Infinitely Repeated Games with Public Monitoring and Monetary Transfers

Susanne Goldlücke* and Sebastian Kranz†

This Version, October 2011‡

Abstract

This paper studies infinitely repeated games with imperfect public monitoring and the possibility of monetary transfers. It is shown that all public perfect equilibrium payoffs can be implemented with a simple class of stationary equilibria that use stick-and-carrot punishments. A fast algorithm is developed that exactly computes the set of pure strategies equilibrium payoffs for all discount factors.

JEL-Codes: C73, D82

Keywords: repeated games, computational methods, monetary transfers, imperfect public monitoring

*Department of Economics, University of Mannheim, L7, 3-5, 68131 Mannheim, Germany. Phone: +49 621 181 3072 Email: goldluecke@uni-mannheim.de

†Department of Economics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany. Email: skranz@uni-bonn.de. Parts of the paper have been written, while I was visiting the Cowles Foundation and Yale University with a research grant from the DFG.

‡Financial support from Deutsche Forschungsgemeinschaft through SFB-TR 15 is gratefully acknowledged. We would like to thank a referee and an associate editor as well as Dirk Bergemann, Mehmet Ekmekci, Paul Heidhues, Johannes Hörner, Jon Levin, Larry Samuelson, Philipp Strack, Juuso Välimäki and Joel Watson for very helpful comments and discussions.

1 Introduction

The theory of infinitely repeated games is used to address a wide range of topics in economics and social sciences, including long-term (financial) contracting, international agreements, employment relations, and cartels. Results that help to find equilibria in these games and numerical procedures to quickly calculate the payoff set in examples are therefore of great importance. In this paper, we study general infinitely repeated games with arbitrary discount factors, imperfect public monitoring, and observable monetary transfers. We first show that all public perfect equilibrium (PPE) payoffs can be implemented with a simple class of stationary equilibria that use stick-and-carrot punishments, in which a deviation from a required monetary transfer is punished by playing a punishment action profile for exactly one period. Based on this result, we develop an algorithm that allows quick computation of optimal stationary equilibria and the set of pure strategy equilibrium payoffs for all discount factors.[1] A variation of the algorithm computes inner approximations of the set of mixed strategy PPE payoffs.

Our results contribute to two streams of literature. First, they contribute to the existing literature that studies repeated games with monetary transfers. Examples include Levin [16, 17] and Malcomson and MacLeod [20] on principal agent games and employment relations, Doornik [6] and Rayo [21] on partnerships and team production, Atkeson [4] and Kletzer and Wright [12] on sovereign lending, Klimenko, Ramey and Watson [13] on international trade agreements, or Harrington and Skrzypacz [10] on cartels that use monetary transfers. The literature typically restricts attention to stationary equilibria in which a single action profile is repeated in every period and deviations from a required payment are punished by reverting to a Nash equilibrium of the stage game. Levin [17] shows that such equilibria can indeed implement every Pareto efficient payoff in principal agent games, and Rayo [21] extends the result to a class of team production problems with multiple agents that commonly observe signals that are independently distributed when conditioning on the played action profile. Our paper provides a characterization for general games. This includes cases in which infinite Nash reversion is not an optimal punishment and cases with signals that are not conditionally

[1] A software package, programmed in R, that implements the algorithm is available on the second author's website http://www.wiwi.uni-bonn.de/kranz/software.htm

For a description of the software and several examples, see Kranz [14].

independently distributed.

Second, our paper contributes to the literature that develops general methods to compute the set of pure strategy equilibrium payoffs for repeated games with arbitrary discount factors. Abreu, Pearce and Stacchetti ([3], henceforth APS), develop a conceptual algorithm to compute the payoff sets for repeated games with imperfect public monitoring and arbitrary discount factors. It is based on the repeated application of a contraction operator on a set of candidate payoffs. Yet, one cannot directly implement this algorithm on a computer because the structure of equilibrium payoff sets can become extremely complicated as is nicely illustrated for a prisoners' dilemma game by Mailath, Obara, and Sekiguchi [18]. Judd, Yeltekin and Conklin [11] analyze games of perfect monitoring and assume a public randomization device, which guarantees that payoff sets are convex polyhedrons. They develop a method to compute upper and lower approximations for the set of pure strategy subgame perfect equilibrium payoffs and to construct strategy profiles that can support payoffs from the lower approximation. Their method is still limited in so far that finding fine approximations for the equilibrium payoff sets for several discount factors remains computationally expensive and it is restricted to games with perfect monitoring.

