$, *and* $p_\phi = \Psi^{pl}(\Pi^{pl}\left(\wp^{pl}(\pi,\phi)\right), \bar{\Pi}^{pl}\left(\wp^{pl}(\pi,\phi)\right))$, *where*

$$\wp^{pl}(\pi,\phi) = \begin{cases} \hat{\pi}^{pl}_\phi & if \quad \phi < \Phi^{pl}(\pi) \\ \pi & if \quad \phi \geq \Phi^{pl}(\pi) \end{cases},$$

$\Phi^{pl}(\pi) \equiv \tilde{\phi}(P^{-1}(\pi))$, *and* $\hat{\pi}^{pl}_\phi = \left(\Phi^{pl}\right)^{-1}(\phi)$*. The sovereign attains the utility* $W^{pl}\left(\hat{\pi}^{pl}_\phi\right) = \alpha - \phi$ *if* $\phi < \Phi^{pl}(\pi)$ *(i.e., if the PC binds), and* $W^{pl}(\pi) = \alpha - \Phi^{pl}(\pi)$ *if* $\phi \geq \Phi^{pl}(\pi)$ *(i.e., if the PC is slack). The value function satisfies:*

$$\mu^{pl}(\pi) = \alpha - \left(1 - F\left(\Phi^{pl}(\pi)\right)\right)\Phi^{pl}(\pi) - \int_{\phi_{\min}}^{\Phi^{pl}(\pi)} \phi \, dF(\phi).$$

When the PC binds, consumption and promised future expected profits are pinned down by ϕ and are history-independent. The sovereign's discounted utility equals $\alpha - \phi$. Note that, in order to keep notation compact, we have defined a pseudo-expected profit $\hat{\pi}^{pl}_\phi$ corresponding to the lowest initial expected profit such that, conditional on the realized ϕ, the PC would have been slack. This change of variable allows us to pin down the optimal choices by simply applying the functions $\Pi^{pl}, \bar{\Pi}^{pl}$, and C^{pl} to $\hat{\pi}^{pl}_\phi$ instead of π. When the PC is slack, consumption and promised future expected profits are history-dependent, and the sovereign receives the utility $\alpha - \Phi^{pl}(\pi)$, irrespective of ϕ. The value function has therefore a very simple formulation.

As a first step towards a market decentralization of the COA, consider the following interpretation of the primal problem. In the initial period, the risk-neutral principal is endowed with a claim whose expected value is π. This claim is backed by a (possibly, negative) output transfer from the sovereign in the current period and by rolling over the claim to the next period. The transfer resembles the returns of a state-contingent bond. The principal receives the agreed return $b^{pl}(\pi)$ in all states in which the PC is slack. Otherwise, she takes a haircut $b^{pl}\left(\hat{\pi}^{pl}_\phi\right) < b^{pl}(\pi)$. In addition, the planner adjusts optimally the future claims (contingent on the aggregate state), as if the contract had undergone a renegotiation. The case without moral hazard is especially intuitive. If the PC is slack, the planner sets $\pi' = \pi$, or identically keeps the obligation constant at its initial level. If the PC binds, she sets $\pi' < \pi$, i.e., she reduces the future obligation so as to keep the sovereign in the contract. Under moral hazard, π changes over time even when the PC is slack in order to provide the optimal dynamic incentives for effort provision.

3 Market Equilibrium

In this section, we show that the COA can be decentralized by a market allocation where the sovereign issues one-period defaultable bonds and sells them to risk-neutral international creditors. In the market economy, the planner is replaced by a syndicate of international investors (the *creditors*) who can borrow and lend at the gross interest rate R. We assume that the sovereign can only issue one-period securities (GDP-linked bonds) whose return is contingent on the aggregate state. The first ($\bar{b}$) pays one unit of good if the economy, currently in recession, switches to a normal state. The second (b) pays one unit of good if the economy remains in recession. Effort is not verifiable, namely, there is no market for securities whose return is conditional on the effort level. Moreover, there is no market to insure against the realization of ϕ. Interestingly, despite the stark restrictions on the market structure, the Markov equilibrium decentralizes the COA.

We restrict attention to Markov-perfect equilibria, where agents condition their strategies on pay-off relevant state variables. In particular, we rule out reputational mechanisms. We view this assumption as realistic in the context of sovereign debt. The reason is that it is generally difficult for creditors to commit to punishment strategies, especially when new lenders can enter and make separate deals with the sovereign.

We label the two securities *recession-contingent debt* and *recovery-contingent debt* and denote their prices by $Q\left(b,\bar{b}\right)$ and $\bar{Q}\left(b,\bar{b}\right)$. At the beginning of each recession period, the sovereign observes the realization of the default cost ϕ and decides whether to honor the recession-contingent debt that reaches maturity or to announce default. When default is announced, a renegotiation protocol is triggered, described in detail below. Since debt is honored in normal times, no arbitrage implies that $\bar{Q} = pR^{-1}$. If the country could commit to pay its debt also in recession, we would have, similarly, that $Q = (1-p)R^{-1}$. However, due to the risk of renegotiation, recession-contingent debt sells at a discount, $Q \leq (1-p)R^{-1}$.

We now describe the renegotiation protocol. If the sovereign announces default, the syndicate of creditors can offer a take-it-or-leave-it haircut that we assume to be binding for all creditors.[17] By accepting this offer, the sovereign averts the default cost. In equilibrium, a haircut is offered only if the default threat is credible, i.e., if the realization of ϕ is sufficiently low to make the sovereign prefer default to full repayment.[18] Note that the creditors have, *ex-post*, all the bargaining power, and their offer makes the sovereign indifferent between an outright default and the proposed haircut.[19]

The timing of a debt crisis can be summarized as follows: The sovereign enters the period with the pledged debt b, observes the realization ϕ, and then decides whether to announce default on all its debt. If the default threat is credible, the creditors offer a haircut $\hat{b} \leq b$. Next, the country decides whether to accept or decline this offer. Then, the sovereign issues new debt subject to the period budget constraint $c = Q\left(b',\bar{b}'\right) \times b' + \bar{Q}\left(b',\bar{b}'\right) \times \bar{b}' + \underline{w} - \hat{b}$. To facilitate comparison with

[17] This is a strong assumption. Note that in our environment there would be no reason for a subset of creditors to deviate and seek a better deal. In reality, this issue may arise if deviants may hope that some courts would rule more favorably for them, as in the dispute involving the Argentinean government vs. a group of vulture funds led by Elliott Management. In Section 4 below, we show that ruling out renegotiation altogether reduces welfare, ex-ante. Therefore, our theory emphasizes the value of making haircut agreements binding for all creditors.

[18] By assumption, the sovereign has always the option to simply honor the debt contract. Thus, the creditors' take-it-or-leave offer cannot demand a repayment larger than the face value of outstanding debt.

[19] Our focus on renegotiable debt with a stochastic outside option is in line with the evidence of Sturzenegger and Zettelmeyer (2008) who document that even within a relatively short period (1998-2005) there are very large differences between average investor losses across different episodes of debt restructuring (see also Panizza *et al.* 2009 and Reinhart and Trebesch 2016).

the COA we start by assuming that the outside option following outright default is α. Thus, in the out-of-equilibrium event that the sovereign declines the offered haircut, the default cost ϕ is triggered, the debt is canceled, and realized utility is $\alpha - \phi$. We will later endogenize α by assuming that the sovereign can subsequently resume access to financial markets.

3.1 Markov Equilibrium: Definition

In Markov-perfect equilibria, the set of equilibrium functions depend only on the pay-off relevant state variables, b and ϕ. For technical reasons, we impose that debt is bounded, $b \in [\underline{b}, \tilde{b}]$ where $\underline{b} > -\infty$ and $\tilde{b} = \bar{w}/\left(1 - R^{-1}\right)$ is the natural borrowing constraint in normal times. In equilibrium, these bounds will never bind. We define the equilibrium for an economy starting in a recession with a recession-debt b. We omit the formal definition for an economy in normal time, since then the first-best allocation obtains.

Definition 2 *A Markov-perfect equilibrium is a set of value functions $\{V, W\}$, a threshold renegotiation function Φ, a set of equilibrium debt price functions $\{Q, \bar{Q}\}$, and a set of optimal decision rules $\{\wp, B, \bar{B}, C, \Psi\}$ such that, conditional on the state vector $(b, \phi) \in \left([\underline{b}, \tilde{b}] \times [\phi_{\min}, \phi_{\max}]\right)$, the sovereign maximizes utility, the creditors maximize profits, and markets clear. More formally:*

- *The value function V satisfies*

$$V(b, \phi) = \max\{W(b), \alpha - \phi\}, \tag{26}$$

where $W(b)$ is the value function conditional on the debt level b being honored,

$$W(b) = \max_{(b', \bar{b}') \in [\underline{b}, \tilde{b}]^2} u\left(Q\left(b', \bar{b}'\right) \times b' + \bar{Q}\left(b', \bar{b}'\right) \times \bar{b}' + \underline{w} - b\right) + Z\left(b', \bar{b}'\right),$$

continuation utility Z is defined as

$$Z\left(b', \bar{b}'\right) = \max_{p \in [\underline{p}, \bar{p}]} \left\{-X(p) + \beta\left(p \times \bar{\mu}\left(\bar{b}'\right) + (1-p) \times \mu\left(b'\right)\right)\right\}, \tag{27}$$

the value of entering normal times with debt $\bar{b}$ is

$$\bar{\mu}(\bar{b}) = u\left(\bar{w} - \left(1 - R^{-1}\right)\bar{b}\right) / (1 - \beta) \tag{28}$$

and $\mu(b) = \int_{\aleph} V(b, \phi)\, dF(\phi)$.

- *The threshold renegotiation function Φ satisfies*

$$\Phi(b) = \alpha - W(b). \tag{29}$$

- *The recovery and recession-contingent debt price functions satisfy the following arbitrage conditions:*

$$\bar{Q}\left(b', \bar{b}'\right) \times \bar{b}' = \Psi\left(b', \bar{b}'\right) R^{-1} \times \bar{b}' \tag{30}$$

$$Q\left(b', \bar{b}'\right) \times b' = \left[1 - \Psi\left(b', \bar{b}'\right)\right] R^{-1} \times \Pi\left(b'\right) \tag{31}$$

where $\Pi(b')$ *is the expected repayment of the recession-contingent bonds conditional on next period being a recession,*

$$\Pi(b) = (1 - F(\Phi(b)))\, b + \int_{\phi_{\min}}^{\Phi(b)} \hat{b}(\phi) \times dF(\phi), \tag{32}$$

and where $\hat{b}(\phi) = \Phi^{-1}(\phi)$ *is the new post-renegotiation debt after a realization* ϕ.

