Queen's Economics Department Working Paper No. 1050

Optimal Dynamic Risk Sharing when Enforcement is a Decision Variable

Thorsten Koeppl
Department of Economics, Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

1-2005

Optimal Dynamic Risk Sharing when Enforcement is a Decision Variable

Thorsten V. Koeppl *

Queen's University

Department of Economics

Dunning Hall 210

94 University Ave.

Kingston, Ontario K7L 3N6

thor@qed.econ.queensu.ca

Phone: (613) 533 2271, Fax: (613) 533 6668

*I would like to thank Beth Allen, Narayana Kocherlakota and Jan Werner for their encouragement and many discussions on the subject. I have benefitted from discussions with James MacGee, Cyril Monnet, Erwan Quintin, Aldo Rustichini and Andrew Winton. I thank seminar participants at the SED meetings in Stockholm, the NBER Decentralization Conference at Georgetown University, the European Meetings of the Econometric Society in Venice, the Universities of Minnesota, Texas-Austin, Iowa, Toronto, Bocconi, ITAM and Vienna, and the Federal Reserve Banks of Minneapolis and Atlanta for their comments. I acknowledge financial support through the Walter W. Heller Memorial Dissertation Fellowship.

Abstract

Societies provide institutions that are costly to set up, but able to enforce long-run relationships. We study the optimal decision problem of using self-governance for risk sharing or governance through enforcement provided by these institutions. Third-party enforcement is modelled as a costly technology that consumes resources, but permits the punishment of agents who deviate from ex-ante specified allocations. We show that it is optimal to employ the technology whenever commitment problems prevent first-best risk sharing, but never optimal to provide incentives exclusively via this technology. Commitment problems then persist and the optimal incentive structure changes dynamically over time with third-party enforcement monotonically increasing in the relative inequality between agents.

Keywords: Limited Commitment, Risk Sharing, Third-party Enforcement.

JEL Classifications: C73, D60, D91, K49.

1 Introduction

Modern societies have developed institutions such as official legal systems or private arbitration systems that are costly to set up, but able to enforce contracts or agreements between people. In many situations, these enforcement institutions play a central role in governing contractual relationships. This is despite the fact that the contracting parties have the choice of self-governance directly through the structure of their contract. Our objective is here to study the problem of choosing self-governance vs. governance through a third party.

Economic transactions within long-term relationships are carried out by self-interested parties only if there is mutual interest in continuing the relationship. All transactions must, therefore, incorporate proper incentives to ensure that all parties continue to participate over time. These incentives are usually costly in the sense that they make it necessary to deviate from transactions that are optimal for both agents from an ex ante point of view. It is here that institutions can improve upon welfare by providing third-party enforcement: Agents involved in a long-run relationship are free to *choose* whether to rely on such institutions rather than on incentives through the structure of their agreement.

To govern relationships, third-party institutions (such as the legal system) are costly to set up as well. In essence, these institutions offer a threat of punishment in the form of fines or physical harm (e.g., imprisonment) in response to contractual violations, but cannot force performance of the contract itself. Their efficacy is based upon the ability to credibly commit to inflicting punishment in an objective manner if necessary. Objectivity arises from equal access as well as equal treatment of the parties involved in a relationship, while enforcement is achieved through the threat rather then the application of punishment. In fact, a strong presence of third-party enforcement manifests itself mainly in the performance of contracts and the absence of actual employment of punishment. Third-party enforcement can then be interpreted as a costly technology that threatens to inflict punishment in case of contract violations, even though this view is abstracting from other important factors such as limited effectiveness, information problems or the incentives for these institutions.

Given that these institutions are available but costly to set up, the question then arises as to what extent it is optimal for people to base incentive structures on these institutions. Are commitment problems persistent in the sense that the parties of a relationship do not want to rely exclusively on these institutions? Does the importance of outside (i.e., third-party) enforcement change dynamically over time? If so, what are the fundamentals that shape the dynamic evolution? Our contribution is to provide answers to these questions by analyzing the optimal use of costly outside enforcement in a long-run relationship.

We study a dynamic risk sharing problem between two risk averse agents where commitment is a priori limited in the spirit of Kehoe and Levine [6] and Kocherlakota [7]. Each period the agents face idiosyncratic income shocks. From an ex ante point of view, it is then optimal to transfer income ex post from an agent with high income realization to an agent with low income realization. We assume, however, that both agents cannot commit to make transfers they have agreed upon ex ante: At any point in time, each agent can choose to renege on the transfer and leave the risk sharing arrangement. In our set-up, incentives for the agents to honor transfers can be provided in two ways. First, agents can use the structure of the risk sharing arrangement itself to provide these incentives. Specifically, an agent can be induced to make a transfer of resources today if she is promised more expected utility in the future. Second, agents can rely on a "punishment" technology: Each period they can invest part of the overall resources in this technology. If investment occurs, the technology allows one to punish any agent who decides not to honor the transfer. This threat of punishment yields - for a resource cost - enforcement of transfers.

We show that - as long as the technology has convex costs and no fixed costs - it is optimal to employ the technology whenever the transfers necessary to support first-best risk sharing are not incentive compatible. It is never optimal, however, to provide incentives exclusively via this technology: The agents will always rely upon varying future promised utility - or, equivalently, the consumption profile - over time. Commitment problems are then only *partially* mitigated by using the technology and, thus, are persistent in this sense.

This implies that the enforcement choice (as represented by the investment decision) depends on the sequence of income shocks. Therefore, the optimal choice of punishment is history-dependent and inherently a dynamic one. For the case of two possible income realizations, we show analytically and numerically that more resources are spent on punishment as the difference in promised utility for the agents increases. Hence, we exhibit a positive relationship between inequality in future promised utility - or the relative position of the agents - and the use of third-party enforcement. In the long run, when no first-best allocation is incentive compatible, promised future utility is then equalized irrespective of the initial level of inequality between agents.

Existing work on dynamic risk sharing with limited commitment[1] takes the lack of commitment as exogenously given and focuses exclusively on the effects of optimally designed incentives that arise within the risk sharing relationship. The structure of these incentives is well understood. Kocherlakota [7] characterizes efficient risk sharing by relying on reversion to autarky as the appropriate punishment if an agent reneges on a risk sharing arrangement: Autarky is a credible punishment in the sense that it characterizes the set of subgame-perfect allocations in bilateral risk sharing environments.[2] More recently, Genicot and Ray [5] extend these results to a framework of risk sharing within coalitions of agents. This paper goes further than this existing literature by studying how agents *choose* optimally between internal incentives or incentives provided through

[1]See for example Phelan [12], Kocherlakota [7], Alvarez and Jermann [1] and Ligon et al. [11] among others.

[2]Gauthier et al. [4] show that optimally designed ex ante payments between agents can help reduce commitment problems. Ligon et al. [10] investigate the role of self-insurance in form of storage on the incentives to share risk over time.

enforcement by a third party from outside the relationship. Hence, we study per se the optimal degree of commitment within a risk sharing relationship.

Our research is related to the just emerging literature on contractual intermediaries. Parallel to our approach, Dixit [3] outlines a theory of enforcement intermediaries. He focuses on the role of third party enforcement in achieving cooperative outcomes in a prisoner's dilemma framework with random matching. The intermediary is modelled close to our approach as a player that can inflict punishments on other players for some positive fee. Ramey and Watson [13] investigate the optimal form of contractual intermediation or conflict resolution in a repeated prisoner's dilemma. Whereas we take the outside enforcement as given and investigate its optimal use by the contracting parties, these authors concentrate on understanding the existing design of such intermediation.[3]

The paper proceeds as follows: Section 2 presents the environment. In Section 3, we describe the optimal contracting problem and derive its recursive formulation. Section 4 characterizes the optimal contract and contains the main results. In Section 5, we present numerical examples concerning the optimal use of the punishment technology. Finally, Section 6 concludes by discussing our modelling choices and puts our contribution into a wider research context. All proofs appear in Appendix A, while Appendix B contains a formal analysis of a result discussed in Section 4.

2 Environment

Consider the following environment where time is discrete and indexed by $t = 0, 1, \ldots$. There are two infinitely lived agents $i = 1, 2$, who receive each period a stochastic

[3]It is useful to distinguish our paper from Krasa and Villamil [9] who study a static investment problem with differential information, where enforcement of the financial contract is a decision problem for the lender. Enforcement of the contract is costly and the contracting parties will try to avoid it via renegotiating the original contract whenever the lender cannot commit to seek enforcement of its terms. While studying the optimal form of the financial contract, the authors take the lack of commitment to be exogenous (i.e., not to be a choice variable).

endowment of a single good. Let $\omega = \{\omega_1, \omega_2, \dots\}$ be a sequence of independently and identically distributed random variables each having finite support $\Omega = \{1, 2, \dots, S\}$ and denote the probability of ω_t equaling s by $\pi_s > 0$ for all $s \in \Omega$. Define a t-history of ω by $\omega^t = \{\omega_1, \omega_2, \dots, \omega_t\}$ and let Ω^t be the set of all possible t-histories of ω. The endowment for agent $i = 1, 2$ in period t is determined by the realization of ω_t and denoted by $y^i_{t,s} \in \{y_1, y_2, \dots, y_S\}$ when $\omega_t = s$ for $t = 0, 1, \dots$. We assume that $y^1_{t,s} \neq y^2_{t,s}$, $\sum_{i=1}^2 y^i_{t,s} = Y > 0$ for all $s \in \Omega$ and $t = 0, 1, \dots$ and that the joint distribution of the endowment is symmetric; i.e., for every $s \in S$ there exists $s' \in S$ such that $y^i_{t,s} = y^j_{t,s'}$ and $\pi_s = \pi_{s'}$.

Preferences for both agents are described over ω^t-measurable consumption processes $c^i \in C = \{\{c^i_t\}_{t=0}^\infty | c^i_t : \Omega^t \longrightarrow [0, Y]\}$ and represented by the utility function

$$E_t \left[\sum_{\tau=0}^{\infty} \beta^\tau u(c^i_{t+\tau}) \right], \tag{1}$$

where $\beta \in (0, 1)$ and E_t expresses the expectation conditional on a history of shocks at date t. We assume that u is increasing, concave and twice continuously differentiable. Furthermore, u is bounded from below with normalization $u(0) = 0$ and $\lim_{c \to 0} u'(c) = \infty$.

Since the agents are risk averse and face idiosyncratic income shocks, there is an incentive to share income risk. We assume, however, that enforcement of arrangements to share risk is limited in the following sense: Each period, after uncertainty in period t is resolved and the current endowment $(y^1_{t,s}, y^2_{t,s})$ is known, an agent i can choose to remain in autarky forever. In this case, the agent will consume her endowment forever and will be excluded from future trade, thereby obtaining a utility of

$$u(y^i_{t,s}) + E_t \left[\sum_{\tau=1}^{\infty} \beta^\tau u(y^i_{t+\tau}) \right] \equiv u(y^i_{t,s}) + \beta V_{aut}, \tag{2}$$

where V_{aut} expresses the future expected utility from autarky which is independent of the realized history of shocks.