Allowing for monetary transfers, we develop a much quicker algorithm that can also be applied to games with imperfect public monitoring. We show that a single number can describe all the information contained in a candidate set of equilibrium payoffs and the discount factor that is relevant to determine which action profiles and payoffs can be implemented. This number has a natural interpretation as the totally available liquidity in a static game with enforceable payments and exogenous liquidity constraints. Re-optimization techniques for linear programming allow to quickly determine the implementable payoffs for all possible levels of liquidity. One implication is that our algorithm directly computes the payoff sets for the whole interval $[0, 1)$ of possible discount factors.

The algorithm is particularly powerful in settings in which closed form solutions for the static problems with exogenous liquidity constraints exist. One such case are games with perfect monitoring, as long as attention is restricted to pure strategy equilibria. To compute the sets of pure strategy SPE payoffs for all discount factors under perfect monitoring, it essentially suffices to calculate stage-game best-reply payoffs and sort the stage game action profiles.[2] For the characterization, one

[2]See also Kranz and Ohlendorf [15], where we derive a related result for two player games with

can restrict attention to stationary equilibria whose continuation equilibria before every payment stage, except possibly the very first payment stage, are all Pareto efficient. We exemplify in Appendix A that this result does not extend to mixed strategy equilibria in games with perfect monitoring. A positive probability of inefficient continuation play, e.g. in the form of money burning, before payment stages on the equilibrium path can be required for optimal mixed strategy stationary equilibria.

For games with imperfect public monitoring, there also sometimes exist analytical closed-form solutions for the static problems with exogenous liquidity constraints. We exemplify how one can obtain closed form solutions for the equilibrium payoff sets of the repeated game using a noisy prisoners' dilemma game without conditional independent signal distribution. If no closed-form solutions can be obtained, one can compute the sets of equilibrium payoffs by solving a series of linear optimization problems.

For the largest part of the paper, we assume that players can burn money, i.e. make transfers to a non-involved third party. Money burning is a very explicit way of generating inefficient continuation play, which can be necessary in optimal equilibria after certain signals. Other forms of inefficient continuation play can, of course, serve a similar function. To better understand the role of money burning, we also characterize the payoff set in repeated games in which players do not burn money but have access to a public correlation device. In this framework, every equilibrium payoff can be implemented by a modification of stationary equilibria: with some probability, which can depend on the realized signal, there will be a transition to a collective punishment state. We show how the equilibrium payoff set for the case without money burning can be computed by considering stationary equilibria that allow for money burning but satisfy an additional constraint on the maximal amount of money burning. In general, the set of equilibrium payoffs can shrink if money burning is not possible. If, however, the stage game has a Nash equilibrium that gives each player her min-max payoff, the possibility of money burning does not enlarge the equilibrium payoff set of the repeated game. For pure

perfect monitoring in order to study renegotiation-proofness. It is also interesting to compare our results to Cronshaw and Luenberger [5], who provide a characterization of the set of strongly symmetric pure strategy subgame perfect equilibria for repeated symmetric games with perfect monitoring and a public randomization device. While in their set-up strong symmetry allows a simple characterization of the payoff set, in our set-up monetary tranfers allow a characterization that is almost as simple.

strategy equilibria in games with perfect monitoring, money burning can only be necessary to implement a Pareto dominated equilibrium payoff

The remainder of the paper is organized as follows: Section 2 describes the model and stationary strategy profiles. Section 3 derives the main results. In Section 4, we show how the results simplify for games with perfect monitoring and illustrate the resulting algorithm with a simple Cournot game. Section 5 illustrates for a noisy prisoners' dilemma game how closed-form analytical solutions can be obtained for games with imperfect public monitoring. In Section 6, we explore the case without money burning. Section 7 briefly concludes. Proofs are relegated to the Appendix.