- *The set of optimal decision rules comprises:*

  1. *A take-it-or-leave-it debt renegotiation offer:*

  $$\wp(b, \phi) = \begin{cases} \hat{b}(\phi) & \text{if} \quad \phi \le \Phi(b), \\ b & \text{if} \quad \phi > \Phi(b). \end{cases} \tag{33}$$

  2. *An optimal debt accumulation and an associated consumption decision rule:*

  $$\begin{aligned} &\langle B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\rangle = \\ &\arg \max_{(b',\bar{b}')\in[\underline{b},\tilde{b}]^2} \left\{ u\left(Q\left(b',\bar{b}'\right) \times b' + \bar{Q}\left(b',\bar{b}'\right) \times \bar{b}' + \underline{w} - \wp(b,\phi)\right) + Z\left(b',\bar{b}'\right)\right\}, \end{aligned} \tag{34}$$

  $$\begin{aligned} C(\wp(b,\phi)) \quad = \quad & Q\left(B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\right) \times B(\wp(b,\phi)) + \\ & \bar{Q}\left(B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\right) \times \bar{B}(\wp(b,\phi)) + \underline{w} - \wp(b,\phi). \end{aligned} \tag{35}$$

  3. *An optimal effort decision rule:*

  $$\Psi\left(b',\bar{b}'\right) = \arg \max_{p\in[\underline{p},\bar{p}]} \left\{-X(p) + \beta\left(p \times \bar{\mu}\left(\bar{b}'\right) + (1-p) \times \mu\left(b'\right)\right)\right\}. \tag{36}$$

- *The equilibrium law of motion of debt is* $\left(b',\bar{b}'\right) = \langle B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\rangle$.

- *The probability that the recession ends is* $p = \Psi\left(b',\bar{b}'\right)$.

Equation (26) implies that there is renegotiation if and only if $\phi < \Phi(b)$. Since, *ex-post*, creditors have all the bargaining power, the discounted utility accruing to the sovereign equals the value that she would get under outright default. Thus

$$V(b,\phi) = W(\wp(b,\phi)) = \begin{cases} W(b) & \text{if} \quad b \le \hat{b}(\phi), \\ \alpha - \phi & \text{if} \quad b > \hat{b}(\phi). \end{cases}$$

Consider, next, the equilibrium price functions. Since creditors are risk neutral, the expected rate of return on each security must equal the risk-free rate of return. Then, the arbitrage conditions (30) and (31) ensure market clearing in the security markets and pin down the equilibrium price of securities in recession. The function $\Pi(b)$ defined in Equation (32) yields the market value of the outstanding debt b conditional on the aggregate state being a recession but before the realization of ϕ. The obligation b is honored with probability $1 - F(\Phi(b))$, where $\Phi(b)$ denotes the largest realization of ϕ such that the sovereign can credibly threaten to default. With probability $F(\Phi(b))$, debt is renegotiated to a level that depends on ϕ denoted by $\hat{b}(\phi)$. The haircut $\hat{b}(\phi)$ keeps the sovereign indifferent between accepting the creditors' offer and defaulting. Consequently, $W\left(\hat{b}(\phi)\right) = \alpha - \phi$.

Consider, finally, the set of decision rules. Equations (34)-(35) yield the optimal consumption-saving decisions while Equation (36) yields the optimal reform effort. The effort depends on b' and $\bar{b}'$, since it is chosen after the new securities are issued. Note also that since $F(\Phi(b)) = 0$ for $b' \ge (\Phi)^{-1}(\phi_{\max})$, the bond price function (32) implies that debt exceeding this level will not raise any debt revenue. Thus, it is optimal to choose $b' \le (\Phi)^{-1}(\phi_{\max})$.

3.2 Decentralization, Existence, and Uniqueness

We now establish a key result of the paper, namely, that the Markov-perfect equilibrium with one-period renegotiable bonds decentralizes the COA with unobservable effort. This result embeds an existence proof for the Markov-perfect equilibrium.

In the equilibrium allocation, the outstanding debt level b replaces the expected profit π in the primal planning problem as the endogenous state variable. The equilibrium debt price function identifies a one-to-one relationship between b and π through the function $\Pi(b)$, see (32) – recall that $\Pi(b)$ is the expected repayment of the recession-contingent bond before ϕ is realized. Since Π is an increasing function, we can invert it and write $b(\pi) = \Pi^{-1}(\pi)$, where $b(\pi)$ satisfies

$$b(\pi) = \frac{\pi - \int_{\phi_{\min}}^{\Phi(b(\pi))} \hat{b}(\phi)\, dF(\phi)}{1 - F(\Phi(b(\pi)))}, \quad \hat{b}(\phi) = W^{-1}(\alpha - \phi).$$

Note that $b(\pi)$ has the same form as $b^{pl}(x)$ in Definition 1. The counterpart in normal times is $\bar{b}(\bar{\pi}) = \bar{\pi}$, due to full commitment. The decentralization result will be stated under the condition that $\pi = \Pi(b)$ (or, identically, that $b = b(\pi)$), namely, the sovereign's initial obligation is the same in the COA and in the market equilibrium.

Proposition 5 *Suppose the sovereign's outside option is α. Then, there exists a Markov equilibrium that decentralizes the COA of Section 2.3. Namely, given equilibrium price functions $\{Q, \bar{Q}\}$ consistent with (30)–(31), the equilibrium policy functions for consumption and effort are the same as in the COA, $C(\wp(b(\pi), \phi)) = C^{pl}(\wp^{pl}(\pi, \phi))$, $\Psi(b(\pi), \bar{b}(\bar{\pi})) = \Psi^{pl}(\pi, \bar{\pi})$; the threshold functions for debt renegotiation is the same as the threshold for which the PC binds in the planning problem, $\Phi(b(\pi)) = \Phi^{pl}(\pi)$; the equilibrium law of motion of debt ($b' = B(\wp(b(\pi), \phi))$ and $\bar{b}' = \bar{B}(\wp(b(\pi), \phi))$) is consistent with the law of motion of promised profits in the COA ($\pi' = \Pi^{pl}(\pi)$ and $\bar{\pi}' = \bar{\Pi}^{pl}(\pi)$); and the sovereign earns the same discounted utility, $\mu(b(\pi)) = \mu^{pl}(\pi)$ and $\bar{\mu}(\bar{b}(\bar{\pi})) = \bar{\mu}^{pl}(\bar{\pi})$.*

The proof establishes that the the program solved by the sovereign and by the creditors in the competitive equilibrium (including market clearing) is the same as the primal planning problem of Section 2.3 (the strategy of establishing equivalence between the two programs is similar to Aguiar *et al.* 2017). The equilibrium characterization is therefore the same as the COA of Section 2.2.2 under the assumption that $\pi = \Pi(b) = P(\nu)$. Since both Π and P are monotonic functions, one can invert them and write $\nu = P^{-1}(\Pi(b))$. Note that the proposition does not establish the uniqueness of the equilibrium, only that there exists an equilibrium that decentralizes the COA.

The decentralization result hinges on the equilibrium price functions Q and $\bar{Q}$. In turn, these require that (i) there are two GDP-linked bonds; (ii) the bonds are renegotiable; (iii) renegotiation entails no cost; and (iv) the renegotiation protocol implies that creditors have full ex-post bargaining power.[20] Under these price functions, when she issues one-period bonds, the sovereign sequentially promises creditors an expected future profit that takes into account the probability of renegotiation. This is equivalent to the promise made by the social planner under full commitment. There are two interesting properties of this result. First, although there is a continuum of states of nature, two securities are sufficient to decentralize the COA. This is due to the state-contingency embedded in the renegotiable nature of bonds. Second, the Markov equilibrium cannot use a reputational mechanism,

[20] If any of these assumptions fail, the market equilibrium will in general be suboptimal. The assumption of full commitment under normal time is instead inessential, as we show in an earlier version of this paper.

while the planner has full commitment. Yet, the Markov equilibrium with one-period bonds provides the right dynamic incentives.

The equivalence result of Proposition 5 holds for any exogenous outside option α. Next, we endogenize α. In a sovereign debt setting it is natural to focus on a case where the sovereign can resume participation in financial markets after a default. For simplicity, we assume that access to new borrowing is immediate, in which case $\alpha = W(0)$ in the Markov equilibrium. This value for α offers also a natural interpretation of the planning problem, namely, that the sovereign can leave the optimal contract and revert to the market with zero debt after suffering the punishment ϕ.

With some abuse of notation let $W(b;\alpha)$ denote the value function conditional on honoring debt b in an economy with outside option α. Then, we look for an allocation satisfying the fixed point condition $\alpha = W(0;\alpha)$. The next Corollary shows that there exists one and only one such fixed point.

Corollary 1 *There exists a unique COA such that the outside option is equal to the value of starting with zero debt in the market equilibrium, $\alpha = W(0;\alpha)$.*

It is useful to note here that if the planner had access to a better technology to punish deviations (i.e., $\alpha < W(0)$), such as forcing the sovereign to autarky or imposing stronger exclusion restrictions, then the COA would yield higher utility than the market. We return to this point in Section 4 below.

Finally, we return briefly to the discussion about one- vs. two-sided commitment in the planning problem of Section 2.2.1. In footnote (7), we argued that commitment on behalf of the principal is not an issue as long as ν is sufficiently low. In the equilibrium, this amounts to assuming that $b \geq 0$. In this case, recession debt will always remain positive along the equilibrium path. Since this claim has non-negative value, the creditors would never unilaterally terminate existing contracts (hence, the principal would have no commitment problem in the planning program).

3.3 Debt Overhang and Debt Dynamics

The equivalence result in Proposition 5 implies that the Markov equilibrium inherits the same properties as the COA. A binding PC in the planning problem corresponds to an episode of sovereign debt renegotiation in the Markov equilibrium. Thus, as long as the recession continues and debt is not renegotiated, consumption falls. Debt dynamics are the mirror image of the dynamics of promised utility in the COA. A fall in promised utility corresponds to an increase of sovereign debt in the corresponding state. In the COA of Section 2.2.2, ν decreases over time during a recession unless the PC binds and the promised utilities $\omega(\nu)$ and $\bar{\omega}(\nu)$ are increasing functions (where, recall, $\nu' = \omega(\nu)$ if the recession continues and the PC is slack). Correspondingly, as long as debt is honored and the recession continues, both recession- and recovery-contingent debt are increasing over time.[21] When debt is renegotiated down, consumption and discounted utilities increase relative to the case in which the PC is slack.