When sharing income risk, the agents also have access to a "punishment" technology that reduces an agent's current and future utility in case this agent decides to remain in autarky. Specifically, if this technology is operated at a level $d_t \in [0,1]$ in period t, the agent loses a fraction d_t of her current and future autarkic utility if she decides in period t to remain in autarky forever.[4] Operating this technology in period t at a level d_t requires an investment of resources equal to $\psi(d_t)$ in period t which depreciates fully after one period. We assume that the cost function $\psi(\cdot)$ is increasing, strictly convex and does not include any fixed costs:

Assumption 2.1. *1. $\psi' \geq 0$ and $\psi'' > 0$.*

2. $\psi(0) = 0$ and $\psi'(0) = 0$.

We assume further that the level of the punishment technology in any period t, d_t, is set before the current shock ω_t is realized. Therefore, the level of punishment in period t is *independent* of the current realization ω_t but can depend on the past history of realizations ω^{t-1}.[5] Formally, we denote the ω^{t-1}-measurable process of punishment levels by $d \in D = \{\{d_t\}_{t=0}^{\infty} | d_t : \Omega^{t-1} \longrightarrow [0,1]\}$, where Ω^{-1} is defined to contain a single element.

[4]Note that the severity of current *and* future punishment depends only on the level of d_t, i.e., on the level of punishment in the period when an agent decides to switch to autarky. Hence, a level of punishment chosen in future periods has no influence on punishments for switching to autarky in earlier periods.

[5]Third-party enforcement does then condition only on the fact whether contract violations occur or not. In the formulation chosen here neither the identity of the violator nor her particular situation - such as current income - matters for outside enforcement.

3 Describing Optimal Allocations

Before formulating the problem that describes optimal risk sharing between the agents we introduce some terminology. An *allocation* $(c^1, c^2, d) \in C \times C \times D$ is given by a consumption process for each agent and a process of punishment levels. An allocation is *feasible* if

$$c^1(\omega^{t-1}, s) + c^2(\omega^{t-1}, s) + \psi(d_t(\omega^{t-1})) \leq Y \text{ for all } t, (\omega^{t-1}, s). \tag{3}$$

An agent will switch to autarky for a given state s at time t if the continuation utility offered by an allocation is less than the value of autarky given the current level of punishment. Specifically, an agent i will honor the allocation if and only if

$$u(c^i(\omega^{t-1}, s)) + E_t\left[\sum_{\tau=1}^{\infty} \beta^{\tau} u(c^i_{t+\tau})\right] \geq (1 - d_t(\omega^{t-1}))\left[u(y^i_{t,s}) + \beta V_{aut}\right] \tag{4}$$

for all t, (ω^{t-1}, s).

Definition 3.1. *An allocation $(c^1, c^2, d) \in C \times C \times D$ is ex post incentive compatible if it satisfies inequality (4) for $i = 1, 2$ for all t, s. An allocation is incentive feasible if it is feasible for all t, s and ex post incentive compatible for $i = 1, 2$ for all t, s.*

We denote the set of incentive feasible allocations by $\Gamma \subset C \times C \times D$. Then, by Assumption 2.1, Γ is convex[6] and compact in the product topology. Next, let $\mathcal{U}$ be the set of joint utility levels that can be attained by an allocation in Γ and denote by $\mathcal{U}_i$ the range of utility levels of consumer i that is consistent with some allocation in Γ. The following lemma establishes properties of the set of attainable utility levels. All proofs are relegated to the appendix.

Lemma 3.2. *1. $\mathcal{U} \subset \mathbb{R}^2$ is compact.*

[6]Convexity follows from the concavity of u, the convexity of ψ and the fact that the ex post incentive compatibility constraints at t are linear in d_t.

2. $\mathcal{U}_i \subset I\!R$ *is compact and* $\mathcal{U}_1 = \mathcal{U}_2$*.*

Proof. See Appendix. □

A short remark concerning incentive feasibility is in place. The ex-post incentive compatibility constraint (4) compares the expected utility of an allocation with the utility obtain by choosing autarky forever and being punished by losing a fraction d_t of current and future utility. Remarkably, it is neither specified who pays the costs $\psi(d_t)$ if nobody reverts to autarky nor who pays the costs if some agent does.

As long as neither of the agents chooses autarky, the distribution of costs is irrelevant since for the utility attained by an allocation only the distribution of resources net of costs $\psi(d_t)$ matters. This implies that it is always possible to recover the costs for operating the punishment technology as long as the agents are participating. Of concern is then that, *given* an agent chooses autarky, it might be optimal for the other agent to choose autarky as well with the result that nobody would pay for the technology and it would not be feasible to operate the technology. When describing incentive feasible allocations this strategic interaction is, however, implicitly taken into account here since $d_t = 0$ is *always* feasible.[7]

The concept of incentive feasibility allows us to define optimal allocations. An allocation $(c^1, c^2, d) \in C \times C \times D$ is *optimal* if there exists no other incentive feasible allocation that provides both agents with at least as much expected utility at period 0 and at least one of them with strictly more expected utility at period 0.

[7]Indeed it is possible to make these arguments analytically precise at considerable costs of complication: one can motivate constraint (4) by identifying the set of incentive feasible allocations with the outcomes of a class of games formalizing repeated bargaining with voluntary participation where the distribution of costs among the agents is specified and non-participating agents do not have to bear any of the costs ex-post. This is achieved by establishing pay-off equivalence between the set of equilibria of all possible games and $\mathcal{U}$, i.e. the set of utility levels attainable through incentive feasible allocations. For details see Koeppl [8].

Define $V_{min} \equiv \min \mathcal{U}_i$ and $V_{max} \equiv \max \mathcal{U}_i$. We can then set up a modified Pareto-problem that describes optimal allocations taking into account incentive feasibility. Define the function $V : [V_{min}, V_{max}] \longrightarrow [V_{min}, V_{max}]$ as the solution to the problem (SP):

$$V(u_0) = \max_{(c^1,c^2,d)} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t^1) \right]$$

subject to

$$(c^1, c^2, d) \in \Gamma$$

$$E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t^2) \right] \geq u_0.$$

The function V refers then to the maximum level of expected utility agent 1 can obtain for any incentive feasible utility level $u_0 \in [V_{min}, V_{max}]$ that must be guaranteed for agent 2. Provided V is well defined it is clearly decreasing, since any incentive feasible allocation at $\hat{\hat{u}}_0 > \hat{u}_0$ is also incentive feasible at $\hat{u}_0$. Concavity of this function follows immediately from the convexity of ψ, the concavity of u and the fact that $V(u_0)$ is the maximum utility given u_0. V is then also continuous and differentiable almost everywhere. The next proposition shows that V is indeed well defined and strengthens some of these immediate results.

Proposition 3.3. *1. For all $u_0 \in [V_{min}, V_{max}]$, a solution to problem (SP) exists.*

2. There is an interval $[\underline{V}, \overline{V}] \subseteq \mathcal{U}_2$, where $\underline{V} < V_{aut} < \overline{V} = V_{max}$, such that V is strictly decreasing and strictly concave.

Proof. See Appendix. □

We now restrict V to the subset $[\underline{V}, \overline{V}]$ of its domain where it is strictly decreasing. By symmetry, $V : [\underline{V}, \overline{V}] \longrightarrow [\underline{V}, \overline{V}]$ and V describes the Pareto-frontier. Hence, any solution of the problem (SP) for given $u_0 \in [\underline{V}, \overline{V}]$ is an optimal allocation. Since u is strictly concave, for every $u_0 \in [\underline{V}, \overline{V}]$ there exists a unique optimal allocation. Furthermore, for any solution of problem (SP) the promise keeping constraint is strictly binding; i.e.,

$$E_0\left[\sum_{t=0}^{\infty}\beta^t u(c_t^2)\right] = u_0 \tag{5}$$

for all $u_0 \in [\underline{V}, \overline{V}]$.

These facts allow us to use the methods introduced by Spear and Srivastava [15] and Thomas and Worrall [17] to formulate the problem (SP) recursively. The state variable for this approach is given by the level u_0 of promised utility for agent 2.

Definition 3.4. *A contract is given by a collection of functions* $(\{c_s, u_s\}_{s=1}^S, d)$*, where* $d : [\underline{V}, \overline{V}] \longrightarrow [0, 1]$*,* $c_s : [\underline{V}, \overline{V}] \longrightarrow [0, Y]$ *for all* $s \in S$ *and* $u_s : [\underline{V}, \overline{V}] \longrightarrow [\underline{V}, \overline{V}]$ *for all* $s \in S$*.*

A contract consists of functions that determine the current level of consumption and the future expected promised utility for agent 2 for each state s, denoted by c_s and u_s respectively, as well as the level of punishment, denoted by d, in terms of the state variable u_0. The Pareto-frontier can then be determined recursively with the optimal allocation being described by a contract.

Proposition 3.5. *V satisfies the following functional equation (FE):*

$$V(u_0) = \max_{(\{c_s,u_s\}_{s=1}^S, d)} \sum_{s=1}^{S} \pi_s \left[u(Y - c_s - \psi(d)) + \beta V(u_s)\right]$$

subject to

$$\sum_{s=1}^{S} \pi_s \left[u(c_s) + \beta u_s\right] = u_0$$

$$u(Y - c_s - \psi(d)) + \beta V(u_s) \geq (1 - d)[u(y_s^1) + \beta V_{aut}] \quad \forall s$$

$$u(c_s) + \beta u_s \geq (1 - d)[u(y_s^2) + \beta V_{aut}] \quad \forall s$$

$$u_s \in [\underline{V}, \overline{V}] \quad \forall s.$$

Proof. See Appendix. □

Since the value function V is strictly concave and the constraint set describing the functional equation (FE) is convex, the solution to the above maximization problem is unique for any state u_0. Applying the Theorem of the Maximum, the optimal contract can then be described by continuous functions for d, c_s and u_s.

Proposition 3.6. *There exists a unique optimal contract* $(\{c_s^*, u_s^*\}_{s=1}^S, d^*)$*. Furthermore, the functions* d^**,* c_s^* *and* u_s^* *are continuous on* $[\underline{V}, \overline{V}]$*.*

Proof. See Appendix. □

4 Optimal Contracts

4.1 Persistence of Limited Commitment

We can now use the problem (FE) to characterize the optimal contract and, in particular, the decision concerning the use of the punishment technology. Let λ be the multiplier on the promise-keeping constraint and μ_s^i the multiplier on the ex post incentive compatibility constraint for agent i in state s. Assuming that the function V is differentiable everywhere with respect to u_0, we obtain the following set of first order conditions which are necessary and sufficient for the optimal contract on $(\underline{V}, \overline{V})$:

$$-(\pi_s + \mu_s^1)u'(Y - \psi(d) - c_s) + (\lambda\pi_s + \mu_s^2)u'(c_s) = 0 \tag{6}$$

$$(\pi_s + \mu_s^1)\beta V'(u_s) + (\lambda\pi_s + \mu_s^2)\beta = 0 \tag{7}$$

$$\sum_{s\in S} \mu_s^1[u(y_s^1) + \beta V_{aut}] + \mu_s^2[u(y_s^2) + \beta V_{aut}] - (\pi_s + \mu_s^1)u'(Y - \psi(d) - c_s)\psi'(d) \leq 0 \tag{8}$$

$$d \geq 0$$

$$\text{and} \tag{9}$$

$$d \cdot \left[\sum_{s\in S} \mu_s^1[u(y_s^1) + \beta V_{aut}] + \mu_s^2[u(y_s^2) + \beta V_{aut}] - (\pi_s + \mu_s^1)u'(Y - \psi(d) - c_s)\psi'(d)\right] = 0.$$

A brief comment about equations (6)-(9) is in order as we omit some of the corresponding Kuhn-Tucker conditions on the decision variables. For $u_0 \in (\underline{V}, \overline{V})$ it is optimal to make current consumption strictly positive for both consumers for all states (i.e., $Y - \psi(d) > c_s > 0$), and hence boundary conditions will never bind for c_s. Hence, it is never optimal to set $d = 1$ and we can restrict attention to $d \in [0, 1)$. Finally, rearrange equation (8) to obtain an expression for $\psi'(d)$ which shows that this expression will always be non-negative. Hence, even if $d = 0$, equation (8) will hold with equality. With respect to u_s the Kuhn-Tucker conditions are standard and, hence, omitted here.