2 Model and Stationary Strategy Profiles

2.1 The game

We consider an infinitely repeated n-player game with imperfect public monitoring and common discount factor $\delta \in [0,1)$. The timing in each period is as follows: at the beginning of a period, there is a payment stage in which the players have the opportunity to make non-negative monetary transfers to each other or to burn money. In a subsequent action stage, the players play a simultaneous move stage game, and then there is again a payment stage in which they can make monetary transfers.[3]

The stage game played in the action stage has the following structure. Each player i has a finite pure action space A_i.[4] The set of stage game action profiles is given by $A = A_1 \times ... \times A_n$. We denote by $a \in A$ a pure action profile and by $\alpha \in \triangle A_1 \times ... \times \triangle A_n$ a mixed action profile of the stage game. The support of a mixed action α_i of player i, i.e., the set of pure actions that player i plays with positive probability, is denoted by $supp(\alpha_i)$.

After an action profile $a \in A$ is chosen, nature draws a commonly observed signal y from a finite signal space Y. The probability distribution of signals depends on

[3] That we allow two payment stages emphasizes that players can make transfers at any point in the game, and it simplifies some formulae. However, the set of equilibrium payoffs stays the same if payments can be made only at the beginning of a period.

[4] Many results for pure strategy equilibria extend to action spaces that are compact subsets of $\mathbb{R}^m$.

the selected action profile a, and is given by a function $\phi(y|a)$ with

$$\phi(y|a) \geq 0 \text{ for all } y \in Y, a \in A,$$
$$\sum_{y \in Y} \phi(y|a) = 1 \text{ for all } a \in A.$$

Player i's stage game payoff depends only on the signal y and on her action a_i and is given by a function $\widehat{g}_i : Y \times A_i \to \mathbb{R}$. We denote by $g_i(\alpha)$ player i's expected payoff given a mixed action profile α. The joint payoff from an action profile α is denoted by

$$G(\alpha) = \sum_{i=1}^{n} g_i(\alpha).$$

In contrast to the action choices, we assume that all transfers are commonly observable. All players choose their monetary transfers simultaneously. We also allow the players to burn money (one can think of the possibility to give money to charity or any other non-interested third party). To have a bounded action space, we assume for convenience that there exists an upper bound on a player's transfers. However, this upper bound shall be sufficiently large, so that we essentially consider a situation of unlimited liability. Players are risk-neutral and utility is linear in money and stage game payoffs. Thus, a player's payoff in a period in which action profile a has been played and signal y has been realized is given by $\widehat{g}_i(y, a_i)$ minus the sum of the net payments that player i has made in the two payment stages.

Unless stated otherwise, our results apply for the case that attention is restricted to pure strategy equilibria and also for the case that players can mix over actions. The variable $\mathcal{A}$ shall denote the set of pure action profiles A in the former case and the set of mixed action profiles $\triangle A_1 \times ... \times \triangle A_n$ in the latter. We assume that the stage game has a Nash equilibrium in $\mathcal{A}$.

A public history h of the repeated game is a list of all monetary transfers and public signals that have occurred before a given point in time. A public strategy σ_i of player i in the repeated game maps every public history that ends before the action stage into an action $\alpha_i \in \mathcal{A}_i$, and every public history that ends before a payment stage into a vector of monetary transfers. A public perfect equilibrium is a profile of public strategies that constitutes mutual best replies after every public history. We will restrict attention to public perfect equilibria.[5]

[5]If attention is restricted to pure strategy equilibria, the restriction to public perfect equilibria

Payoffs and continuation payoffs of the repeated game are defined as average discounted expected payoffs, i.e. as the expected discounted sum of future payoffs multiplied by $(1-\delta)$. We denote by $u^0(\sigma)$ the vector of payoffs in the repeated game given a strategy profile σ.

2.2 Stationary strategy profiles

In this section, we introduce a class of stationary strategy profiles that allow a simple characterization of PPE payoffs for every discount factor. These stationary strategy profiles have the feature that the same action profile is played in every period on the equilibrium path and punishments have a simple stick-and-carrot structure.