The left panel of Figure 3 shows the equilibrium policy function for recession- and recovery-contingent debt (solid lines) conditional on no renegotiation and on the economy being in recession. Recession-contingent debt increases until $b^{\max}$. At this level, debt is renegotiated with certainty, and issuing more debt would not increase the expected repayment $\Pi(b)$. Note that $b^{\max}$ corresponds to the lower bound on promised utility $\underline{\nu}$ discussed in Section 2.2.2 and displayed in Figure 2. The figure

[21] This result hinges on moral hazard. In the planning problem without moral hazard, consumption and promised utilities are constant during recession when the PC is slack. Hence, both recession- and recovery-contingent debt would remain constant if effort were contractible.

also shows the level of debt b^{+} such that with this initial debt level the sovereign is constrained in that the optimal debt issuance is maximized, $b' = b^{\max}$ (cf. ν^{-} in Figure 2). Note that for any initial debt $b \geq b^{+}$ the sovereign would issue $b' = b^{\max}$, although newly issued recovery-contingent debt $\bar{b}'$ is increasing in b in this range.

The right panel of Figure 3 shows the equilibrium effort as function of b (solid line). This is the mirror image of the lower right panel in Figure 2: it is increasing at low debt levels and starts falling as debt reaches a sufficient high level. Intuitively, for low debt levels, a higher debt increases the desire for the sovereign to escape recession (this is the only force in the first best of Proposition 1 where, recall, effort is decreasing in promised utility). However, as debt increases, the probability that debt is fully honored in recession falls, increasing the share of benefits from recovery accruing to the creditors and making moral hazard more severe. In this region, there is a *debt overhang* problem. This is reminiscent of Krugman (1988), although in our paper it is not optimal for the creditors to forgive debt when the economy is in the debt overhang region. Moreover, in our model debt overhang is an equilibrium outcome, namely, a sequence of realizations may lead the sovereign to choose debt in this region and creditors to rationally buy it. Creditors are willing to buy recession-contingent debt even when the country issues $b^{\max}$ although they know that it will be renegotiated for sure. The reason is that renegotiation at uncertain terms will grant them a positive expected return equal to the market return R.

3.4 Contractible effort

To conclude the analysis of decentralization, we discuss a market arrangement that decentralizes the COA with observable effort of Section 2. This decentralization requires that reform effort be observable and verifiable, and that there exist a market for securities whose return is contingent on the reform effort. In particular, assume that there exists a security b_e that pays one unit of good in the next period if the reform effort is lower than the constrained efficient effort level denoted by $p^{*}(b)$.[22]. We label this security *effort-deviation debt* and denote its price by $Q_e\left(b', \bar{b}', b'_e\right)$. This security is also defaultable. In particular, if the sovereign fails to deliver $p^{*}(b)$ at time t, then at time t+1, after observing ϕ, the sovereign can either honor the debt b_e or announce default and trigger the usual renegotiation protocol.

Since along the equilibrium path the sovereign exerts the effort level p^{*}, the effort-deviation debt will be priced at $Q_e = 0$ in equilibrium. In the proposed equilibrium, the sovereign issues the maximum feasible effort-deviation debt $b_e = \tilde{b}$ and raises no additional revenue. After issuing b_e, the sovereign will not find it profitable to deviate and set $p < p^{*}$. To see why, consider one such deviation. Since, with a positive probability the economy would recover (even if effort were set to zero, $p_{\min} > 0$), then the effort-deviation debt would become due yielding an arbitrarily low consumption and expected utility.[23] Therefore, deviations are never profitable, and the allocation is equivalent to the COA of Proposition 2 in which the planner controls the effort. The assumption that reform effort is verifiable

[22] Here, we define $p^{*}(b) = p\left(P^{-1}(\Pi(b))\right)$. Recall that $p(\nu)$ denotes the constraint efficient effort level when effort is observable and that $\nu = P^{-1}(\Pi(b))$.

While one could in principle allow for a richer set of securities whose return is conditional on effort and aggregate state, the availability of such assets would be not affect the allocation.

[23] The assumption that there is full enforcement in normal times simplifies the argument but is not essential. The decentralization could alternatively be attained if the sovereign could issue two GDP-linked effort-deviation debt instruments. Then, the sovereign would issue the maximum sustainable recession-linked effort-deviation debt. The expected value of a deviation would in this case be $\alpha - E(\phi)$ which is a lower bound to the continuation value under no deviation (cf. Equation 4 in the COA). Therefore, the sovereign would prefer to exert the effort level p^{*}.

is very strong: it requires that international courts can accept and verify evidence about insufficient reform effort when ruling about the breech of contractual agreements. In reality, we do not see such contracts, arguably because the extent to which a country passes and, especially, enforces reforms is opaque.

4 Less Complete Markets

In this section, we consider a competitive (Markov) equilibrium subject to more severe market frictions: in the spirit of Eaton and Gersovitz (1981), the sovereign can issue only a one-period non-state-contingent bond. We first maintain the same renegotiation protocol as in the earlier sections; then, as an extension, we rule out renegotiation as in the original Eaton-Gersovitz model. It is fruitful to analyze this market environment because of its empirical appeal: in the real world, government bonds typically promise repayments that are independent of the aggregate state of the economy. As is common in the sovereign debt literature, we do not investigate the microfoundations of this market incompleteness. Rather, we take it as exogenous and study its effect on the allocation. All other assumptions remain unchanged.

The one-asset economy does not attain the COA. In the COA, the planner can optimally weigh the gains from risk sharing against the moral hazard issue. We proved that two defaultable securities are sufficient to replicate the COA. However, in the one-asset economy, the shortage of securities forces a particular correlation structure between future consumption in recovery and recession that is generally suboptimal, resulting in less risk sharing in equilibrium.

A formal definition, proof of existence, and characterization of equilibrium is deferred to the appendix (Propositions 6 and 7).[24] Here, we emphasize the salient features of the equilibrium. Consider, first, the equilibrium policy function for effort $\Psi(b)$. The first-order condition from (36) yields $X'(\Psi(b')) = \beta[\bar{\mu}(b') - \mu(b')]$. The function Ψ is falling in b for b sufficiently large. To see why, note that there exists a threshold debt $b^{EG} = \Phi^{-1}(\phi_{\max})$ such that debt is renegotiated with probability one if $b' \geq b^{EG}$ and the recession continues, while the debt is honored with positive probability if $b' < b^{EG}$. Recall that if the recession ends debt is repaid even for $b' \geq b^{EG}$. Thus, for $b' \geq b^{EG}$ the difference $\bar{\mu}(b') - \mu(b')$ is decreasing in b' implying a decreasing effort, i.e., debt overhang. This feature extends to a contiguous range of debt below b^{EG}. Conversely, it can be shown that $\Psi(b)$ is increasing for b sufficiently low. We conclude that the effort function is non-monotonic in debt thereby sharing the same qualitative properties as the benchmark allocation with two securities.

Consider, next, consumption dynamics. Note, first, that the risk of repudiation introduces some state contingency, since debt is repaid with different probabilities under recession and normal times. This provides a partial substitute for state-contingent contracts, although it does not restore the COA in the absence of GDP-linked securities (see e.g. Aguiar and Amador 2011). Recall that in the benchmark economy consumption was determined by two CEEs, (15) and (19). Intuitively, the planner could control the promised utility in each of the two states. In the decentralized equilibrium, the sovereign attained the same allocation through optimally issuing the two securities. This is not feasible in the one-asset economy: there is only one CEE, which pins down the *expected* marginal rate of substitution in consumption conditional on debt being honored. In the appendix (Proposition 7),

[24] The definition of the Markov equilibrium with non-state-contingent debt is parallel to Definition 2, except that the equilibrium policy function for effort is now a function $\Psi(b)$, where b denotes the non-contingent debt. Similarly, the price of this bond is denoted $Q(b)$.

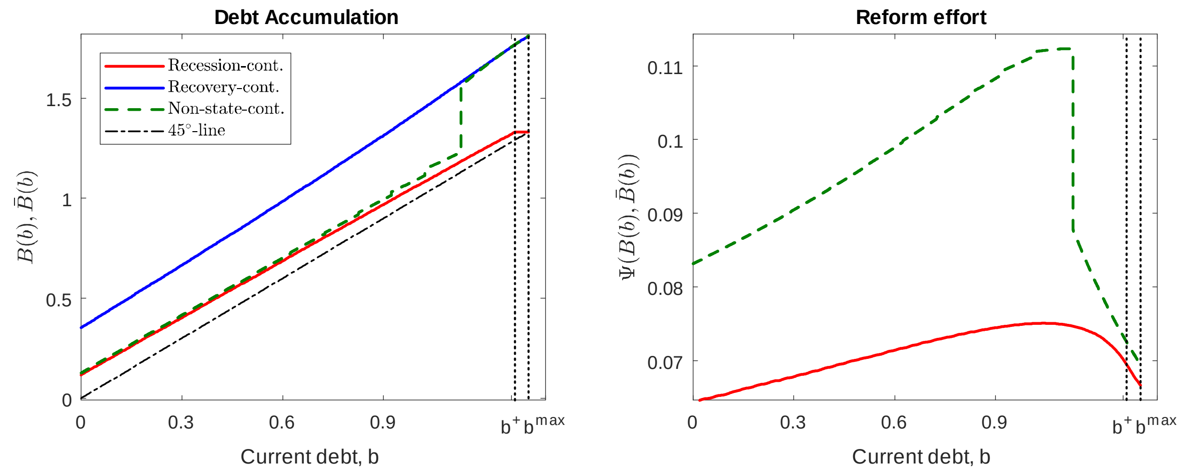

Figure 3: Policy functions conditional on the maximum cost realization $\phi_{\max}$. The solid lines show the policy functions for the constrained efficient Markov equilibrium. For comparison, the dashed line in each panel shows the policy functions of the one-asset economy.

we show that the CEE with non-state-contingent debt takes the form

$$E\left\{\frac{u'(c')}{u'(c)}|_{\text{debt is honored at t+1}}\right\} = 1 + \frac{\Psi'(b') \times [b' - \Pi(b')]}{\Pr(\text{debt is honored at t+1})} \tag{37}$$

where $b' - \Pi(b')$ is the difference between the expected debt repayment if the economy recovers or remain in recession.