We can reduce equations (6) and (7) to a single equation in the three decision variables given by

$$-V'(u_s) = \frac{u'(Y - \psi(d) - c_s)}{u'(c_s)}. \tag{10}$$

It is immediate that *given* d, $u_s > u_{s'}$ if and only if $c_s > c_{s'}$. Hence, u_s^* is an increasing function of c_s^*, or, equivalently, current consumption and future utility are varying together across states. A major complication arises from the fact that this equation depends also on the choice variable d. If d were constant over the state space $[\underline{V}, \overline{V}]$, this equation (together with the ex post incentive compatibility constraints for state s) would determine the dynamic evolution *independently* for each state $s \in S$. If d varies, however, the system of equations becomes genuinely dependent in the sense that one cannot conduct the analysis for each state separately.

The evolution of the state variable u_0 depends on which ex post incentive compatibility constraints are binding for a given state s. The following lemma summarizes results concerning the law of motion of u_0.

Lemma 4.1. *Let $u_0 \in (\underline{V}, \overline{V})$ and suppose that V is differentiable at u_0. Then the following hold:*

1. *If $\mu_s^i(u_0) = 0$ for all i, then $u_s^*(u_0)$=u_0.*

2. *If $\mu_s^1(u_0) > 0$ and $\mu_s^2(u_0) = 0$, then $u_s^*(u_0) < u_0$.*

3. *If $\mu_s^2(u_0) > 0$ and $\mu_s^1(u_0) = 0$, then $u_s^*(u_0) > u_0$.*

4. *Suppose $-V'(u_0) \leq 1$ and $\mu_s^1(u_0)\mu_s^2(u_0) > 0$. If $y_s^2 > y_s^1$, then $u_s^*(u_0) > u_0$.*

5. *Suppose $-V'(u_0) \geq 1$ and $\mu_s^1(u_0)\mu_s^2(u_0) > 0$. If $y_s^1 > y_s^2$, then $u_s^*(u_0) < u_0$.*

Proof. See Appendix. □

This determines the optimal variation of future promised utility except for cases where the ex post incentive compatibility constraints are binding for both agents simultaneously in some state s. In this case, the direction of the movements for u_0 can be ambiguous. Based on Lemma 4.1 it is possible to describe at least partially which agent's ex post incentive compatibility constraint is binding: If only one of the agents faces a binding constraint at some income level, he receives more future utility than he was promised initially. Since u_s^* is increasing in c_s^*, this agent must receive even more future utility at higher income levels. This is compatible with the first order conditions only if the agent is constrained at higher income levels. Hence, agents tend to be constrained when their income is high and, thus, have a strong reason to choose autarky over staying with the contract. This intuition is formally summarized in the lemma below.

Lemma 4.2.

1. *Suppose $u_s^*(u_0) > u_0$ for some s. If $y_{s'}^2 > y_s^2$, then $\mu_{s'}^{*2}(u_0) > 0$.*

2. *Suppose $u_s^*(u_0) < u_0$ for some s. If $y_{s'}^1 > y_s^1$, then $\mu_{s'}^{*1}(u_0) > 0$.*

Proof. See Appendix. □

Two main questions arise concerning the use of the punishment technology within the optimal contract. First, under what circumstances and to what extent is it optimal to use the punishment technology to achieve better risk sharing among the agents? Second, how does the decision concerning the use of the punishment technology vary endogenously over time?

From Lemma 4.1 it is clear that the state variable remains unchanged for some state $s \in S$ as long as none of the incentive constraints in this state is binding. We can distinguish two cases depending on whether the first-best allocation at u_0 is incentive feasible or not. For the first case, $\mu_s^i = 0$ for all i and s and, hence, from equation (8), $d^* = 0$. Turning to the case where the first best allocation at u_0 is not incentive feasible, at least some ex post incentive feasibility constraint is binding. Again by equation (8), it follows that $d^* > 0$ as long as $\mu_s^i > 0$ for some i and some s.[8] Beyond these straightforward observations it is possible to give a stronger result on the use of the punishment technology.

Theorem 4.3. *Let $u_0 \in (\underline{V}, \overline{V})$ and suppose that V is differentiable at u_0. Then there exists $s \in S$ such that $u_s^*(u_0) \neq u_0$ if and only if $d^*(u_0) > 0$.*

Proof. See Appendix. □

This theorem makes several important points. First, the agents will never rely exclusively on the technology that provides punishment to deal with limited commitment. Enforcement problems are always mitigated by a *combination* of using the explicit threat of punishment ($d^* > 0$) and implicit incentives provided through variations in future promised utility ($u_s^* \neq u_0$). Hence, any optimal contract will retain the commitment problem to a certain degree and counteract it by the intertemporal allocation of consumption between the agents. In this sense, commitment problems are persistent.

[8] Note that assuming $\psi'(0) = 0$ is essential for this result. In the case that $\psi'(0) > 0$, one might not want to use the punishment technology when a first-best allocation is not incentive feasible, but rather rely exclusively on internal incentives by setting $u_s^* \neq u_0$.

Second, the state variable u_0 will change with the realization of income shocks even though the punishment technology is employed. Thus, the distribution of wealth as summarized by u_0 varies over time and does not remain fixed. This implies that decisions concerning the use of the punishment technology are path dependent and vary over time due to changes in the wealth distribution. Therefore, the choice of enforcement is inherently a dynamic problem and cannot be treated as an ex ante static problem.

4.2 Punishment and Inequality

We turn now to the second question of how the level of punishment changes over time as the state variable u_0 evolves endogenously. Since the environment is symmetric with respect to the characteristics of the two agents, it is possible to restrict attention to the case where $u_0 \leq \overline{u}$ or $V(u_0) \geq u_0$, where $\overline{u} \in [\underline{V}, \overline{V}]$ such that $\overline{u} = V(\overline{u})$. If the first-best allocation for u_0 is incentive feasible, the contract is completely characterized by Theorem 4.3. We therefore turn to the case where the first-best allocation at u_0 is not incentive feasible.[9]

As inequality increases - i.e., as $|u_0 - \overline{u}|$ increases - it is more difficult to sustain efficient risk sharing since the outside option of leaving the arrangement becomes more attractive on average. Risk sharing has then to be supported by stronger incentives. These can be provided in two different ways: One can either increase $|u_s - u_0|$ (i.e., provide more indirect incentives via future promised utility) or one can invest more in the punishment technology. However, using more indirect incentives decreases future risk sharing on average. One should therefore expect that investment in the punishment technology would rise to at least partially counteract the negative effects on risk sharing. In other words, punishment should behave like a "normal" good in terms of inequality (in symbols, $|u_0 - \overline{u}|$) and substitution between the ways to provide incentives should not take place.

[9]It is straightforward to show that - independent of income shocks - for all $\beta \in (0, 1)$, there exists some level u_0 for which the first-best allocation is not incentive feasible.

Unfortunately, this question is too complex to be analyzed in full generality. We therefore assume for the reminder of the analysis in this section that there are only two states - i.e., $S = 2$ with $S = \{H, L\}$ - representing the current level of income for agent 2, where $y_H^2 > y_L^2$. Before characterizing the optimal choice of punishment as a function of the state variable u_0, we derive the following lemma:[10]

Lemma 4.4. *Suppose* $S = \{H, L\}$ *and* $u_0 < \overline{u}$. *If at* u_0 *the first-best allocation is not incentive feasible, then for any optimal contract* $\mu_s^{*1} = 0$ *and* $\mu_H^{*2} > 0$.

Proof. See Appendix. □

Whenever agent 2 is promised less utility than agent 1, at least one of her incentive compatibility constraints must be binding. Since there are only two states, Lemma 4.2 implies that her constraint when she has high income must necessarily bind. This fact allows us to prove the following monotonicity result for d^* which confirms the intuition outlined above for the case in which for some u_0 the first-best allocation is incentive feasible.

Theorem 4.5. *If* $S = 2$, *the policy function* $d^*(u_0)$ *is monotone on* $[\underline{V}, \overline{u}]$.

Proof. See Appendix. □

Corollary 4.6. *Suppose that for some* u_0 *the first-best allocation is incentive feasible. If* $S = 2$, *the policy function* $d^*(u_0)$ *is monotonically decreasing on* $[\underline{V}, \overline{u}]$ *and monotonically increasing on* $[\overline{u}, \overline{V}]$.

Proof. See Appendix. □

When inequality increases, it is optimal to decrease overall consumption and devote more resources to ensure enforcement of the risk sharing arrangement. Even though

[10]Even though the value function is *not* differentiable at $u_0 = \overline{u}$ if $S = 2$, none of the results in this section is affected by this non-differentiability. Moreover, if $S = 2$, $u_s^*(\overline{u}) = \overline{u}$, which shows that differentiability is necessary for the validity of Theorem 4.3.

Corollary 4.6 establishes this result only for the case when some first-best allocation is incentive feasible, numerical solutions described in more detail below confirm this result for the general case.

This result can be interpreted in a slightly different way. Suppose that one of the agents has higher bargaining power than the other. Then this agent has an interest in maintaining her position and is willing to spend more resources on outside enforcement. This enables her to at least partially lock in the relative position by keeping u_s "closer" to u_0. When the difference between the relative positions (i.e., $|u_0 - \overline{u}|$) increases, it is harder to maintain the current position, and more resources are spent on outside enforcement. Interestingly, however, Theorem 4.3 shows that outside enforcement is always too costly for the agents to maintain a current advantage in their bargaining power over time.

4.3 Long-run Implications of Optimal Contracts

After characterizing properties of the optimal contract, the question arises how the relationship between the agents develops in the long run. Of particular interest is how the relative position of the two agents adjusts in the long run and whether convergence to an invariant distribution over the state space occurs. We focus first on the two-state case. Later, we discuss what assumptions are necessary to derive a slightly weaker result for the case of an arbitrary finite number of states.

Before stating the main result of this section, it is necessary to introduce some notation. The stochastic process $\{\omega_t\}_{t=0}^{\infty}$ can be defined over the probability space $(\Omega^\infty, \mathcal{F}^\infty, \Pi^\infty)$, where an event is a particular sample path of the process, the σ-algebra $\mathcal{F}^\infty$ is generated by the cylinder sets of the process, and Π^∞ is the product measure based on the probabilities $\{\pi_1, \pi_2, \ldots, \pi_S\}$.