A repeated game strategy specifies gross amounts $\tilde{p}_{ij} \geq 0$ that player i transfers to player j and an amount $\tilde{p}_{i0}$ that player i burns. For convenience, however, we will describe all monetary transfers in stationary strategy profiles by the net payments players make. It is straightforward that for any vector of net payments $p = (p_1, ..., p_n) \in \mathbb{R}^n$ with $\sum_{i=1}^n p_i \geq 0$, one can find corresponding gross transfers $\tilde{p}_{ij}$ generating these net payments, i.e.

$$p_i = \tilde{p}_{i0} + \sum_{j \neq i}^{n} \left(\tilde{p}_{ij} - \tilde{p}_{ji}\right), \tag{1}$$

while there is no player who simultaneously makes and receives gross transfers.[6]

A stationary strategy profile is characterized by $n+2$ states. Play starts in the *up-front payment state*, in which players are required to make up-front payments

is without loss of generality. The set of pure strategy PPE payoffs is the same as the set of pure strategy sequential equilibrium payoffs (see APS). This equivalence does not hold for mixed strategies, however.

[6]Concretely, we can specify gross transfers as follows. We denote by $I_P = \{i|p_i > 0\}$ the set of net payers and by $I_R = \{i|p_i \leq 0\} \cup \{0\}$ the set of net receivers including the sink for burned money. For any receiver $j \in I_R$, we denote by

$$s_j = \frac{|p_j|}{\sum_{j \in I_R} |p_j|}$$

the share she receives from the total amount that is transferred or burned and assume that each net payer distributes her gross transfers according to these proportions

$$\tilde{p}_{ij} = \begin{cases} s_j p_i & \text{if } i \in I_P \text{ and } j \in I_R \\ 0 & \text{otherwise.} \end{cases}$$

p^0. Afterward, play can be in one of $n+1$ states, which we index by $k \in K = \{e, 1, 2, ..., n\}$. We call the state $k = e$ the *equilibrium state* and $k = i \in \{1, ..., n\}$ the *punishment state* of player i. A stationary strategy profile specifies for each state $k \in K$ an action profile $\alpha^k \in \mathcal{A}$ that will be played in the action stage. Furthermore, it specifies for each state $k \in K$ a payment function $p^k : Y \to \mathbb{R}^n$ that maps the signal y from the preceding action stage into a required vector of payments. Payments in the beginning of the period only occur in the up-front payment state in the first period.

The state transitions are as follows: If no player unilaterally deviates from a required payment, the new state becomes the equilibrium state: $k = e$. If player i unilaterally deviates from a required payment, the new state becomes the punishment state of player i, i.e. $k = i$. In all other situations the state does not change.

A stationary strategy profile σ is completely characterized by a vector of up-front payments p^0, its *action plan* $(\alpha^k)_{k \in K}$ that specifies one action profile for every state k and its *payment plan* $(p^k)_{k \in K}$ that specifies a payment function for every state k. For a given discount factor δ, we call a stationary strategy profile σ a *stationary equilibrium* if σ constitutes a public perfect equilibrium of the repeated game. We denote by $(\alpha^k, p^k)_{k \in K}$ a stationary strategy-profile without up-front payments and by Σ^0 the set of stationary equilibria without up-front payments.

The following definitions are useful for the characterization of stationary equilibria. For any payment function p, we let $E[p_i|\alpha]$ denote the expected payments of player i given action profile α; expectations are taken over the signal distribution and mixing probabilities. For any stationary strategy profile player i's expected repeated game payoff at the beginning of a period in the equilibrium state is

$$u_i(\sigma) = g_i(\alpha^e) - E[p_i^e|\alpha^e].$$

Whenever the equilibrium in question is clear from the context, we will suppress the dependence on σ. Similarly, the joint equilibrium state payoff is given by

$$U(\sigma) = G(\alpha^e) - \sum_{i=1}^{n} E[p_i^e|\alpha^e],$$

where the sum on the right hand side denotes the expected amount of money that is burned on the equilibrium path. Player i's continuation payoff at the beginning of his punishment state is denoted by

$$v_i(\sigma) = (1-\delta)(g_i(\alpha^i) - E[p_i^i|\alpha^i]) + \delta u_i.$$

We call v_i player i's *punishment payoff.* We denote the sum of punishment payoffs by

$$V(\sigma) = \sum_{i=1}^{n} v_i.$$

3 Main Results

3.1 Characterization with Optimal Stationary Equilibria

Using the one shot deviation principle, we now establish the constraints that a stationary strategy profile without up-front payments $\sigma = (\alpha^k, p^k)_k$ has to satisfy to be a stationary equilibrium. There are three types of constraints, which we call payment constraints, budget constraints, and action constraints.