Note, first, that in the one-asset economy consumption would increase at the end of the recession even if the probability that the recession ends were exogenous, i.e., $\Psi' = 0$. In this case, the CEE (37) would require that the expected marginal utility in the CEE be equal to the current marginal utility. For this to be true, consumption growth must be positive if the recession ends and negative if it continues. Thus, the missing market for GDP-linked securities curtails consumption smoothing even in the absence of moral hazard.

In the general case where $\Psi' \neq 0$, the market incompleteness interacts with the moral hazard friction introducing a novel strategic motive for debt. When the outstanding debt is low, then $\Psi' > 0$ and right-hand side of (37) is larger than unity. In this case, issuing more debt strengthens the *ex-post* incentive to exert reform effort. The CEE implies then higher debt accumulation and lower future consumption growth than in the absence of moral hazard. In contrast, in the region of debt overhang ($\Psi' < 0$) more debt strengthens the *ex-post* incentive to exert reform effort. To remedy this, the sovereign issues less debt than in the absence of moral hazard. Thus, when debt is large the moral hazard friction magnifies the reduction in consumption insurance.

Figure 3 shows the policy functions for debt and effort in the one-asset economy (dashed lines).[25] Debt is always higher than recession-debt in the two-security economy, indicating that there is less consumption smoothing. Effort is hump-shaped in both the one-asset economy and in the COA. Conditional on debt, effort is higher in the one-asset economy reflecting the fact that there is less

[25] Note that the policy rules for consumption and debt feature discontinuities. In the appendix, we prove that such discontinuities only arise in correspondence of debt levels that are never chosen in equilibrium (unless there is renegotiation). Moreover, the generalized Envelope Theorem in Clausen and Strub (2016) ensures that the FOCs are necessary despite the fact that the decision rules are discontinuous (see Proposition 7).

insurance against the continuation of a recession, making the sovereign more eager to leave the recession itself.

4.1 No Renegotiation

Next, we consider the effect of shutting down renegotiation. Namely, we assume that there are no possibilities to avert default when the sovereign finds it optimal not to honor the debt (like in Eaton and Gersovitz 1981). This alternative environment has three implications. First and foremost, there is costly default in equilibrium. The real costs suffered by the sovereign yields no benefit to creditors. This is in contrast to the equilibrium with renegotiation, where no real costs accrue and creditors recover a share of the face value of debt. Second, renegotiation affects the price function of debt, and thus the incentive for the sovereign to accumulate debt. In particular, the bond price now becomes $Q^{NR} = \Psi^{NR}(b)/R + \left[1 - \Psi^{NR}(b)\right]\left(1 - F\left(\Phi^{NR}(b')\right)\right) b'$. When renegotiation is ruled out the lender expects a lower repayment. Thus, the risk premium associated with each debt level is higher, and it becomes more costly for the sovereign to roll over debt. Therefore, the sovereign will be more wary of debt accumulation. This limits consumption smoothing and lowers welfare. Third, conditional on the debt level, the range of ϕ for which debt is honored is different across the two economies.[26] While this is in general ambiguous, our numerical analysis suggests that ruling out renegotiation reduces the likelihood that a given level of debt is honored.

Figure 4 in the appendix compares the policy functions in an Eaton-Gersovitz economy without renegotiation (solid lines) with a one-asset economy with renegotiation (dashed lines). Ruling out renegotiation implies lower consumption for each debt level than in the economy where debt can be renegotiated. Effort is larger when renegotiation is ruled out, reflecting the fact that the debt price is lower in recession and that there is less insurance. Finally, the probability of full debt repayment is lower without renegotiation than when debt can be renegotiated.

4.2 Assistance Program

The possibility of market failures discussed in the previous section provides scope for policy intervention. Consider an assistance program conducted by an international institution, e.g., the IMF. The assistance program mimics the COA through a sequence of transfers to the sovereign during recession in exchange of the promise of a repayment once the recession is over. The IMF has full commitment, but (like the planner) has access to limited instruments to punish deviations: it can inflict the same stochastic punishment (ϕ) as can markets.

The program has two key features. First, the terms of the program are renegotiable: whenever the country credibly threatens to abandon it, the IMF sweetens the deal by increasing the transfers and reducing the payment the country owes when the recession ends. When there is no credible default threat, the transfer falls over time, implying the constrained optimum sequence of declining consumption and time-varying reform effort prescribed by the COA. Second, when the recession ends, the IMF receives a payment from the sovereign, financed by issuing debt in the market. This payment depends on the length of the recession and on the history of renegotiations.

More formally, let ν^o denote the discounted utility guaranteed to the sovereign when the program is first agreed upon. At the beginning of that period, the IMF buys the debt b_0 with an expenditure

[26] More formally, in the equilibrium with renegotiation the sovereign renegotiates if $\phi < W(0) - W(b)$ whereas in the no-renegotiation equilibrium she defaults if $\phi < W^{NR}(0) - W^{NR}(b)$, where W^{NR} is the value function under no renegotiation and recession. As long as W^{NR} is falling more steeply in b than W, then, conditional on the debt level, the sovereign is more likely to honor the debt in the benchmark equilibrium than in the no-renegotiation equilibrium.

$\Pi(b_0)$.[27] Then, the IMF transfers to the country $T(\nu^o) = c(\nu^o) - \underline{w}$ where $c(\nu^o)$ is as in the COA of Proposition 3 unless the realization of ϕ makes the sovereign want to terminate the program, in which case the country receives $T_\phi = c_\phi - \underline{w}$, where, again, c_ϕ is as in Proposition 3. In the subsequent periods, consumption and promised utility evolve in accordance with the law of motion of the COA. In other words, the IMF commits to a sequence of state-contingent transfers that mimics the COA in the planning allocation. Note that this implies, by construction, that the sovereign exerts the incentive-compatible reform effort level. As soon as the recession ends, the country owes a debt b_N to the international institution, determined by the equation $R^{-1} \times \bar{b}_N = \bar{c}(\bar{\omega}_N) - \bar{w} + \bar{b}_N$, where the discounted utility $\bar{\omega}_N$ depends on the duration of the recession and on the history of realizations of ϕ. Note that $\bar{c}$ is first-best consumption, exactly like in the COA, and that in normal time the country resumes to markets to refinance its debt at the constant level b_N.

How large ν^o can be depends on how many resources the IMF is willing to commit to sustain the assistance program. A natural benchmark is to pin down ν^o at the level that allows the IMF to break even in expectation. Whether, ex-post, the IMF makes net gains or losses hinges on the duration of the recession and on the realized sequence of ϕ's.

Can such a program improve upon the market allocations? The answer hinges on the extent of market and contract incompleteness. If there exists a market for GDP-linked bonds and if the renegotiation process is frictionless and efficient, the assistance program cannot improve upon the market allocation. More formally, $\nu^o = P^{-1}(\Pi(b_0))$. This follows immediately from our decentralization result in Section 3. However, in the one asset economies (with or without renegotiation), the assistance program yields higher efficiency than the competitive equilibrium. Interestingly, the assistance program would in this case yield higher utility than the equilibrium with GDP-linked securities because the market incompleteness provides the IMF with a more powerful threat to discipline the sovereign's behavior by making deviations less attractive. The more pervasive the market incompleteness, the closer is the assistance program to first best. Formally, the outside option $\alpha = W(0)$ is lower the more incomplete are markets.

The assistance program would be even more powerful if the IMF could observe and verify the reform effort (e.g., by taking temporarily control over the reform process). In this case, the policy intervention could implement the COA without moral hazard of Section 2.2.1 acting as a stand-in for the missing market of effort-deviation debt discussed in Section 3.4. Clearly, policies infringing on country's sovereignty run into severe politico-economic implementability constraints whose discussion goes beyond the scope of our paper.

4.3 Welfare Effects

In a previous version of the paper (Müller *et al.* 2015), we quantified the effects of the different informational frictions and of assistance programs when markets are incomplete. The quantitative analysis required some generalization of the stylized theory in order to align the model with the data. In particular, we relaxed the assumption that there exists an absorbing state and assumed, instead, that in normal times the economy falls into a recession with an exogenous probability. Additionally, we relaxed the assumption that $\beta R = 1$, and emphasize the case in which $\beta R < 1$ since this ensures that the competitive equilibrium has a non-degenerate stationary distribution. Then, we calibrated the model economy to match salient moments of observed debt-to-GDP ratios and default premia for Southern European countries.

[27] Π depends on the market structure. If there are two securities, then Π is given by equation (32). If there is only one asset, the corresponding definition given in the appendix applies.

The analysis showed that the calibrated model is successful in replicating a number of moments in the data, including key non-calibrated moments such as the bond spread, the frequency of renegotiations, the average haircut of the debt's face value and its variance across renegotiation episodes. We found that an assistance program can deliver sizeable welfare gains, especially if it can alleviate the moral hazard problem. The welfare effects of completing markets in a one asset economy are significant when renegotiation is ruled out in the market economy. The details are available upon request.

5 Concluding Remarks

This paper proposes a novel theory of sovereign debt dynamics under limited enforcement and moral hazard. A sovereign country issues debt to smooth consumption during a recession whose duration is uncertain and endogenous. The expected duration of the recession depends on the intensity of (costly) structural reforms. Both elements – the risk of repudiation and the need for structural reforms – are salient features of the European debt crisis during the last decade.

A key result is that a Markovian competitive equilibrium with renegotiable one-period GDP-linked securities implements the constrained optimum allocation. The crux of this result is that, under the assumption that creditors have all the ex-post bargaining power, the renegotiable securities are a stand-in for missing markets for state-contingent debt. In addition, these markets provide the optimal trade-off between insurance and incentives that a fully committed planner subject only to informational constraints can achieve. A surprising element is that the market needs no reputational mechanism to discipline the sovereign's effort provision over time. In fact, it attains the constrained optimum with very "simple" instruments, i.e., two one-period securities.

We also study the effect of additional exogenous restrictions on market arrangements, including assuming that the sovereign can only issue non-contingent debt and, in the spirit of Eaton-Gersovitz (1981), ruling out renegotiation. In this case, the market equilibrium is not constrained efficient. We discuss an assistant program that can restore efficiency and the associated welfare gains.