Given the optimal contract and an initial condition u_0, for every sample path $\omega \in \Omega^\infty$ it is possible to construct a sequence $\{u_t(\omega; u_0)\}_{t=0}^{\infty}$ of promised future utility levels for agent 2. Set $u_1(\omega; u_0) = u_s^*(u_0)$ if $s \in S$ is realized in period 0. Define $u_t(\omega; u_0)$ recursively by setting $u_t(\omega; u_0) = u^*_s(u_{t-1})$ if $s \in S$ is realized in period t for all $t > 0$. Moreover, denote the set of promised utility levels for which some first-best allocations is incentive feasible by $[u_{FB}, u^{FB}] \subset [\underline{V}, \overline{V}]$. Suppressing the arguments of u_t, we can then prove the following result on the long-run behavior of the optimal contract.

Theorem 4.7. *Let $S = 2$ and suppose that u^*_s is non-decreasing.*

1. *If there exists a first-best allocation that is incentive feasible, then for any optimal contract, $\lim_{t\to\infty} u_t = u_{FB}$ Π^∞-a.s. whenever $u_0 < u_{FB}$ and $\lim_{t\to\infty} u_t = u^{FB}$ Π^∞-a.s. whenever $u_0 > u^{FB}$.*
2. *If there does not exist a first-best allocation that is incentive feasible, then for any optimal contract, $\lim_{t\to\infty} u_t = \bar{u}$ Π^∞-a.s. for every $u_0 \in [\underline{V}, \overline{V}]$, where $\bar{u}$ satisfies $\bar{u} = V(\bar{u})$.*

Proof. See Appendix. □

Provided that there are only two states, for any initial condition u_0 the stochastic process for u_t converges with probability 1 to a unique point distribution. Hence, the availability of outside enforcement does not prevent the equalization of wealth between the agents over time or, for the case that the set of incentive feasible first-best allocations is non-empty, convergence to the "closest" element of this set.[11]

It is possible to give a slightly weaker result for the case that there are more than two states and no first-best allocation is incentive feasible. The optimal contract and the exogenous process of shocks define a Markov transition function. Continuity of the policy functions u^*_s establishes the Feller Property for this transition function. Moreover, the transition function will satisfy a mixing condition whenever the value function V

[11] It is straightforward to show that d^* being monotonically increasing in wealth inequality is a necessary condition for u^*_s to be increasing. Moreover, the monotonicity assumption on u^*_s seems rather weak since numerical solutions given below indicate that these functions are indeed increasing for a wide range of parameterizations.

is differentiable everywhere. Then standard results on weak convergence of Markov processes (e.g., Stokey, Lucas with Prescott [16, Theorem 12.12] yield convergence to a long-run stationary distribution of wealth independent of initial conditions, provided u_s^* is an increasing function of the state variable u_0. We defer details of this argument to the appendix.

To summarize our contribution, we have established three important theoretical results. First, commitment problems are persistent and not completely resolved by the use of costly third-party enforcement. Second, more unequally distributed bargaining power leads to greater reliance upon third-party enforcement. Last, the presence of third-party enforcement never prevents adjustments to a long-run, possibly equal, distribution of wealth across agents.

5 Numerical Solutions

The main analytical results of this section are derived under certain restrictions. We now provide further support for the generality of these results by presenting some numerical solutions for optimal contracts. Before presenting these results, we outline the algorithm used to solve for the Pareto frontier and the optimal contract, and describe how this algorithm can be implemented computationally.

The algorithm is based upon dynamic programming techniques. These methods are generally not applicable when solving incentive constrained problems, since the value function of the problem itself will influence the constraint set directly as can easily be observed from problem (FE). Hence, the constraint set will change with every iteration of the value function when solving the functional equation (FE). More importantly, the domain of the state variables for which the maximization problem is well defined will change with each iteration as well. Rustichini [14] demonstrates that one can modify standard dynamic programming methods in a straightforward way to handle these problems. He shows analytically that one can iterate directly on a guess for the value function in order to obtain convergence to the true value function of the incentive constrained problem. Conditions for this result are that the value function iteration starts with the value function of the unconstrained problem as an initial guess and that one adjusts the domain of the state variables in an appropriate way. Given these conditions, convergence is then monotonic from above to the true solution of the functional equation (FE). The details of the algorithm we employ are as follows:

Step 1: Calculate the initial guess J_0 for the value function V.

Step 2: Adjust the domain $\mathcal{D}_n$ of the state variable u_0 given the guess J_n for the value function V.

Step 3: Solve the static maximization problem for each realization of the state variable u_0 given J_n. Use this result to update the guess to J_{n+1}.

Step 4: If $sup_{u_0 \in \mathcal{D}_n}(J_n(u_0) - J_{n+1}(u_0)) > \epsilon > 0$, go back to Step 2.

Step 5: Use J_{n+1} to calculate policy functions and find the law of motion on $\mathcal{D}_n$.

To calculate the initial guess start with the Pareto frontier (which can be calculated analytically in a straightforward manner for any given utility function u) describing the first best solution of the risk sharing problem. Then define a new maximization problem (PRE) by deleting the ex post incentive compatibility constraints for consumer 1 that contain the value function V from problem (FE). Solve (PRE) by iterating over the value function of this problem with the Pareto frontier as the initial guess to obtain the initial guess for Step 1 of the algorithm above.[12]

[12]By using Blackwell's sufficient conditions (e.g., Stockey, Lucas with Prescott [16, Theorem 3.3]) it is

To implement the algorithm described above, we discretize the state space for u_0 and, hence, solve the functional equation for a finite number of values for u_0 in each iteration. The static maximization routine uses a linear quadratic approximation of the maximization problem with a cubic spline interpolation of the value function to guarantee twice continuous differentiability of the objective function. Finally, when computing the optimal contract, we perform a grid search over the decision variables of the maximization problem taking the solution of the value function as given.

Below we present the output of two examples that show the value function and the optimal decision with respect to the level of punishment d^* as functions of the state variable u_0. The utility function chosen is CES, $u(c) = \theta c^{1-\sigma}/(1-\sigma)$, where $\sigma \in (0,1)$ and $\theta > 0$, to satisfy the assumptions of Section 2. Costs are described by $\psi = \chi \cdot d^\zeta$, where $\zeta > 1$ and $\chi > 0$.

The first example exhibits a situation where some first-best allocation is incentive feasible. The cost function is given by $\psi = 4d^2$ and the Bernoulli utility function is $u(c) = \sqrt{c}$. Other parameters are given by $\beta = 0.8$ and $y_s \in \{1.8, 0.2\}$. Figure 1 compares the frontier of first-best allocations with the value function of problem (FE). Whereas both functions coincide for first-best allocations that are incentive feasible, the Pareto-frontier for the incentive constrained problem is bent inward and does not extend to the axes. Nevertheless, it extends beyond the value of autarky which is given by $V_{aut} \approx 4.4721$ units of utility.

The enforcement choice is depicted in Figure 2. Note that $d^* = 0$ for the region where the first-best allocation is incentive compatible. The graph also depicts a lower and an upper bound for the optimal decision d^* on the interval $[\underline{V}, \overline{u}]$. Figures 3 and 4 show the levels of future promised utility u_s^* and the current consumption levels c_s^* as a function of the state variable u_0.

The second example has the same cost function as above. The other parameters are

straightforward to show that iteration over value functions of problem (PRE) is a contraction operator. This ensures convergence to the "right" guess to apply the method of Rustichini [14].

changed to $\beta = 0.6$, $\theta = 1$, $\sigma = 0.4$ and $y_s \in \{1.5, 0.5\}$. For these values, there does not exist a first-best allocation that is incentive feasible. The Pareto-frontier, therefore, shifts inward relative to the value of first-best allocations as shown in Figure 5.

The enforcement choice depicted in Figure 6 is strictly positive. Furthermore, the policy function d^* is increasing in wealth inequality, a result we obtained in our numerical solutions for any parameterization. The non-differentiability of the value function at $\bar{u}$ for $S = 2$ causes some numerical error which is reflected in the small difference between the law of motion of both states at $u_0 = \bar{u}$ (cf. Figure 7).

Last, we stress that Figures 3 and 7 show that u_s^* is an increasing function of the state u_0, a result that can be confirmed in numerical experiments for a wide range of parameters. This gives us confidence that the results concerning the long-run properties of the optimal contract are true quite generally as the assumptions of Theorem 4.7 seem to be satisfied with wide generality.

6 Concluding Remarks

Our analysis demonstrates that commitment problems persist even though the parties sharing risk have access to costly third-party enforcement. This result is strong in the sense that we impose rather weak restrictions on the cost structure, thereby giving the use of enforcement the best possible chance. More importantly, even though the presence of fixed costs will introduce a barrier to using third-party enforcement, persistence depends only on the fact that costs are increasing in the use of punishments. As long as this is the case, there are always incentives to avoid part of these costs by relying also on intertemporal features of the contract. Since commitment problems become more severe with increasing differences in the relative position of the agents, the monotonicity property of optimal enforcement is not too surprising. However, it is striking that the costs of keeping fixed a specific positive level of inequality *always* outweigh the existing incentives to do so; the technology is never "abused" to lock in a specific level of inequality.

We have assumed that enforcement cannot depend on the current realization of the income shock. This can be justified along two lines. First, impartial punishment is based on the violation of the contract (i.e., leaving the arrangement) disregarding other circumstances like differences in current income. Second, if punishment depends on the current realization of the shock, the incentives of the two agents are not properly aligned. Whoever has a high income realization prefers a strictly lower punishment level than the other agent. Hence, communicating the current income distribution to the outside would be difficult if not impossible. This problem does not occur if punishment next period depends only on the new level of promised utility set endogenously by the agents in the previous period. Future work should concentrate on modelling a non-cooperative game between the agents and a third agent providing enforcement. It is then possible to study not only the incentives of the third party, but also difficulties in the communication between agents and the outside party.

By using a dynamic contracting approach for our analysis we are silent about any initial condition that would pin down the dynamic evolution of the long-run relationship. Since our description of the optimal contract is independent of any initial conditions, the outcome of any ex ante bargaining procedure would simply consist of the optimal contract described here evaluated at an initial condition reflecting the relative bargaining power of the agents. By construction, there would be no incentives for the agents to violate this contract at any later time.

A final remark concerns decentralizing the environment. Optimal contracts could be decentralized as a financial markets equilibrium with complete markets and portfolio constraints. These constraints mimic how stringent the incentive compatibility or participation constraints for the optimal contract are. Since the agents choose the set of feasible allocations in our problem, the value of the portfolio constraints must vary dynamically over time as uncertainty is resolved. The decentralization should reflect the optimal choice of enforcement and, hence, offers a *conceptually* genuine theory of endogenous portfolio constraints.[13] The main difficulty clearly arises from the problem of distributing the enforcement costs among the agents. The requirement here is to construct either a market mechanism or a direct mechanism that distributes the costs without disturbing the properly decentralized financial decisions of the agents.