Payment constraints Given that player i has an equilibrium payoff of u_i and a punishment payoff of v_i, he is never willing to make a higher payment than

$$p_i^{\max} \equiv \frac{\delta}{1-\delta}(u_i - v_i).$$

A stationary equilibrium thus must satisfy the following payment constraints for all states $k \in K$:

$$p_i^k(y) \leq p_i^{\max} \text{ for all } i, y. \tag{PC-k}$$

Budget constraints Even though players can burn money, they cannot get any outside funding. In every state k, the following budget constraints must therefore be satisfied:

$$\sum_{i=1}^{n} p_i^k(y) \geq 0 \;\text{ for all } y. \tag{BC-k}$$

Action constraints There are no incentives to deviate from any pure action in the support of the mixed action profiles for state $k \in K$ if and only if

$$g_i(a_i, \alpha_{-i}^k) - E[p_i^k | a_i, \alpha_{-i}^k] \geq g_i(\widehat{a}_i, \alpha_{-i}^k) - E[p_i^k | \widehat{a}_i, \alpha_{-i}^k], \tag{AC-k}$$
$$\text{for all } i, a_i \in supp(\alpha_i^k) \text{ and } \widehat{a}_i \in A_i.$$

These action constraints imply that player i must have the same expected payoff for all pure actions in the support of α_i^k.

Up-front payments Next, we describe how the possibility of up-front payments transforms the set of feasible payoffs. Up-front payments are incentive compatible if they do not exceed $p_i^{\max} = \frac{\delta}{1-\delta}(u_i - v_i)$ for any player. Incentive compatible up-front payments allow any distribution of the joint equilibrium payoff that guarantees every player at least his punishment payoff. This leads to the following straightforward result:

Proposition 1 *If there exists a stationary equilibrium σ with joint equilibrium payoffs U and punishment payoffs $v = (v_1, ..., v_n)$ then every payoff in the simplex*

$$\{u^0 \in \mathbb{R}^n \mid \sum_{i=1}^{n} u_i^0 \leq U \text{ and } u_i^0 \geq v_i \text{ for all } i\} \tag{2}$$

can be achieved by some stationary equilibrium that differs from σ only in the up-front payments.

Optimal stationary equilibria We will show that every public perfect equilibrium can be implemented by a set of optimal stationary equilibria, which are defined as follows:

Definition 1 *We say a stationary equilibrium σ is optimal if there does not exist another stationary equilibrium that implements a higher difference $U(\sigma) - V(\sigma)$ of joint equilibrium state payoffs and joint punishment payoffs. An action plan and payment plan are optimal if they are the action plan and payment plans of an optimal stationary equilibrium.*

By maximizing the difference $U - V$, optimal stationary equilibria simultaneously achieve the highest possible joint payoffs U and the lowest punishment payoffs v_i for every player i that can be achieved with any stationary equilibrium. Formally this result will be an implication of Theorem 1, but we want to provide an intuition here. Looking at the constraints for a stationary equilibrium, one finds that across states the constraints are solely linked by the right hand sides of the payment constraints $\frac{\delta}{1-\delta}(u_i - v_i)$, which depend on both equilibrium state and punishment payoffs. This means that if one can implement a lower punishment payoff v_i for some player i, the only effect on the constraints for the other states $k \neq i$ is that the payment constraints of player i are relaxed. This means that similar to Abreu's [1] simple strategy profiles for games with perfect monitoring, a harsher punishment of player i facilitates the implementation of harsher punishments for other players and the implementation of more efficient equilibrium play.

As will be seen more clearly in Section 3.2, a related logic applies for the equilibrium state: if a change in α^e allows to implement a higher joint equilibrium state payoff U, it is always possible to implement harsher punishments. A rough intuition for why only the joint equilibrium payoffs matter, is that payments functions can always be structured such that the payment constraints between players are smoothed.

We can now state our main result.