To retain tractability, we make important assumptions that we plan to relax in future research. First, in our theory the default cost follows an exogenous stochastic process. In a richer model, this would be part of the equilibrium dynamics. Strategic delegation is a potentially important extension. Voters may have an incentive to elect a government that undervalues the cost of default. In our current model, however, the stochastic process governing the creditor's outside option is exogenous, and is outside of the control of the sovereign and creditors.

Second, again for simplicity, we assume that renegotiation is costless, that creditors can perfectly coordinate and that they have full bargaining power in the renegotiation game. Each of these assumptions could be relaxed. For instance, in reality the process of negotiation may entail costs. Moreover, as in the recent contention between Argentina and the so-called vulture funds, some creditors may hold out and refuse to accept a restructuring plan signed by a syndicate of lenders. Finally, the country may retain some bargaining power in the renegotiation. All these extensions would introduce interesting additional dimensions, and invalidate some of the strong efficiency results (for instance, the result that the market economy attains the constrained optimum in the absence of income fluctuations). However, we are confident that the gist of the results is robust to these extensions.

Finally, by focusing on a representative agent, we abstract from conflicts of interest between different groups of agents within the country. Studying the political economy of sovereign debt would be an interesting extension. We leave the exploration of these and other avenues to future work.

References

Abraham, Arpad, Eva Carceles-Poveda, Yan Liu, and Ramon Marimon (2017). "On the optimal design of a financial stability fund," Mimeo, European University Institute.

Abraham, Arpad, and Nicola Pavoni (2008). "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation," *Review of Economic Dynamics* 11(4), 781-803.

Aguiar, Mark, and Manuel Amador (2011). "Growth in the Shadow of Expropriation," *Quarterly Journal of Economics* 126(2), 651-697.

Aguiar, Mark, Manuel Amador, Hugo Hopenhayn, and Iwan Werning (2017). "Take the Short Route: Equilibrium Default and Maturity," Mimeo, Princeton University.

Aguiar, Mark, and Manuel Amador (2014). "Sovereign Debt: a Review," *Handbook of International Economics* 4, 647-87.

Aguiar, Mark, and Gita Gopinath (2006). "Defaultable Debt, Interest Rates, and the Current Account," *Journal of International Economics* 69 (1), 64–83.

Aguiar, Mark, Manuel Amador, and Gita Gopinath (2009). "Investment Cycles and Sovereign Debt Overhang," *Review of Economic Studies*, 76(1), 1-31.

Alvarez, Fernando, and Urban J. Jermann (2000). "Efficiency, equilibrium, and asset pricing with risk of default," *Econometrica* 68(4), 775–797.

Arellano, Cristina (2008). "Default risk and income fluctuations in emerging economies," *American Economic Review* 98(3), 690–712.

Arellano, Cristina, Lilia Maliar, Serguei Maliar, and Viktor Tsyrennikov (2014). "Envelope Condition Method with an Application to Default Risk Models," Mimeo, http://dx.doi.org/10.2139/ssrn.2470009.

Asonuma, Tamon, and Christoph Trebesch (2016). "Sovereign debt restructurings: preemptive or post-default," *Journal of the European Economic Association* 14(1) 175–214.

Atkeson, Andrew (1991). "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica* 59(4), 1069-1089.

Blanchard, Olivier, and Francesco Giavazzi (2003). "Macroeconomic Effects of Regulation and Deregulation in Goods and Labor Markets," *Quarterly Journal of Economics* 118(3), 879–907.

Bolton, Patrick, and Olivier Jeanne (2007). "Structuring and restructuring sovereign debt: the role of a bankruptcy regime," *Journal of Political Economy* 115(6), 901–924.

Broner, Fernando A., Alberto Martin, and Jaume Ventura (2010). "Sovereign risk and secondary markets," *American Economic Review* 100(4), 1523-1555.

Bulow, Jeremy, and Kenneth Rogoff (1989). "A constant recontracting model of sovereign debt," *Journal of Political Economy* 97(1), 155–178.

Chatterjee, Satyajit and Burcu Eyigungor, (2012). "Maturity, Indebtedness, and Default Risk," American Economic Review, 102(6), 2674-2699.

Clausen, Andrew, and Carlo Strub (2016). "A General and Intuitive Envelope Theorem," Mimeo. University of Edinburgh, URL:https://andrewclausen.net/Clausen_Strub_Envelope.pdf.
Conesa, Juan Carlos, and Timothy J. Kehoe (2015). "Gambling for redemption and self-fulfilling debt crises," Research Department Staff Report 465, Federal Reserve Bank of Minneapolis.

Dovis, Alessandro (2017). "Efficient Sovereign Default," Mimeo, University of Pennsylvania.

Eaton, Jonathan, and Mark Gersovitz (1981). "Debt with potential repudiation: Theoretical and empirical analysis," *Review of Economic Studies* 48(2), 289–309.

Hatchondo, Juan Carlos, Leonardo Martinez, and César Sosa Padilla (2014). "Voluntary sovereign debt exchanges," *Journal of Monetary Economics* 61 32–50.

Hopenhayn, Hugo A., and Juan Pablo Nicolini (1997). "Optimal Unemployment Insurance," *Journal of Political Economy* 105(2), 412-438.

Ilzkovitz Fabienne, and Adriaan Dierx (2011). "Structural Reforms: A European Perspective," Reflets et perspectives de la vie économique 3/2011 (Tome L), 13–26.

Jeanne, Olivier (2009). "Debt Maturity and the International Financial Architecture," *American Economic Review* 99(5), 2135–2148.

Kehoe, Patrick J., and Fabrizio Perri (2002). "International business cycles with endogenous incomplete markets," *Econometrica* 70(3), 907–928.

Kocherlakota, Narayana R. (1996). "Implications of efficient risk sharing without commitment," *Review of Economic Studies* 63(4), 595–609.

Krueger, Dirk, and Harald Uhlig (2006). "Competitive risk sharing contracts with one-sided commitment," *Journal of Monetary Economics* 53(7), 1661–1691.

Krugman, Paul (1988). "Financing vs. forgiving a debt overhang," *Journal of Development Economics* 29(3), 253–268.

Ljungqvist, Lars, and Thomas J. Sargent (2012). *Recursive Macroeconomic Theory*, Third Edition, Cambridge, MA: MIT Press.

Marcet, Albert, and Ramon Marimon (1992). "Communication, commitment, and growth," *Journal of Economic Theory* 58(2), 219-249.

Mendoza, Enrique G., and Vivian Z. Yue (2012). "A General Equilibrium Model of Sovereign Default and Business Cycles," *Quarterly Journal of Economics* 127(2), 889–946.

Müller, Andreas, Kjetil Storesletten, and Fabrizio Zilibotti (2016). "The Political Color of Fiscal Responsibility," *Journal of the European Economic Association* 14(1), 252–302.

Panizza, Ugo, Federico Sturzenegger, and Jeromin Zettelmeyer (2009). "The economics and law of sovereign debt and default," *Journal of Economic Literature* 47(3), 651–698.

Phelan, Christopher (1995). "Repeated Moral Hazard and One-Sided Commitment," *Journal of Economic Theory* 66(2), 488-506.

Reinhart, Carmen M., and Christoph Trebesch (2016). "Sovereign debt relief and its aftermath," *Journal of the European Economic Association* 14(1), 215–251.

Rogerson, William P. (1985), "Repeated Moral Hazard," *Econometrica*, 53(1), 69–76.

Song, Zheng, Kjetil Storesletten, and Fabrizio Zilibotti (2012). "Rotten Parents and Disciplined Children: A Politico-Economic Theory of Public Expenditure and Debt," *Econometrica* 80(6), 2785–2803.

Spear, Stephen E., and Sanjay Srivastava (1987). "On repeated moral hazard with discounting," *Review of Economic Studies* 54(4), 599–617.

Sturzenegger, Federico, and Jeromin Zettelmeyer (2008). "Haircuts: estimating investor losses in sovereign debt restructurings, 1998–2005," *Journal of International Money and Finance* 27(5), 780–805.

Thomas, Jonathan, and Tim Worrall (1988). "Self-enforcing wage contracts," *Review of Economic Studies* 55(4), 541–554.

Thomas, Jonathan, and Tim Worrall (1990). "Income fluctuation and asymmetric information: An example of a repeated pincipal-agent model," *Journal of Economic Theory* 51(2), 367–390.

Tomz, Michael, and Mark L.J. Wright (2007). "Do countries default in bad times?," *Journal of the European Economic Association* 5(2-3), 352–360.

Yue, Vivian Z. (2010). "Sovereign default and debt renegotiation," *Journal of International Economics* 80(2), 176–187.

Online Appendixes of "Sovereign Debt and Structural Reforms"

Andreas Müller, Kjetil Storesletten and Fabrizio Zilibotti

A Appendix A: Proofs of main lemmas, propositions and corollaries

Proof of Lemma 1. To prove this result, we ignore technicalities related to the continuum of states. Namely, we assume that there exist N states with associated positive probabilities, and view the continuum as an approximation of the discrete state space for $N \to \infty$. This is without loss of generality (one can replace statements about single states ϕ by statements about small intervals of positive measure). The proof involves two steps.

1. We start by proving that, if $\phi_2 > \phi_1$, then, $\mu_{\phi_1} > \alpha - \phi_1 \Rightarrow \mu_{\phi_2} = \mu_{\phi_1} > \alpha - \phi_2$. To see why, recall that the planner's objective is to deliver the promised utility ν given by (2) in a profit-maximizing way. Since u, P, and $\bar{P}$ are strictly concave, and X is strictly convex, then, profit maximization is attained by setting $c_{\phi_2} = c_{\phi_1}$, $p_{\phi_2} = p_{\phi_1}$, $\omega_{\phi_2} = \omega_{\phi_1}$, and $\bar{\omega}_{\phi_2} = \bar{\omega}_{\phi_1}$. This planning choice implies that $\mu_{\phi_2}(\nu) = \mu_{\phi_1}(\nu)$. This is feasible since the assumption that the PC is slack in state ϕ_1 implies that *a fortiori* the PC is slack in state ϕ_2. Two cases are then possible: either $\exists \tilde{\phi}$ such that $\mu_{\phi}(\nu) = \alpha - \tilde{\phi}(\nu)$ for all $\phi \geq \tilde{\phi}(\nu)$; or $\mu_{\phi}(\nu) = \mu(\nu) > \alpha - \phi_{\min}$ $\forall \phi \in [\phi_{\min}, \phi_{\max}]$. The latter case would imply that no PC ever binds and can be ignored.