Appendix A

Proof of Lemma 3.2:

1. Since $\mathcal{U} \in I\!R^2$, $\mathcal{U}$ is compact if and only if $\mathcal{U}$ is closed and bounded. Obviously, $\mathcal{U} \subset [0, 1/(1-\beta)u(Y)]^2$ is bounded. Let u_n be a convergent sequence such that $u_n \in \mathcal{U}$ for all n and denote its limit by $\hat{u}$. Then, there is a sequence of allocations (c_n^1, c_n^2, d_n) in Γ such that the n-th allocation generates the utility levels u_n for all n. Since Γ is compact in the product topology, there exists a subsequence that converges to an allocation $(\hat{c}^1, \hat{c}^2, \hat{d}) \in \Gamma$. Since u_n converges to $\hat{u}$, every subsequence of u_n also converges to $\hat{u}$. We can restrict the function $u(\cdot)$ in (1) to the interval $[0, Y]$; the utility function defined by (1) is then continuous in the product topology. Hence, $\hat{u}$ is generated by the allocation $(\hat{c}^1, \hat{c}^2, \hat{d}) \in \Gamma$. Thus, $\hat{u} \in \mathcal{U}$ and $\mathcal{U}$ is closed.

2. For $i = 1, 2$, $\mathcal{U}_i$ is the projection of $\mathcal{U}$ into $I\!R$. Hence, $\mathcal{U}_i$ is compact. By symmetry, $\mathcal{U}_1 = \mathcal{U}_2$.

□

Proof of Proposition 3.3:

[13] Alvarez and Jermann [1] suggest a decentralization of an economy with exogenously given participation constraints. Their borrowing constraint are "endogenous" only to the extent that they are not completely arbitrary, but rather determined by the fundamentals of the economy.

1. Let $u_0 \in [V_{min}, V_{max}]$ and consider the following maximization problem (UP):

$$\begin{aligned} &\max u^1 \\ &\text{subject to} \\ &\quad (u^1, u^2) \in \mathcal{U} \\ &\quad u^2 \geq u_0. \end{aligned}$$

   Since $\mathcal{U} \subset \mathbb{R}^2$ is compact, the constraint set of (UP) is compact and by continuity of the objective function, this problem has a solution in $\mathcal{U}$. Thus, there exists an incentive feasible allocation that attains these utility levels. Hence, the problem (SP) has a solution for all $u_0 \in [V_{min}, V_{max}]$.

2. Suppose V is not strictly decreasing over $[V_{min}, V_{max}]$. Since V is concave and continuous, V is either constant over $[V_{min}, V_{max}]$ or constant over a subinterval starting from V_{min} and strictly decreasing over the remainder of the interval. It is therefore sufficient to show that V is strictly decreasing at V_{aut}, which is clearly an element of $\mathcal{U}_2$.

   Let $u_0 = V_{aut}$. Suppose first that at the optimal allocation some ex post incentive constraint for agent 2 is not binding in period $t = 0$. Then, for some $s \in S$,

$$u(c_{0,s}^2) + E_1 \left[\sum_{t=1}^{\infty} \beta^t u(c_t^2) \right] > (1 - d_0) \left[u(y_s^2) + \beta V_{aut} \right].$$

   Hence, we can decrease $c_{0,s}^2$ and increase $c_{0,s}^1$ slightly without violating incentive feasibility. Thus, there exists $\tilde{u}_0 < V_{aut}$ such that $V(\tilde{u}_0) > V(V_{aut})$.

   Suppose now that for the solution to (SP) given V_{aut}, all ex post incentive compatibility constraints bind for agent 2 at $t = 0$. Since $u_0 = V_{aut}$, we have

$$u_0 = \sum_{s \in S} \pi_s (1 - d_0)[u(y_{t,s}^2) + \beta V_{aut}] = (1 - d_0) V_{aut},$$

   which implies that $d_0 = 0$. We construct an allocation for some $u_0 < V_{aut}$ that gives agent 1 a utility which is strictly higher than $V(V_{aut})$. Define the following two functions for a given $s \in S$ and given the optimal allocation:

$$f_1(\epsilon) = u(c^2_{0,s} - \psi(\epsilon)) + E_1\left[\sum_{t=1}^{\infty} \beta^t u(c^2_t)\right]$$

and

$$f_2(\epsilon) = (1-\epsilon)\left[u(y^2_{t,s}) + \beta V_{aut}\right].$$

Then, $f_1'(\epsilon) = -\psi'(\epsilon)u'(c^2_{0,s} - \psi(\epsilon))$ and $f_2'(\epsilon) = -\left[u(y^2_{t,s}) + \beta V_{aut}\right]$. Define $B \equiv u'(\frac{1}{2}c^2_{0,s})$. Optimality of the allocation and $\lim_{c\to 0} u'(c) = \infty$ yields $c^2_{0,s} > 0$ and hence $B < \infty$. Since $\psi'(0) = 0$, for ϵ close to 0, we obtain

$$f_1'(\epsilon) < -B\psi'(\epsilon) < f_2'(\epsilon) < 0.$$

Hence, there exists an incentive feasible allocation that gives $u_0 < V_{aut}$ to agent 2 and $V(V_{aut})$ to agent 1 such that some ex post incentive compatibility constraint for agent 2 is not binding at $t = 0$. Thus, we can construct an allocation where agent 2 obtains $\tilde{u}_0 < V_{aut}$ and $V(\tilde{u}_0) > V(V_{aut})$. By concavity, V must then be strictly decreasing on $[\tilde{u}_0, V_{max}]$.

Let $\hat{u}, \hat{\hat{u}} \in [\underline{V}, \overline{V}]$ and $\hat{u} < \hat{\hat{u}}$. Let $(\hat{c}^1, \hat{c}^2, \hat{d})$ and $(\hat{\hat{c}}^1, \hat{\hat{c}}^2, \hat{\hat{d}})$ be the corresponding solutions to problem (SP). Since V is strictly decreasing on $[\underline{V}, \overline{V}]$, after some history ω^t, $\hat{c}^1_{t+1,s} < \hat{\hat{c}}^1_{t+1,s}$. Strict concavity of u implies strict concavity of V.

□

Proof of Proposition 3.5:

Let $u_0 \in [\underline{V}, \overline{V}]$ be given and let $(\hat{c}^1, \hat{c}^2, \hat{d})$ be an optimal allocation. Define $(\hat{c}^1_{1,s}, \hat{c}^2_{1,s}, \hat{d}_{1,s})$ as the continuation allocation from $t = 1$ onwards when state $s \in S$ occurred in period $t = 0$.

Claim: The continuation allocation $(\hat{c}^1_{1,s}, \hat{c}^2_{1,s}, \hat{d}_{1,s})$ from period $t = 1$ onwards given $s \in S$ occurred in period $t = 0$ is an optimal allocation.

Suppose not. Then after $s \in S$ occurs in period $t = 0$, there exists a continuation allocation $(\tilde{c}^1_{1,s}, \tilde{c}^2_{1,s}, \tilde{d}_{1,s})$ from period $t = 1$ that is feasible and yields at least as much utility for both agents and strictly more utility for one agent than $(\hat{c}^1_{1,s}, \hat{c}^2_{1,s}, \hat{d}_{1,s})$, the one specified in the optimal allocation. Define a new allocation by replacing the part of the old allocation after the event s occurs in the first period by $(\tilde{c}^1_{1,s}, \tilde{c}^2_{1,s}, \tilde{d}_{1,s})$. This allocation is clearly incentive feasible. Furthermore, it delivers at least as much utility to both agents as the optimal allocation and strictly more expected utility for one agent. Hence, $(\hat{c}^1, \hat{c}^2, \hat{d})$ is not optimal, which is a contradiction.

Define $\tilde{V}(u_0)$ to be the value of the solution to the right hand side of the objective function in (FE).

Claim: $V(u_0) \leq \tilde{V}(u_0)$.

Let $(\hat{c}^1, \hat{c}^2, \hat{d})$ be the optimal allocation given u_0. Similarly, let $(\hat{c}^1_{1,s}, \hat{c}^2_{1,s}, \hat{d}_{1,s})$ be the continuation allocation of the optimal allocation at $t = 1$ after $s \in S$ occurred in period $t = 0$. Define

$$\hat{u}_s = E_1 \left[\sum_{t=1}^{\infty} \beta^{t-1} u(\hat{c}^2_{1,s}) \right]$$

for all $s \in S$. By the previous claim, for all $s \in S$ the continuation allocation is optimal lies; i.e.,

$$V(\hat{u}_s) = E_1 \left[\sum_{t=1}^{\infty} \beta^{t-1} u(\hat{c}^1_{1,s}) \right].$$

Consider now the contract $(\{\hat{c}^2_{0,s}, \hat{u}_s\}_{s=1}^S, \hat{d}_0)$. This contract is clearly feasible and ex post incentive compatible for (FE). Furthermore, by the definition of $\hat{u}_s$,

$$\sum_{s \in S} \pi_s [u(\hat{c}_{0,s}) + \beta \hat{u}_s] = u_0.$$

Since $\{\hat{c}^1, \hat{c}^2, \hat{d}\}$ attains a utility of $V(u_0)$ for agent 1, it follows that $\tilde{V}(u_0) \geq V(u_0)$.

Claim: $V(u_0) \geq \tilde{V}(u_0)$.

Let $(\{\hat{c}_s, \hat{u}_s\}_{s=1}^S, \hat{d}_0)$ be the solution to the right hand side of the objective function in (FE) yielding $\tilde{V}(u_0)$. Since $\hat{u}_s \in [\underline{V}, \overline{V}]$ for all $s \in S$, there exists an optimal allocation that yields $\hat{u}_s$ for agent 2 and $V(\hat{u}_s)$ for agent 1 for every $s \in S$. Call this allocation $(\hat{c}^1_{1,s}, \hat{c}^2_{1,s}, \hat{d}_{1,s})$.

Consider the allocation $(\{Y - \psi(\hat{d}) - \hat{c}_s\}_{s=1}^S, \hat{c}^1_{1,s}, \{\hat{c}_s\}_{s=1}^S, \hat{c}^2_{1,s}, \hat{d}, \hat{d}_{1,s})$. The allocation is incentive feasible for the problem (SP) and, since $\sum_{s\in S} \pi_s[u(\hat{c}_s) + \beta\hat{u}_s] = u_0$, agent 2 receives utility u_0. Since agent 1 receives $\sum_{s\in S} \pi_s[u(\hat{c}_s) + \beta V(\hat{u}_s)] = \tilde{V}(u_0)$ from the allocation, $V(u_0) \geq \tilde{V}(u_0)$.

□

Proof of Proposition 3.6:

Since the constraint correspondence is compact-valued and continuous, the Theorem of the Maximum (Debreu [2, Theorem 1.8 (4)]) applies. As V is strictly concave and the constraint set is convex, the solution of the maximization problem is unique and, therefore, given by unique policy functions d and c_s, u_s for all $s \in S$.

□

Proof of Lemma 4.1:

1. Using equation (7), $\mu^i_s = 0$ for all i implies that $-V'(u_s) = \lambda$. By the envelope theorem, $\lambda = -V'(u_0)$ and the result follows from the fact that V is strictly decreasing.

2. Using the envelope theorem, equation (7) reduces to

$$-V'(u_s) = -V'(u_0)\frac{\pi_s}{\pi_s + \mu^1_s}.$$

   Hence, $-V'(u_s) < -V'(u_0)$. Since V is strictly decreasing and strictly concave, $u_s < u_0$.