Theorem 1 *All public perfect equilibrium payoffs can be implemented by optimal stationary equilibria that only differ in their up-front payments.*

In the proof of Theorem 1, the optimal stationary equilibrium is derived from PPE equilibria that implement the highest joint PPE payoff and the lowest PPE payoffs for each player, respectively. An optimal action plan is formed from the first period action profiles of these equilibria. A payment plan is chosen such that continuation payoffs in the stationary equilibrium match the continuation payoffs of these equilibria. Because the proof is done without knowledge of the structure of these PPE equilibria, it is not constructive and does not show how optimal action and payment plans can be found. The next subsection addresses this problem.

3.2 Finding Optimal Stationary Equilibria

A brute force method

Optimal action and payment plans solve the following optimization problem:

$$\max_{(\alpha^k)_{k\in K},(p^k)_{k\in K}} U - V$$
$$\text{s.t. (PC-k), (BC-k), (AC-k), for all states } k \in K.$$

If one fixes an action plan $(\alpha^k)_k$, one can find a corresponding payment plan that maximizes $U - V$ by solving the linear program

$$\max_{\{p^k\}_{k\in K}} U - V \tag{LP-PP}$$
$$\text{s.t. (PC-k), (BC-k), (AC-k), for all states } k \in K.$$

If we restrict attention to pure strategy equilibria, a brute force method to find optimal stationary equilibria is to solve this optimization problem for all possible (pure) action plans $(a^k)_k \in A^{n+1}$. Unfortunately, the number of action plans can quickly grow very large, which makes this method ill suited for larger stage games.[7] Another drawback is that one has to repeat the computation for every discount factor of interest.

Basic idea of a faster algorithm We will now develop results that yield an algorithm, which can handle quite large stage games and will directly compute the pure strategy equilibrium payoff sets for all discount factors $\delta \in [0, 1)$. For mixed strategy equilibria the results do not yield an exact algorithm but suggest approximation methods. The basic idea is that one first solves a series of static problems with enforceable payments for each state and afterwards combines these solutions to determine optimal stationary equilibria for all discount factors.

The static problem Consider the following static problem. The stage game is played once and there exist enforceable contracts that specify for each player $i = 1, .., .n$ and every signal $y \in Y$ a vector of gross monetary transfers to other players and an amount of money burning. From an incentive perspective, only net payments are relevant. We therefore describe an enforceable contract by a payment function $p(.)$ that specifies the net payments $p_i(y)$ of player i if signal y realizes.

The possible payments that player i can make shall be bounded by an exogenously given liquidity constraint $\lambda_i L \geq 0$, with $L \geq 0$, $\lambda_i \geq 0$ and $\sum_{i=1}^{n} \lambda_i = 1$. This means the totally available liquidity across all players is given by L and λ denotes the liquidity distribution.

We say that an action profile $\alpha \in \mathcal{A}$ can be implemented with a payment function $p(.)$ in the static problem given liquidity allocation λL, if the following payment, budget and action constraints hold:

$$p_i(y) \leq \lambda_i L \quad \text{for all } i, y, \tag{PC}$$

$$\sum_{i=1}^{n} p_i(y) \geq 0 \text{ for all } y, \tag{BC}$$

$$g_i(a_i, \alpha_{-i}) - E[p_i|a_i, \alpha_{-i}] \geq g_i(\widehat{a}_i, \alpha_{-i}) - E[p_i|\widehat{a}_i, \alpha_{-i}], \tag{AC}$$
$$\text{for all } i, a_i \in supp(\alpha_i) \text{ and } \widehat{a}_i \in A_i.$$

[7] For example, in a discretized Cournot duopoly in which each output profiles are chosen from a 100 x 100 grid, there are ten thousand action profiles but one trillion different action plans.

Whether an action profile α can be implemented with some payment function and how much money needs to be burned, does not depend on the liquidity distribution λ, but only on the total liquidity L. More precisely, we have the following straightforward result:

Lemma 1 *If the payment function p can implement an action profile α for the liquidity allocation λL then the payment function $\widetilde{p}$ with*

$$\widetilde{p}_i(y) = p_i(y) + \left(\widetilde{\lambda}_i - \lambda_i\right) L \tag{3}$$

can implement α for the liquidity allocation $\widetilde{\lambda} L$.