2. Next, we prove that, if $\hat{\phi}_2 > \hat{\phi}_1$, then, $\mu_{\hat{\phi}_2} = \alpha - \hat{\phi}_2$ implies that $\mu_{\hat{\phi}_1} = \alpha - \hat{\phi}_1$. To derive a contradiction, suppose that this is not the case, and that $\mu_{\hat{\phi}_1} > \alpha - \hat{\phi}_1 > \mu_{\hat{\phi}_2}$ (the opposite inequality would violate the PC and is not feasible). Then, the planner could deliver to the agent the promised utility ν by uniformly increasing ω_2, and $\bar{\omega}_2$ and reducing ω_1, and $\bar{\omega}_1$ also uniformly, so as to keep ν unchanged, while leaving consumption and effort constant (it is easy to check that this is feasible). The strict concavity of P and $\bar{P}$ guarantees that this change increases profits. Let $\tilde{\phi}(\nu)$ denote the largest ϕ such that $\mu_{\tilde{\phi}(\nu)}(\nu) = \alpha - \tilde{\phi}(\nu)$. Then, $\mu_{\phi}(\nu) = \alpha - \phi$ for all $\phi \leq \tilde{\phi}(\nu)$.

Parts 1. and 2. above jointly establish that the threshold $\tilde{\phi}(\nu)$ is unique. ■

Proof of Proposition 3. The proof that P is differentiable at the interior of its support is provided in Appendix B (Lemma 2).

For the characterization result, consider the following program. The planner solves (1) subject to

(2)-(5), and (14). The Lagrangian yields

$$\begin{aligned}\mathcal{L} &= \int_{\aleph} \left[\underline{w} - c_\phi + \beta\left((1-p_\phi) P(\omega_\phi) + p_\phi \bar{P}(\bar{\omega}_\phi)\right)\right] f(\phi) d\phi \\ &+\vartheta\left(\int_{\aleph} \left(u(c_\phi) - X(p_\phi) + \beta\left((1-p_\phi)\omega_\phi + p_\phi\bar{\omega}_\phi\right)\right) f(\phi) d\phi - \nu\right) \\ &+\int_{\aleph} \lambda_\phi \left(u(c_\phi) - X(p_\phi) + \beta\left((1-p_\phi)\omega_\phi + p_\phi\bar{\omega}_\phi\right) - [\alpha-\phi]\right) d\phi \\ &+\int_{\aleph} \chi_\phi\left(-X'(p_\phi) + \beta(\bar{\omega}_\phi - \omega_\phi)\right) d\phi + \int_{\aleph} \theta_\phi\left(\omega_\phi - (\alpha - E[\phi])\right) d\phi, \end{aligned} \tag{38}$$

where $\vartheta \geq 0$, $\lambda_\phi \geq 0$, χ_ϕ, and $\theta_\phi \geq 0$ denote the multipliers. The first-order conditions (FOCs) with respect to c_ϕ, $\bar{\omega}_\phi$, ω_ϕ, p_ϕ and χ_ϕ yield

$$\begin{aligned} 0 &= -f(\phi) + [\vartheta f(\phi) + \lambda_\phi] u'(c_\phi), \\ 0 &= \beta p_\phi \bar{P}'(\bar{\omega}_\phi) f(\phi) + [\vartheta f(\phi) + \lambda_\phi]\beta p_\phi + \chi_\phi \beta \\ 0 &= \beta(1-p_\phi)P'(\omega_\phi) f(\phi) + [\vartheta f(\phi) + \lambda_\phi]\beta(1-p_\phi) - \chi_\phi\beta, \quad \forall \omega_\phi > \alpha - E[\phi] \\ 0 &= \beta\left[\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi)\right] f(\phi) + [\vartheta f(\phi) + \lambda_\phi]\left(-X'(p_\phi) + \beta(\bar{\omega}_\phi - \omega_\phi)\right) - \chi_\phi X''(p_\phi) \\ 0 &= -X'(p_\phi) + \beta(\bar{\omega}_\phi - \omega_\phi) \end{aligned}$$

Lemma 3 in Appendix B rules out the possibility of corner solutions establishing that the FOCs are necessary optimality conditions. The envelope condition yields $-P'(\nu) = \vartheta$ for all $\nu > \alpha - E[\phi]$ (note that P is only differentiable at the interior support of ν), while the slackness condition for θ_ϕ implies $0 = \theta_\phi(\omega_\phi - (\alpha - E[\phi]))$. Combining the FOCs and the envelope condition, and noting that $-\bar{P}'(\bar{\omega}_\phi) = 1/u'(\bar{c}(\bar{\omega}_\phi))$, yields

$$\frac{1}{u'(c_\phi)} = -P'(\nu) + \frac{\lambda_\phi}{f(\phi)}, \quad \forall \nu > \alpha - E[\phi], \tag{39}$$

and Equations (15) and (16) in the text. Note that we have also used the facts that $p_\phi = \Upsilon(\bar{\omega}_\phi - \omega_\phi)$ (established in the text, see Equation (14)) and $\Upsilon'(\bar{\omega}_\phi - \omega_\phi) = \beta/X''(\Upsilon(\bar{\omega}_\phi - \omega_\phi))$.

In order to prove $\tilde{\phi}(\nu)$ is strictly decreasing, note that Lemma 1 establishes that all PCs associated with $\phi < \tilde{\phi}(\nu)$ are binding, while all those with $\phi \geq \tilde{\phi}(\nu)$ are slack. Also, note that, since $P'(\nu)$ is strictly decreasing, the PK always binds ($\vartheta > 0$). Hence, the PK can be written as $\nu = \left(1 - F\left(\tilde{\phi}(\nu)\right)\right)\left(\alpha - \tilde{\phi}(\nu)\right) + \int_{\tilde{\phi}(\nu)}^{\phi_{\max}} (\alpha - \phi) dF(\phi)$, which can rearranged to yield equation (13). Standard arguments establish then that $\tilde{\phi}(\nu)$ is strictly decreasing.

Consider part 1. in the proposition. For all $\phi < \tilde{\phi}(\nu)$, the PC is binding and holds with equality. In these cases, the optimal choices are independent of ν and c_ϕ, ω_ϕ, $\bar{\omega}_\phi$, and θ_ϕ are determined by the PC (5) holding with equality, and by the FOCs (15)–(17). Consider, next, part 2. For all $\phi \geq \tilde{\phi}(\nu)$, the PC is slack, and Lemma 1 implies Equation (18). The solution is history dependent: $c_\phi = c(\nu)$, $\omega_\phi = \omega(\nu)$, $\bar{\omega}_\phi = \bar{\omega}(\nu)$, $\theta_\phi = \theta(\nu)$, where $c(\nu)$, $\omega(\nu)$, $\bar{\omega}(\nu)$, and $\theta(\nu)$ are determined by Equations (15)–(18). To prove that $c(\nu)$ is an increasing function, note that, since P and u are strictly concave, then (39) implies that $c(\nu)$ is strictly increasing for $\nu > \alpha - E[\phi]$. Moreover, consumption is also falling at the lower bound, $c(\underline{\nu}) < c(\nu)\ \forall \nu > \underline{\nu}$. This can be proved by a contradiction argument:

Suppose that $c(\underline{\nu}) \geq c(\nu)$. Then, the FOCs imply that the promise-keeping constraint is slack, which is impossible.

Next, we prove that promised utility falls over time, i.e., $\omega_\phi = \omega(\nu) \leq \nu$, with $\omega(\nu) < \nu$ for $\theta_\phi = 0$ (i.e., if ν is sufficiently large). To this aim, we establish the following properties in three steps: (i) $\bar{P}(\bar{\omega}_\phi(\nu)) - P(\omega_\phi(\nu)) > 0$; (ii) if $\nu > \underline{\nu}$, then $\bar{c}(\bar{\omega}_\phi(\nu)) > c_\phi(\nu) \geq c(\nu) > c(\omega(\nu))$ and $\omega(\nu) < \nu$ in all states ϕ in which the PC is slack; (iii) $\bar{\omega}(\nu)$ is strictly increasing and $\omega(\nu)$ constant on the interval $[\underline{\nu}, \nu^-]$.

We first prove that $\bar{P}(\bar{\omega}_\phi(\nu)) - P(\omega_\phi(\nu)) > 0$ for $\nu > \nu^-$ (i.e., when $\omega_\phi(\nu) > \alpha - E[\phi]$ and $\theta_\phi = 0$). Suppose, to derive a contradiction, that $\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi) \leq 0$ for $\theta_\phi = 0$. Then (15), (16), and (17) imply the following consumption ordering

$$\bar{c}(\bar{\omega}_\phi) \leq c_\phi \leq c(\omega_\phi). \tag{40}$$

To arrive at a contradiction suppose that $\omega_\phi \geq \nu$. Once the economy recovers, promised-utility and profits remain constant such that consumption can be written as

$$\begin{aligned}
\bar{c}(\bar{\omega}_\phi) &= \bar{w} - \bar{P}(\bar{\omega}_\phi) + \beta \bar{P}(\bar{\omega}_\phi) = \bar{w} - \bar{P}(\bar{\omega}_\phi) + \frac{p_\phi}{R}\bar{P}(\bar{\omega}_\phi) + \frac{1-p_\phi}{R}\bar{P}(\bar{\omega}_\phi) \\
&> \underline{w} - \bar{P}(\bar{\omega}_\phi) + \frac{p_\phi}{R}\bar{P}(\bar{\omega}_\phi) + \frac{1-p_\phi}{R}\bar{P}(\bar{\omega}_\phi) \\
&\geq \underline{w} - P(\omega_\phi) + \frac{p_\phi}{R}\bar{P}(\bar{\omega}_\phi) + \frac{1-p_\phi}{R}P(\omega_\phi) \\
&\geq \underline{w} - P(\nu) + \frac{p_\phi}{R}\bar{P}(\bar{\omega}_\phi) + \frac{1-p_\phi}{R}P(\omega_\phi) = c_\phi.
\end{aligned}$$

We used $\bar{w} > \underline{w}$ to derive the first inequality, the fact that $[(1-p_\phi)/R - 1]\,\bar{P}(\bar{\omega}_\phi) \geq [(1-p_\phi)/R - 1]\,P(\omega_\phi)$ since $[(1-p_\phi)/R - 1] < 0$ and $\bar{P}(\bar{\omega}_\phi) \leq P(\omega_\phi)$ to derive the second inequality, and $\omega_\phi \geq \nu$ along with $P'(\nu) < 0$ to derive the last inequality. $\bar{c}(\bar{\omega}_\phi) > c_\phi$ contradicts the ordering in (40). Thus, we have proven that $\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi) \leq 0 \Rightarrow \omega_\phi < \nu$ for $\nu > \nu^-$. Next, recall that consumption is strictly increasing in the promise (since P is strictly concave) such that $\omega_\phi < \nu \Rightarrow c(\omega_\phi) < c(\nu) \leq c_\phi(\nu)$. This again contradicts the ordering in (40). Thus, we have shown that $\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi) > 0$ for $\nu > \nu^-$.