3. The proof is analogous to the one given above.

4. If both ex post incentive compatibility constraints are binding in some state s, they must also be binding in the state s' where the income of both agents is reversed. Otherwise, at least one of the agents can be made better off by replicating the contract for state s' in state s. Hence the original contract cannot be optimal.

Thus, the allocations for the pair of states (s, s') must be symmetric in the sense that agent 1 receives agent 2's allocation of state s in state s'. Without loss of generality assume that $y^2_{s'} > y^1_{s'}$. Since u_s is increasing in c_s for given d, we obtain

$$u_{s'} = V(u_s) > u_s = V(u_{s'})$$

and

$$c_{s'} = Y - \psi(d) - c_s > c_s = Y - \psi(d) - c_{s'}.$$

Strict concavity of V and symmetry of the problem imply $-V'(u_s) = 1$ if and only if $V(u_s) = u_s$. Since $u_{s'} > V(u_{s'})$ and V is strictly concave, $-V'(u_{s'}) > 1$. By hypothesis, $-V'(u_0) \leq 1$. Hence, $u_{s'} > u_0$.

5. The proof is analogous to the one given above.

□

Proof of Lemma 4.2:

Let $u_s > u_0$ for some $s \in S$. By Lemma 4.1, $\mu^2_s > 0$ and the ex post incentive compatibility constraint binds for agent 2 in state s. Let s' be any state such that $y^2_{s'} > y^2_s$. Then, since u_s is increasing in c_s, it must be the case that $u_{s'} > u_s > u_0$. Hence, $\mu^2_{s'} > 0$. The second statement is proved by an analogous argument.

□

Proof of Theorem 4.3:

Suppose $d = 0$. Then by equation (8), $\mu^i_s = 0$ for all $i = 1, 2$ and all $s \in S$. By Lemma 4.1, $u_s = u_0$ for all $s \in S$.

Suppose $d > 0$. Suppose further that $u_s = u_0$ for all $s \in S$. By the envelope theorem we have $\lambda = -V'(u_0)$. From equation (7) we obtain $-\mu^1_s V'(u_0) = \mu^2_s$. If $\mu^1_s = \mu^2_s = 0$ for all s, then $d = 0$ by equation (8), which is impossible. Hence, for some s, $\mu^i_s > 0$ for $i = 1, 2$.

Since $Y - \psi(d)$ and u_s are constant across states, it follows from equation (10) that c_s is also constant across states. Thus, the utility levels for both agents are constant across states. This implies that for each agent the ex post incentive compatibility constraint can be binding for at most one income level. Since for all $s \in S$, $y_s^1 \neq y_s^2$, we have $\mu_s^i = 0$ for some $i = 1, 2$ in all states s. A contradiction is therefore obtained.

□

Proof of Lemma 4.4:

Since at u_0 the first-best allocation is not incentive feasible, at least some incentive constraint must be binding and, by Theorem 4.3, $d > 0$. It is also clear that all the incentive constraints cannot be binding; otherwise $u_0 = V(u_0) < V_{aut}$ and the contract cannot be optimal.

<u>*Claim:*</u> $\mu_s^1 \mu_s^2 = 0$ for all $s \in \{H, L\}$.

Suppose not. Then there exists $s \in \{H, L\}$ such that $\mu_s^i > 0$ for $i = 1, 2$. Then, since not all incentive constraints can be binding, $\mu_{s'}^i = 0$ for some i and $s' \neq s$. Without loss of generality, assume $s' = L$ and $\mu_L^1 = 0$. Then,

$$u(Y - \psi(d) - c_L) + \beta V(u_L) > u(c_H) + \beta u_H = (1 - d)[u(y_H) + \beta V_{aut}]$$

and

$$u(c_L) + \beta u_L \geq u(Y - \psi(d) - c_H) + \beta V(u_H) = (1 - d)[u(y_L) + \beta V_{aut}].$$

Therefore, both agents receive higher or equal utility in state $s' = L$ than in state $s = H$. This cannot be optimal since one can replicate the contract for state $s' = L$ in state $s = H$ and make at least some agent better off without making the other worse off.

<u>*Claim:*</u> $\mu_s^1 = 0$ for all $s \in \{H, L\}$.

Suppose first that $\mu_H^1 > 0$ (i.e., agent 1's incentive constraint binds when his income is low). Hence $u_H < u_0$. By Lemma 4.2, $\mu_L^1 > 0$. By the previous claim, $\mu_s^2 = 0$ for all s. Thus, $V(u_0) < u_0$, a contradiction.

Suppose now that $\mu_L^1 > 0$ (i.e., agent 1's incentive constraint binds when his income is high). Then, by the previous claim, $\mu_L^2 = 0$. Furthermore, $\mu_H^1 = 0$, since otherwise

$$V(u_0) = E\left[(1-d)\left(u(y_s^1) + \beta V_{aut}\right)\right] < V_{aut},$$

which contradicts $u_0 < \bar{u}$.

By incentive feasibility,

$$u(c_H) + \beta u_H \geq u(Y - \psi(d) - c_L) + \beta V(u_L) = (1-d)[u(y_H) + \beta V_{aut}].$$

Since $u_0 < \bar{u}$, $V(u_0) > u_0$. Hence, $u(Y-\psi(d)-c_H)+\beta V(u_H) > u(c_L)+\beta u_L$. Therefore, both agents receive higher or equal utility in state $s = H$ than in state $s = L$. This cannot be optimal since one can replicate the contract for state $s = H$ in state $s = L$ and make at least some agent better off without making the other one worse off.

Claim: $\mu_H^2 > 0$.

The previous claim implies that $\mu_s^2 > 0$ for some $s \in S$. If $\mu_L^2 > 0$, $u_L > u_0$. By Lemma 4.2, $\mu_H^2 > 0$ which completes the proof.

□

Proof of Theorem 4.5:

We show that the policy function for d must be monotone on $[\underline{V}, \bar{u}]$. Symmetry implies that it must be monotone - with the sign of the slope reversed - on the other part of its domain. Without loss of generality, we assume throughout the proof that $d(u_0) > 0$ (i.e., that at u_0 the first-best allocation is not incentive feasible). We proceed first with an intermediate result.

Claim: If $\mu_s^2 > 0$ for all $s \in S = \{H, L\}$ at $\hat{u}_0$, then $\hat{\hat{d}} > \hat{d}$ for all $\hat{\hat{u}}_0 < \hat{u}_0$.

Suppose not. Then, there exists $\hat{\hat{u}}_0 < \hat{u}_0$ such that $\hat{\hat{d}} < \hat{d}$. By incentive feasibility,

$$u(\hat{\hat{c}}_s) + \beta\hat{\hat{u}}_s \geq (1 - \hat{\hat{d}})\left[u(y_s^2) + \beta V_{aut}\right].$$

Since $\mu_s^2 > 0$ for all $s \in S = \{H; L\}$ at $\hat{u}_0$, we have

$$u(\hat{c}_s) + \beta \hat{u}_s = (1 - \hat{d}) \left[u(y_s^2) + \beta V_{aut} \right].$$

and

$$\hat{\hat{u}}_0 \geq (1 - \hat{\hat{d}}) V_{aut} > (1 - \hat{d}) V_{aut} = \hat{u}_0.$$

This is a contradiction.

Suppose now that the policy function is not monotone on a subinterval of $[\underline{V}, \bar{u}]$. Continuity implies that there exists $\hat{u}_0 < \tilde{u}_0$ such that $\hat{d} = \tilde{d} > 0$. Since d is the same, strict concavity of V and u imply that $\hat{u}_H = \tilde{u}_H$ and $\hat{c}_H = \tilde{c}_H$. Using equation (8) and the claim above, we obtain

$$\frac{\sum_{s \in S} [V'(\hat{u}_0) - V'(\hat{u}_s)]\gamma_s}{\sum_{s \in S} u'(Y - \psi(\hat{d}) - \hat{c}_s)} = \psi'(\hat{d}) =$$
$$= \psi'(\tilde{d}) = \frac{[V'(\tilde{u}_0) - V'(\hat{u}_H)]\gamma_H}{u'(Y - \psi(\hat{d}) - \hat{c}_H) + u'(Y - \psi(\hat{d}) - \tilde{c}_L)},$$

where γ_s denotes the value of the outside option if income is given by y_s. Since $\hat{u}_0 < \tilde{u}_0$, $\hat{u}_L \geq \hat{u}_0$ and V is strictly concave,

$$\sum_{s \in S} [V'(\hat{u}_0) - V'(\hat{u}_s)]\gamma_s > [V'(\tilde{u}_0) - V'(\hat{u}_H)]\gamma_H.$$

To satisfy $\psi'(\hat{d}) = \psi'(\tilde{d})$, we need $\hat{c}_L > \tilde{c}_L$. Since u_s is an increasing function of c_s, we have $\hat{u}_L > \tilde{u}_L = \tilde{u}_0$. Since the allocation for state $s = H$ is the same for $\hat{u}_0$ and $\tilde{u}_0$, it follows that $\hat{u}_0 > \tilde{u}_0$, which is a contradiction.

□

Proof of Corollary 4.6:

If there exists a first-best allocation that is incentive feasible, $d = 0$ for some interval $[u_{FB}, \overline{u}]$ and $d > 0$ for $[\underline{V}, u_{FB})$. The result then follows.

□

Proof of Theorem 4.7:

1. Let $u_0 \in [\underline{V}, u_{FB})$. Define $\mathcal{A} \equiv \{\omega \in \Omega^\infty | \omega_t = H$ for finitely many t$\}$. Clearly, $\Pi^\infty(\mathcal{A}^c) = 1$. Hence, $\lim_{t\to\infty} u_t = u_{FB}$ Π^∞-a.s. if $\lim_{t\to\infty} u_t = u_{FB}$ for all $\omega \in \mathcal{A}^c$.

   Let $\omega \in \mathcal{A}^c$. By Lemma 4.4 and the assumption that u_s is non-decreasing in u_0, $\{u_t\}_{t=0}^\infty$ is monotonically non-decreasing. Since $u_s(u_{FB}) = u_{FB}$ for all $s \in \{H, L\}$ (cf. Theorem 4.3), the sequence is bounded from above and, hence, must converge to a limit.

   Define $m(u_0) = max_s u_s(u_0)$. Since $\omega \in \mathcal{A}^c$, for all $T \in I\!N$ there exists $t > T$ such that $u_t > m(u_T)$. For $T \to \infty$, $u^* \geq m(u^*)$. From the definition of $m(\cdot)$ and Lemma 4.1, we have $u_0 < m(u_0)$ for all $u_0 \in [\underline{V}, u_{FB})$. Hence, $\lim_{t\to\infty} u_t = u_{FB}$ for $\omega \in \mathcal{A}^c$.

   The argument for $u_0 \in (u^{FB}, \overline{V}]$ is analogous.

2. If $u_s(\bar{u}) = \bar{u}$ for all $s \in \{H, L\}$, the result follows by an argument analogous to the one given above. By Lemma 4.4, for all $u_0 < \bar{u}$, $u_H(u_0) > u_0$ and $u_L(u_0) = u_0$. Conversely, for all $u_0 > \bar{u}$, $u_H(u_0) = u_0$ and $u_L(u_0) < u_0$. Continuity of u_s implies then $u_s(\bar{u}) = \bar{u}$ for all s.