Liquidity Requirement We define the liquidity requirement $\mathcal{L}(\alpha)$ of an action profile α as the minimum total liquidity L that is necessary to implement α in the static problem.[8] Because of Lemma 1, the liquidity requirement is independent of the actual liquidity distribution λ. It is given as the solution to the following linear program:

$$\mathcal{L}(\alpha) = \min_{p, L \geq 0} L \text{ s.t. (PC), (BC), (AC).} \tag{LP-L}$$

To find closed-form solutions for $\mathcal{L}(\alpha)$ in specific examples, it will often be convenient to solve (LP-L) with a liquidity distribution that gives all liquidity to a single player or distributes liquidity equally across players. The liquidity requirement of an action profile α is 0 if and only if α is a Nash equilibrium of the stage game.

Equilibrium state For a given value of total liquidity $L \geq \mathcal{L}(\alpha)$, we denote by $U^e(L, \alpha)$ the maximum expected joint payoff that can be implemented with action profile α:

$$U^e(L, \alpha) = \max_p \left(G(\alpha) - \sum_{i=1}^{n} E[p_i|\alpha] \right) \text{ s.t. (PC),(BC), (AC).} \tag{LP-e}$$

Lemma 1 implies that the solution to the linear program (LP-e) is independent of the chosen liquidity distribution λ. Observe that $U^e(L, \alpha)$ is bounded, weakly increasing and concave in L, and also piece-wise linear with a finite number of kinks.[9] Appendix B explains a method that exploits these attributes in order to quickly compute $U^e(L, \alpha)$ for all values of L.

[8] Note that it depends on the signal structure whether an action profile can be implemented at all. For an action profile that cannot be implemented, the required liquidity would be infinite.

[9] That $U^e(L, \alpha)$ is weakly increasing and bounded is obvious. Concavity and piece-wise linearity follows from standard results on linear optimization.

Punishment states We now define a punishment payoff for player i in the static problem. For any given action profile α, liquidity $L \geq \mathcal{L}(\alpha)$ and some arbitrary liquidity distribution λ, we define

$$v^i(L,\alpha) = \min_p \left(g_i(\alpha) + \lambda_i L - E[p_i|\alpha]\right) \text{ s.t. (PC),(BC), (AC).} \tag{LP-i}$$

Again, because of Lemma 1, $v^i(L,\alpha)$ is independent of the liquidity distribution λ. Note that $v^i(L,\alpha)$ is the lowest expected payoff that can be imposed on player i in the static problem if no liquidity is given to player i.

We denote the stage game best reply payoff or *cheating* payoff of player i by

$$c_i(\alpha) = \max_{\widehat{a}_i \in A_i} g_i(\widehat{a}_i, \alpha_{-i}).$$

Lemma 2 *It holds that $v^i(L,\alpha) \geq c_i(\alpha)$. Furthermore, $v^i(L,\alpha) = c_i(\alpha)$ if $g_i(\alpha) = c_i(\alpha)$.*

This lemma is easiest understood for the case $\lambda_i = 0$. That no punishment payoff below $c_i(\alpha)$ can be implemented is obvious. If all pure actions a_i in the support of α_i are stage game best replies to α_{-i}, then one can always implement player's i stage game payoff $c_i(\alpha)$ by never making any transfer to player i. If some pure action a_i in the support is not a best reply to α_{-i}, then $g_i(\alpha) < c_i(\alpha)$ and player i must receive positive transfers after some signals. These transfers can have the effect that player i's punishment payoff cannot be pushed down to his cheating payoff $c_i(\alpha)$. This can be most clearly seen for a signal distribution with full support, i.e. $\phi(y|a) > 0$ for all $a \in A$ and all $y \in Y$: player i would then receive positive expected payments under every deviation.

Similar to $U^e(L,\alpha)$, the function $v^i(L,\alpha)$ is bounded, weakly decreasing, convex and piece-wise linear in L (with a finite number of kinks); efficient computation techniques are also described in Appendix B.

Relation between static problems and stationary equilibria For a stationary equilibrium σ of the repeated game, it is helpful to think of the sum of the right-hand side of payment constraints,

$$L(\sigma) \equiv \frac{\delta}{1-\delta}\left(U(\sigma) - V(\sigma)\right),$$

as the total liquidity that is endogenously generated by the threat to punish players who fail to make a payment. Similarly, one can think of

$$\lambda_i(\sigma) \equiv \frac{u_i(\sigma) - v_i(\sigma)}{U(\sigma) - V(\sigma)}$$

as the fraction of total liquidity given to player i.