Combining, $\bar{P}(\bar{\omega}_\phi) - P(\omega_\phi) > 0$, $\theta_\phi = 0$, (15), (16), and (17) implies that

$$\bar{c}(\bar{\omega}_\phi(\nu)) > c_\phi(\nu) \geq c(\nu) > c(\omega(\nu)).$$

Since $c(\nu)$ is strictly falling in ν, the last inequality also implies that $\omega(\nu) < \nu$, as long as the recession lasts, the PC remains slack, and the lower bound $\underline{\nu}$ has not been reached. Once $\omega(\nu) = \underline{\nu}$, promised utility and consumption remains constant as long as the recession lasts and the PC remains slack. The first inequality also implies that consumption increases when the recession ends $\bar{c}(\bar{\omega}_\phi(\nu)) > c_\phi(\nu)$.

Finally, we use Topkis's theorem to show that $\bar{\omega}(\nu)$ must be strictly increasing in ν when $\omega(\nu) = \underline{\nu}$ and that $\bar{P}(\bar{\omega}_\phi(\nu)) - P(\omega_\phi(\nu)) > 0$ also for $\nu \leq \nu^-$. Given the threshold property, the planner solves the following maximization problem for all states ϕ with a slack PC

$$\langle \omega(\nu), \bar{\omega}(\nu) \rangle = \max_{\bar{\omega}, \omega \geq [\alpha - E[\phi]} \underline{w} - u^{-1}\left(x(\nu, \omega, \bar{\omega})\right) + \beta \begin{bmatrix} (1 - \Upsilon(\bar{\omega} - \omega))P(\omega) \\ +\Upsilon(\bar{\omega} - \omega)\bar{P}(\bar{\omega}) \end{bmatrix},$$

where $x(\nu, \omega, \bar{\omega}) = \alpha - \tilde{\phi}(\nu) + X(\Upsilon(\bar{\omega} - \omega)) - \beta\left[(1 - \Upsilon(\bar{\omega} - \omega))\omega + \Upsilon(\bar{\omega} - \omega)\bar{\omega}\right]$. The strictly positive cross-derivative of the objective function

$$\frac{u''(x(\nu, \omega, \bar{\omega}))\tilde{\phi}'(\nu)}{u'(x(\nu, \omega, \bar{\omega}))^2} \times \beta \Upsilon(\bar{\omega} - \omega) > 0,$$

shows that the objective is supermodular in $(\nu, \bar{\omega})$ for a given promised-utility ω. Thus, by Topkis's theorem, the promise $\bar{\omega}(\nu)$ is strictly increasing in ν once $\omega(\nu) = \underline{\nu}$. This also implies that the difference in profits is strictly positive, $\bar{P}(\bar{\omega}(\nu)) - P(\underline{\nu}) \geq \bar{P}(\bar{\omega}(\nu^{-})) - P(\underline{\nu}) > 0$, for $\nu \leq \nu^{-}$. Thus, we have shown that, $\bar{P}(\bar{\omega}_{\phi}(\nu)) - P(\omega_{\phi}(\nu)) > 0$ for all ν.

Finally, we establish that effort is non-monotone in ν. We established above that $\bar{\omega}(\nu)$ is strictly increasing in ν and $\omega(\nu)$ is constant for $\nu \leq \nu^{-}$. This implies that effort is strictly increasing in $\nu \in [\underline{\nu}, \nu^{-}]$. Consider, next, a range of high ν such that the PC is never binding in the next period. Then [THIS ARGUMENT IS INCOMPLETE AND MUST BE DEVELOPED] existing results in model with dynamic moral hazard without limited commitment ensure that effort is decreasing in ν. Thus, effort is a non-monotone function of ν.

This concludes the proof of the Proposition. ■

Proof of Proposition 4. Consider first the set of realizations for which the PC is slack, $\phi \geq \Phi^{pl}(\pi)$. In this case, it follows immediately from points (1)–(3) of Definition 1 that the planner's optimal choice is $c_{\phi} = C^{pl}(\pi)$, $\pi'_{\phi} = \Pi^{pl}(\pi)$, $\bar{\pi}'_{\phi} = \bar{\Pi}^{pl}(\pi)$, and $p_{\phi} = \Psi^{pl}(\Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi))$. In this case, the sovereign attains the utility $W^{pl}(\pi) = \alpha - \Phi^{pl}(\pi)$, namely, the outside option utility associated with the largest realization of ϕ for which the PC binds. When the PC binds, the solutions are dictated from the PC. In particular, the sovereign receives a utility $\alpha - \phi$. We can then define $\hat{\pi}^{pl}_{\phi}$ as a pseudo-profit such that, conditional on the realization ϕ and on $\pi = \hat{\pi}^{pl}_{\phi}$, the PC is just binding, i.e., $\hat{\pi}^{pl}_{\phi} = \left(\Phi^{pl}\right)^{-1}(\phi)$ and $W^{pl}\left(\hat{\pi}^{pl}_{\phi}\right) = \alpha - \phi$. Given this definition of $\hat{\pi}^{pl}_{\phi}$ and Definition 1, we can write the planner's optimal choice when $\phi \geq \Phi^{pl}(\pi)$ as $c_{\phi} = C^{pl}\left(\hat{\pi}^{pl}_{\phi}\right)$, $\pi'_{\phi} = \Pi^{pl}\left(\hat{\pi}^{pl}_{\phi}\right)$, $\bar{\pi}'_{\phi} = \bar{\Pi}^{pl}\left(\hat{\pi}^{pl}_{\phi}\right)$, and $p_{\phi} = \Psi^{pl}(\Pi^{pl}\left(\hat{\pi}^{pl}_{\phi}\right), \bar{\Pi}^{pl}\left(\hat{\pi}^{pl}_{\phi}\right))$. The definition of $\wp^{pl}(\pi, \phi)$ allows us to write the solution in a more compact form. The expression for the value function $\mu^{pl}(\pi)$ follows from the fact that the sovereign receives the utility $\alpha - \Phi^{pl}(\pi)$ if the PC is slack and $\alpha - \phi$ if the PC binds. ■

Proof of Proposition 5. We start by guessing that $\Phi(b(\pi)) = \Phi^{pl}(\pi)$. We show that under this guess the equilibrium is characterized by the same program as the COA. We then verify the guess, establishing that there exists a Markov equilibrium that decentralizes the COA.

We start by proving that $\wp(b(\pi), \phi) = b^{pl}\left(\wp^{pl}(\pi, \phi)\right)$, i.e.,

$$\wp(b(\pi), \phi) = \begin{cases} \hat{b}(\phi) = b^{pl}\left(\hat{\pi}^{pl}_{\phi}\right) & \text{if} \quad \phi < \Phi(b(\pi)) = \Phi^{pl}(\pi), \\ b(\pi) = b^{pl}(\pi) & \text{if} \quad \phi \geq \Phi(b(\pi)) = \Phi^{pl}(\pi). \end{cases}$$

Consider, first, the case in which the PC binds. We must show that, then, $\hat{b}(\phi) = b^{pl}(\hat{\pi}_{\phi})$. Setting $b = \hat{b}(\phi)$ in the bond revenue equation (32) for the competitive equilibrium yields the recursive equation

$$\hat{b}(\phi) = \frac{\Pi\left(\hat{b}(\phi)\right) - \int_{\phi_{\min}}^{\phi} \hat{b}(\phi')\, dF(\phi')}{1 - F(\phi)}.$$

This has the same form as the recursion (25) in the primal planning problem, where, if we set $\pi = \hat{\pi}^{pl}_{\phi}$,

$$b^{pl}\left(\hat{\pi}^{pl}_{\phi}\right) = \frac{\hat{\pi}^{pl}_{\phi} - \int_{\phi_{\min}}^{\phi} b^{pl}\left(\hat{\pi}^{pl}_{\phi'}\right) dF(\phi')}{1 - F(\phi)}.$$

Both equations only depend on ϕ. In particular, the two equations are identical for $\hat{b}(\phi) = b^{pl}(\hat{\pi}_\phi)$ and $\hat{\pi}_\phi^{pl} = \Pi\left(\hat{b}(\phi)\right)$.

Consider, next, the case in which the PC is slack. In this case, one obtains

$$b^{pl}(\pi) = \left(\frac{\pi - \int_{\phi_{\min}}^{\Phi^{pl}(\pi)} b^{pl}\left(\hat{\pi}_\phi^{pl}\right) dF(\phi)}{1 - F\left(\Phi^{pl}(\pi)\right)}\right) = \left(\frac{\Pi(b(\pi)) - \int_{\phi_{\min}}^{\Phi(b(\pi))} \hat{b}(\phi)\, dF(\phi)}{1 - F(\Phi(b(\pi)))}\right) = b(\pi).$$

Moreover, conditional on the guess of a common threshold, it is immediate to verify that $\mu(b(\pi)) = \mu^{pl}(\pi)$ and $\bar{\mu}\left(\bar{b}(\bar{\pi})\right) = \bar{\mu}^{pl}(\bar{\pi})$. Then, since the effort function stems from the same IC in the equilibrium and in the planning problem, it follows immediately that $\Psi\left(b(\pi), \bar{b}(\bar{\pi})\right) = \Psi^{pl}(\pi, \bar{\pi})$.