□

Appendix B

We give now a more rigorous analysis of the discussion following Theorem 4.7 in Section 4.3. We assume throughout that there is no incentive feasible first-best allocation and show that the distribution of wealth converges weakly to a unique long-run distribution provided that u_s^* is non-decreasing.

Given an optimal contract, the state variable u_0 follows an endogenous Markov process that reflects the policy functions $(\{c_s^*, u_s^*\}_{s=1}^S, d)$ as well as the exogenous Markov chain of shocks w_t. Formally, we can express this Markov process by a transition function Q^*. Let $\mathcal{B}$ be the Borel σ-algebra on the interval $[\underline{V}, \overline{V}]$. Define $Q^* : [\underline{V}, \overline{V}] \times \mathcal{B} \longrightarrow [0, 1]$ by

$$Q^*(u_{t-1}^*, B) = Prob(B | u_{t-1}^*(u_0), \ldots, u_1^*(u_0), u_0) = Prob(B | u_{t-1}) \qquad \text{(B.11)}$$

for all $B \in \mathcal{B}$. Associated with the Markov transition function is the operator T_{Q^*} that maps the space of all bounded, $\mathcal{B}$-measurable, real-valued functions into itself. This operator is formally given by

$$T_{Q^*} f = \int f(u_t^*) Q^*(u_{t-1}^*, du_t^*) = \sum_{s=1}^{S} f(u_s^*(u_{t-1}^*))\pi_s, \tag{B.12}$$

where the function f is any bounded, $\mathcal{B}$-measurable, real-valued function. To prove our result we use the following mixing condition:

Condition B.1. *There exists $\epsilon > 0$ and $T \in I\!N$ such that $Prob(u_T(\underline{V}) \geq \bar{u}) \geq \epsilon$ and $Prob(u_T(\overline{V}) \leq \bar{u}) \geq \epsilon$.*

This condition can be interpreted in our context as follows. Suppose that $u_0 \in \{\underline{V}, \overline{V}\}$; i.e., in period $t = 0$ we have the highest possible degree of inequality. Given Condition B.1, there is a positive probability that the initial inequality between agents is reversed within in a finite number of periods.

Let F_0 be any distribution function over $[\underline{V}, \overline{V}]$. Furthermore, denote by F_t the distribution function for u_t given F_0. We say that the sequence of distribution functions $\{F_t\}_{t=0}^{\infty}$ *converges weakly* to F (or $F_t \Rightarrow F$) if and only if $\lim_{t\to\infty} F_t(u_0) = F(u_0)$ at every continuity point u_0 of F. The next result formally establishes weak convergence of the wealth distribution to a unique invariant distribution by adopting an argument of Kocherlakota [7].

Lemma B.2. *If V is differentiable everywhere, for all $u_0 \in [\underline{V}, \bar{u}]$ there exists $s \in S$ such that $u_s(u_0) > u_0$.*

Proof. For $u_0 \in [\underline{V}, \bar{u}]$ some ex post incentive feasibility constraint for agent 2 must be binding. Otherwise, by Lemma 4.1 agent 2 would have an expected utility strictly lower than u_0.

If $\mu_s^1(u_0) = 0$ and $\mu_s^2(u_0) > 0$ the result follows. If $\mu_s^1(u_0)\mu_s^2(u_0) > 0$ there exists a state s' such that $\mu_{s'}^1(u_0)\mu_{s'}^2(u_0) > 0$. Otherwise, one could replicate the allocation in s' for s making one of the agents strictly better off. Since $-V'(u_0) \leq 1$ for all $u_0 \in [\underline{V}, \bar{u}]$, the result follows then from the symmetry assumption on the endowment process and Lemma 4.1. □

Theorem B.3. *If u_s^* is non-decreasing and V is differentiable everywhere, there exists a unique distribution F such that $F_t \Rightarrow F$ for any initial distribution F_0.*

Proof. By Proposition 3.6, u_s^* is continuous and, hence, the operator T_{Q^*} satisfies the Feller Property. Furthermore, T_{Q^*} is monotone as u_s^* is assumed to be non-decreasing. Since $[\underline{V}, \overline{V}]$ is compact and T_{Q^*} preserves continuity, by Theorem 12.10 of Stokey, Lucas with Prescott [16] there exists an invariant distribution over $[\underline{V}, \overline{V}]$ under the transition function Q^*. Furthermore, by Theorem 12.12 of Stokey, Lucas with Prescott [16], the invariant distribution is unique and weak convergence from any initial distribution occurs if T_{Q^*} is monotone and if Condition B.1 for the Markov transition function Q^* is fulfilled. To show that Condition B.1 is satisfied, define $m(u_0) = max_s u_s(u_0)$. Define further a sequence $\{w_n\}_{n=1}^{\infty}$ recursively by setting $w_n = m(w_{n-1})$ where $w_0 = \underline{V}$.
Suppose there does not exist $N \in I\!N$ such that $w_N \geq \bar{u}$. Since the sequence is non-decreasing and bounded from above by $\bar{u}$, it must converge to a limit $\bar{w} \leq \bar{u}$. Since m is a continuous function, $m(\bar{w}) = \bar{w}$. By Lemma B.2, we have $m(u_0) > u_0$ for all $u_0 \in [\underline{V}, \bar{u}]$. A contradiction. □

References

[1] F. Alvarez and U. Jermann, Efficiency, equilibrium, and asset pricing with risk of default, *Econometrica* **68** (2000), 775-797.

[2] G. Debreu, *Theory of Value*, Cowles Foundation, Yale Univ. Press, New Haven, 1959.

[3] A. Dixit, On modes of economic governance, *Econometrica* **71** (2003), 449-483.

[4] C. Gauthier, M. Poitevin and P. González, Ex ante payments in self-enforcing risk-sharing contracts, *J. Econ. Theory* **76** (1997), 106-144.

[5] G. Genicot and D. Ray, Group formation in risk-sharing arrangements, *Rev. Econ. Stud.* **70** (2003), 87-113.

[6] T. Kehoe and D. Levine, Debt-constrained asset markets, *Rev. Econ. Stud.* **60** (1993), 865-888.

[7] N. Kocherlakota, Implications of efficient risk sharing without commitment, *Rev. Econ. Stud.* **63** (1996), 595-609.

[8] T. Koeppl (2002), *Risk sharing with endogenous enforcement*, Unpublished PhD Thesis, University of Minnesota.

[9] S. Krasa and A. Villamil, Optimal contracts when enforcement is a decision variable, *Econometrica* **68** (2000), 119-134.

[10] E. Ligon, P. Thomas and T. Worrall, Mutual insurance, individual savings, and limited commitment, *Rev. Econ. Dynam.* **3** (2000), 216-246.

[11] E. Ligon, P. Thomas and T. Worrall, Informal insurance agreements with limited commitment: theory and evidence from village economies, *Rev. Econ. Stud.* **69** (2002), 209-244.

[12] C. Phelan, Repeated moral hazard and One-Sided Commitment, *J. Econ. Theory* **66** (1995), 488-506.

[13] G. Ramey and J. Watson, Contractual Intermediaries, *J. Law, Econ., Organ.* **18** (2002), 362-384.

[14] A. Rustichini, Dynamic Programming Solution of Incentive Constrained Problems, *J. Econ. Theory* **78** (1998), 329-354.

[15] S. Spear and S. Srivastava, On Repeated Moral Hazard with Discounting, *Rev. Econ. Stud.* **54** (1987), 599-617.

[16] N. Stokey and R. Lucas with E. Prescott, *Recursive Methods in Economic Dynamics*, Harvard Univ. Press, Cambridge, 1989.

[17] J. Thomas and T. Worrall, Self-Enforcing Wage Contracts, *Rev. Econ. Stud.* **55** (1988), 541-554.

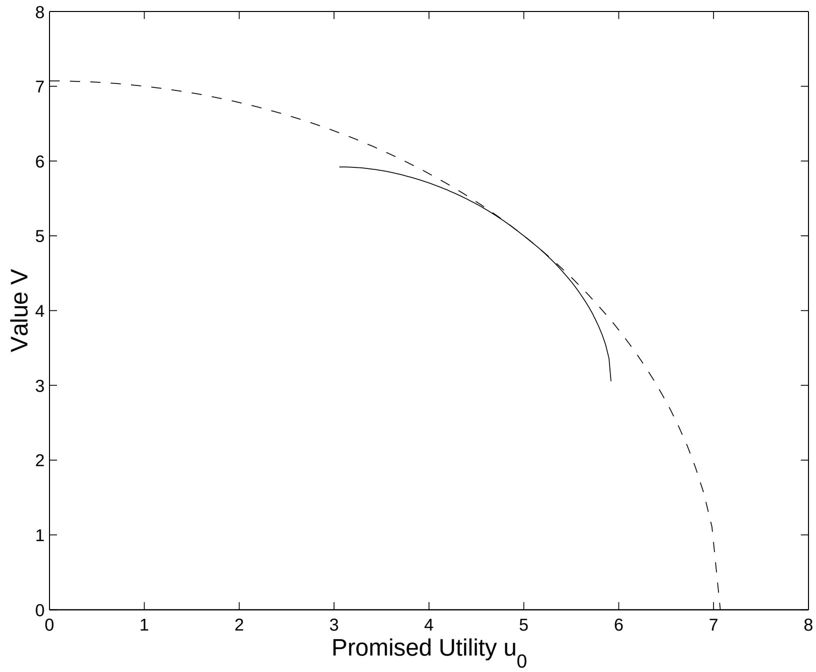

Figure 1: Value Function

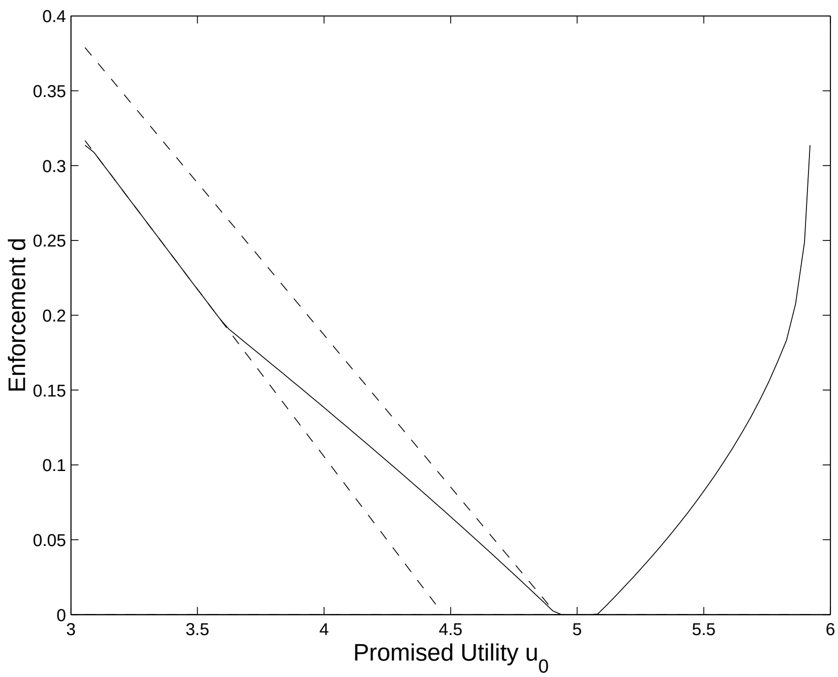

Figure 2: Level of Punishment

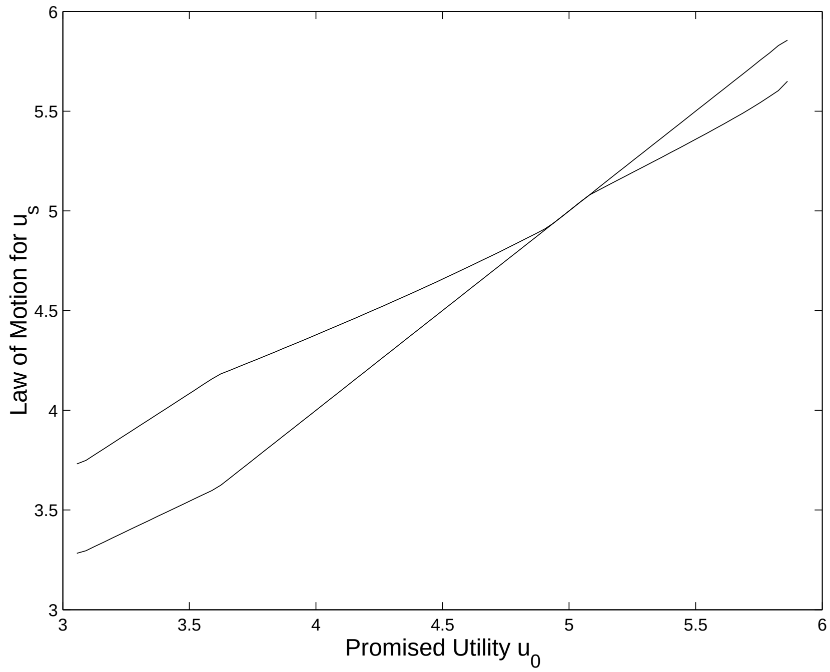

Figure 3: Law of Motion for u_0

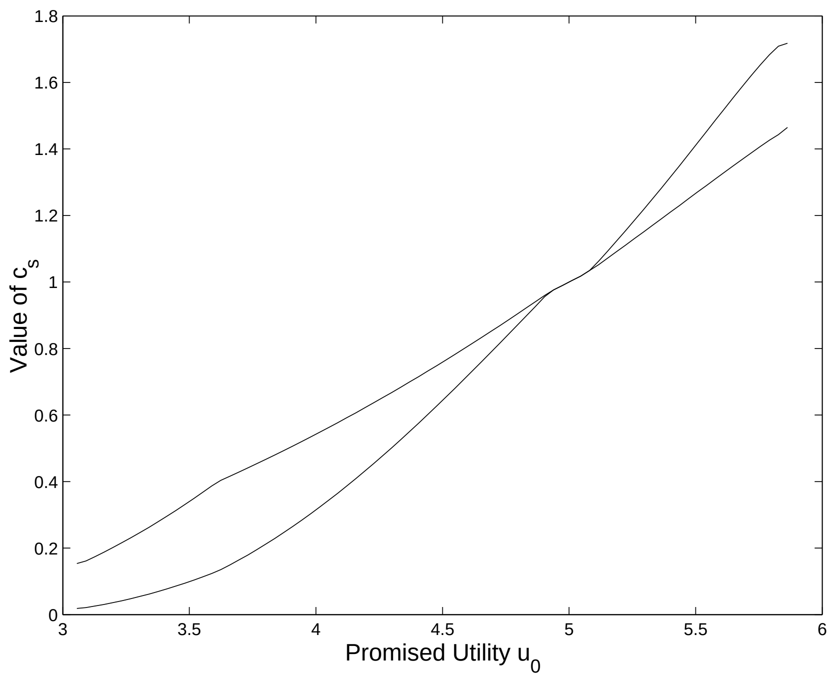

Figure 4: Consumption Levels

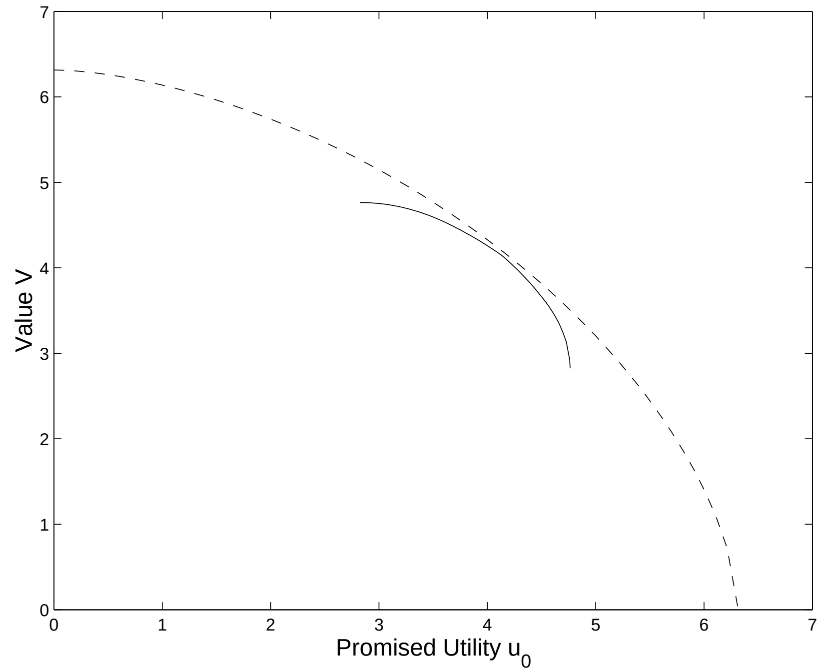

Figure 5: Value Function

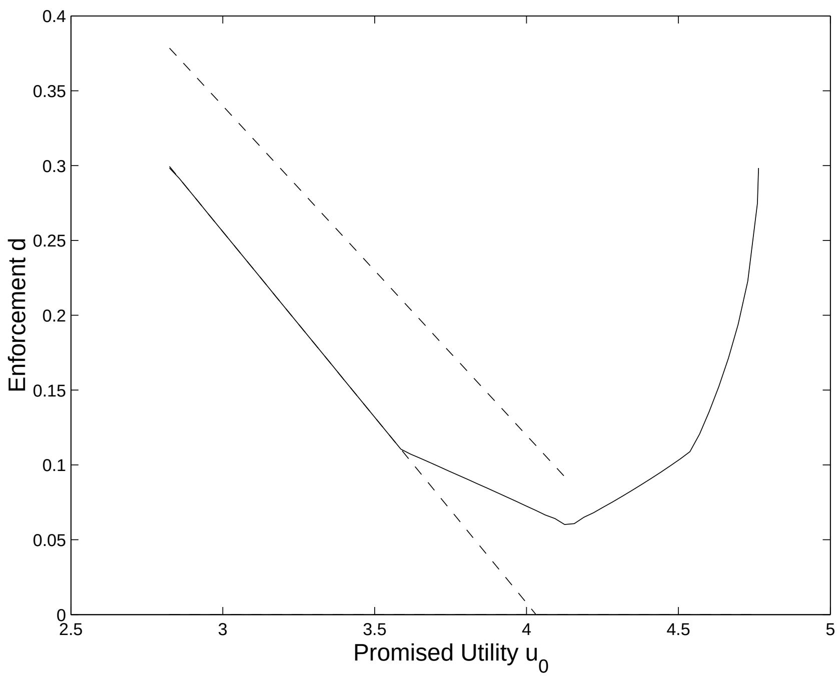

Figure 6: Level of Punishment

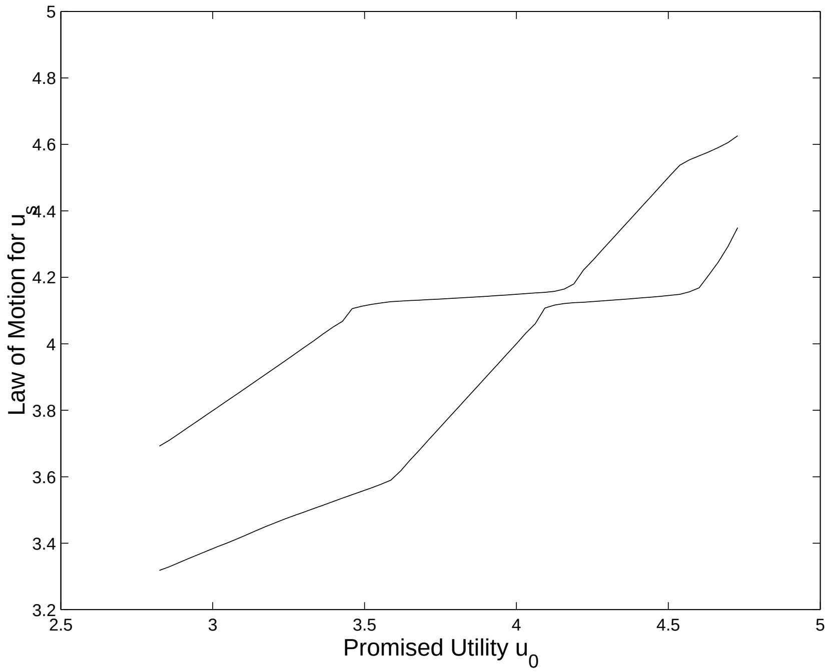

Figure 7: Law of Motion for u_0

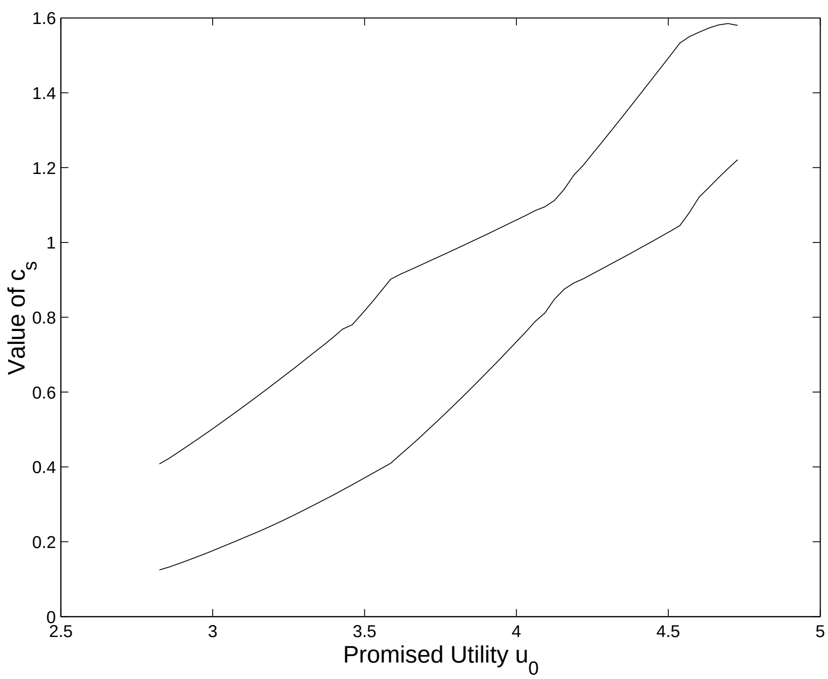

Figure 8: Consumption Levels