Next, we show that $C(b(\pi)) = C^{pl}(\pi)$. The gist of the argument is that we can rewrite the program (34)–(35) as the optimal choice of consumption and future expected debt repayment instead of consumption and future debt. More formally, the equilibrium functions solve

$$\begin{aligned}
&\left\langle B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\right\rangle = \\
&\arg\max_{(b',\bar{b}')\in[\underline{b},\tilde{b}]^2} \left\{ \begin{array}{c} u\left(Q\left(b',\bar{b}'\right)\times b' + \bar{Q}\left(b',\bar{b}'\right)\times \bar{b}' + \underline{w} - \wp(b,\phi)\right) + \\ -X\left(\Psi\left(b',\bar{b}'\right)\right) + \beta\left(\Psi\left(b',\bar{b}'\right)\times\bar{\mu}\left(\bar{b}'\right) + \left(1-\Psi\left(b',\bar{b}'\right)\right)\times\mu(b')\right) \end{array}\right\},
\end{aligned}$$

$$\begin{aligned}
C(\wp(b,\phi)) &= Q\left(B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\right)\times B(\wp(b,\phi)) + \\
&\quad \bar{Q}\left(B(\wp(b,\phi)), \bar{B}(\wp(b,\phi))\right)\times \bar{B}(\wp(b,\phi)) + \underline{w} - \wp(b,\phi).
\end{aligned}$$

Evaluate optimal consumption at $b = b(\pi)$ and use the fact that $Q\left(b',\bar{b}'\right)\times b' + \bar{Q}\left(b',\bar{b}'\right)\times\bar{b}' = \beta\left(\Psi\left(b',\bar{b}'\right)\bar{b}' + \left(1-\Psi\left(b',\bar{b}'\right)\right)\Pi(b')\right)$, to yield

$$\begin{aligned}
C(\wp(b(\pi),\phi)) &= \beta\left(1-\Psi\left(b\left(\Pi^e(\wp(b(\pi),\phi))\right), \bar{b}\left(\bar{\Pi}^e(\wp(b(\pi),\phi))\right)\right)\right) b\left(\Pi^e(\wp(b(\pi),\phi))\right) \\
&\quad + \Psi\left(b\left(\Pi^e(\wp(b(\pi),\phi))\right), \bar{b}\left(\bar{\Pi}^e(\wp(b(\pi),\phi))\right)\right)\bar{b}\left(\bar{\Pi}^e(\wp(b(\pi),\phi))\right) + \underline{w} - \wp(b(\pi),\phi)
\end{aligned}$$

where

$$\left\langle \Pi^e(\wp(b(\pi),\phi)), \bar{\Pi}^e(\wp(b(\pi),\phi))\right\rangle = \arg\max_{\pi',\bar{\pi}'} \left\{ \begin{array}{c} u(c) - X\left(\Psi\left(b(\pi'), \bar{b}(\bar{\pi}')\right)\right) \\ +\beta\left(\begin{array}{c} \Psi\left(b(\pi'), \bar{b}(\bar{\pi}')\right)\times\bar{\mu}\left(\bar{b}(\bar{\pi}')\right) + \\ \left(1-\Psi\left(b(\pi'), \bar{b}(\bar{\pi}')\right)\right)\times\mu(b(\pi')) \end{array}\right) \end{array}\right\}, \tag{41}$$

subject to $b(\pi') \leq (\Phi)^{-1}(\phi_{\max})$ and $c = \underline{w} - \wp(b(\pi),\phi) + \beta\left(1-\Psi\left(b(\pi'),\bar{b}(\bar{\pi}')\right)\right)\pi' + \Psi\left(b(\pi'),\bar{b}(\bar{\pi}')\right)\bar{\pi}'$. Since (41), the competitive equilibrium solves the same program as the planning problem, and $\left\langle \Pi^e(b(\pi)), \bar{\Pi}^e(b(\pi)), C(b(\pi))\right\rangle = \left\langle \Pi^{pl}(\pi), \bar{\Pi}^{pl}(\pi), C^{pl}(\pi)\right\rangle$. The fact that in equilibrium $\pi' = \Pi^e(b(\pi)) = \Pi^{pl}(\pi)$ and $\bar{\pi}' = \bar{\Pi}^e(b(\pi)) = \bar{\Pi}^{pl}(\pi)$ implies that the equilibrium law of motion for discounted profits is the same as in the COA. This set of equivalences also establishes that $W(b(\pi)) = W^{pl}(\pi)$.

To conclude the proof, we must verify the guess that $\Phi(b(\pi)) = \Phi^{pl}(\pi)$. To this aim, note that, in equilibrium $W(b(\pi)) = \alpha - \Phi(b(\pi))$, whereas in the COA $W^{pl}(\pi) = \alpha - \Phi^{pl}(\pi)$. Both are fixed-point conditions, since W depends on Φ and W^{pl} depends on Φ^{pl}. Moreover, we have established that $\Phi(b(\pi)) = \Phi^{pl}(\pi) \Rightarrow W(b(\pi)) = W^{pl}(\pi)$. Therefore, the fixed point condition is the same in the Markov and in the COA. Hence, the COA can be decentralized by the Markov equilibrium. Note that we cannot rule out the existence of multiple fixed points. ■

Proof of Corollary 1. With slight abuse of notation we write $x(0;\alpha) = x(0)$ to make the dependence of a variable x on α explicit. We show first that $\Phi(0;\alpha)$ is strictly increasing in α. Suppose to the opposite that $\partial\Phi(0;\alpha)/\partial\alpha \leq 0$. Then, the indifference condition implies that $\partial W(0;\alpha)/\partial\alpha > 0$ since $\partial\Phi(0;\alpha)/\partial\alpha = 1 - \partial W(0;\alpha)/\partial\alpha$. This also implies that, $\partial EV(0;\alpha)/\partial\alpha > 0$ since $EV(0;\alpha) = \int_{\aleph} \max\{W(0;\alpha), \alpha - \phi\}\, dF(\phi)$. At the same time, $\Pi(0;\alpha) = 0\ \forall\alpha$. However, we know from Proposition 5 that $P(\nu;\alpha) = \Pi(0;\alpha)$ implies that $\nu = EV(0;\alpha)$. Thus, for any $\alpha' > \alpha$ there exists an allocation that yields the same profits $\Pi(0;\alpha) = \Pi(0;\alpha')$ but higher promised-utility $EV(0;\alpha') > EV(0;\alpha)$ than the COA for α. This must be a contradiction, since - for a given α - it is always feasible for the planner to replicate the allocation with outside option α' by relaxing all the PCs by $\alpha' - \alpha$. However, the planner chose not to do so because its not optimal. Thus, $\partial\Phi(0;\alpha)/\partial\alpha > 0$.

Next, by the Theorem of the Maximum $W(0;\alpha)$ (and therefore $\Phi(0;\alpha)$) is continuous in α. Moreover, at the fixed point $W(0;W(0)) = W(0) = W(0) - \Phi(0;W(0))$ the threshold must be zero. Then, since $\Phi(0;\alpha)$ is continuous and strictly increasing in the relevant range of $\alpha \in [u(\underline{w})/(1-\beta), \bar{\mu}(0)]$ and $u(\underline{w})/(1-\beta) < W(0) < \bar{\mu}(0)$, it is enough to state that $\Phi(0;u(\underline{w})/(1-\beta)) < \Phi(0,W(0)) = 0$ and $\Phi(0;\bar{\mu}(0)) > \Phi(0,W(0)) = 0$ to establish that the fixed-point $W(0)$ exists in the interior of α's support and is unique. ■

A.1 Figure 4

In this section we show Figure 4 comparing the properties of the one-asset economy with and without renegotiation. This figure is discussed in Section 4.1 in the text.

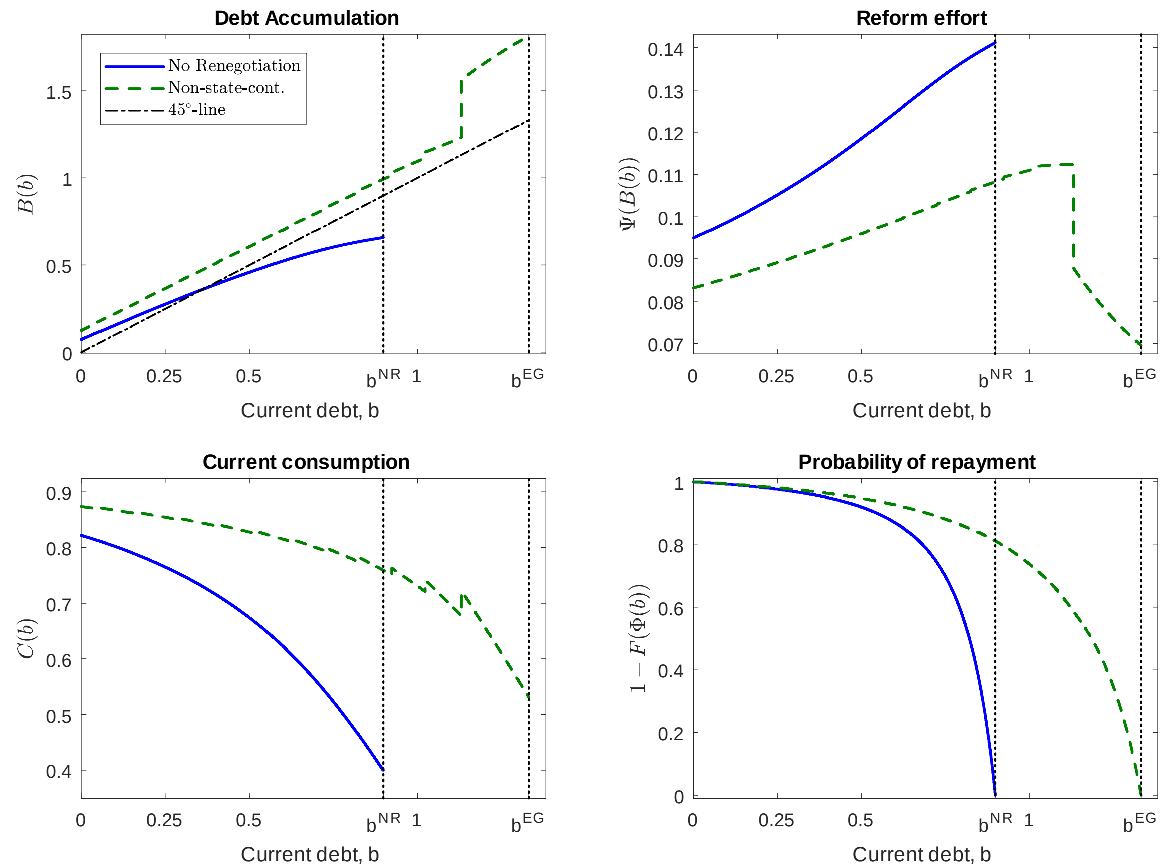

Figure 4: Policy functions of the one-asset economy, conditional on the maximum cost realization $\phi_{\max}$. The dashed lines show the Markov equilibrium with renegotiation, the solid lines the equilibrium where renegotiation is ruled out.