Assume σ is a stationary equilibrium with an action plan $(\alpha^k)_k$. One can then obviously implement a payoff of at least $U(\sigma)$ in the static problem (LP-e) given total liquidity $L(\sigma)$ and action profile α^e. A similar result holds for the punishment payoffs. This means we have $U(\sigma) \leq U^e(L(\sigma), \alpha^e)$ and $v_i(\sigma) \geq v^i(L(\sigma), \alpha^i)$. An implication of Proposition 2 below will be that if σ is an optimal stationary equilibrium, it must be the case that $U(\sigma) = U^e(L(\sigma), \alpha^e)$ and $v_i(\sigma) = v^i(L(\sigma), \alpha^i)$. This result is not obvious for the following reason. If we solve the static problem (LP-e) with total liquidity $L(\sigma)$ and liquidity distribution $\lambda_i(\sigma)$, there is no requirement that the resulting payment function p^e has to satisfy $g(\alpha^e) - p^e = u(\sigma)$. This means a payment plan derived from solving the separate static problems does not necessary lead to the liquidity distribution $\lambda_i(\sigma)$ and there is no guarantee that such a payment plan is consistent with the payment constraints of any stationary equilibrium.

Yet, we will show in Proposition 2 that no matter for which liquidity distributions the static problems have been solved, one can transform the resulting payment plan such that it will be consistent with all constraints of a stationary equilibrium and implement the same joint equilibrium and punishment payoffs as in the static problems.

Optimal action plans Before stating the result, we require a few more definitions in order to link the solutions of the static problems to optimal stationary equilibria. We denote by

$$\overline{\alpha}^e(L) \in \arg \max_{\alpha \in \mathcal{A} | L \geq \mathcal{L}(\alpha)} U^e(L, \alpha)$$

an action profile for which the highest joint payoff can be implemented in a static problem with total liquidity L. The corresponding highest joint payoff is denoted by

$$\bar{U}^e(L) = U^e(L, \overline{\alpha}^e(L)).$$

Similarly, we define for the static problem of player i's punishment state

$$\overline{\alpha}^i(L) = \arg \min_{\alpha \in \mathcal{A} | L \geq \mathcal{L}(\alpha)} v^i(L, \alpha)$$

and

$$\bar{v}^i(L) = v^i(L, \overline{\alpha}^i(L)).$$

From the functions $\bar{U}^e(L)$ and $\bar{v}^i(L)$ we can derive an upper bound on the total liquidity that can be endogenously generated by any stationary equilibrium given discount factor δ. It is implicitly defined by the fixed point condition:

$$\bar{L}(\delta) \equiv \max \left\{ L | L = \frac{\delta}{1-\delta} \left(\bar{U}^e(L) - \sum_{i=1}^{n} \bar{v}^i(L) \right) \right\}. \tag{4}$$

To see that a non-negative $\bar{L}(\delta)$ does always exist, consider the right-hand side, $\frac{\delta}{1-\delta}(\bar{U}^e(L) - \sum_{i=1}^{n} \bar{v}^i(L))$, as a function of L. This function is bounded from below by 0, weakly increasing, and it stays constant at an upper bound once L is sufficiently large. Thus the set of fixed points, in which L equals the right hand side, is nonempty and an application of Tarski's fixed point theorem shows that it has a maximum.

Proposition 2 *Given discount factor δ, there exists an optimal stationary equilibrium σ with action plan $(\overline{\alpha}^k(\bar{L}(\delta)))_k$ joint equilibrium payoffs $U(\sigma) = \bar{U}^e(\bar{L}(\delta))$ and punishment payoffs $v_i = \bar{v}^i(\bar{L}(\delta))$ for every player $i = 1, ..., n$.*

Together with our previous results, Proposition 2 implies:

Corollary 1 *The set of public perfect payoffs for discount factor δ is given by the following simplex*

$$\{u^0 \in \mathbb{R}^n \mid \sum_{i=1}^{n} u_i^0 \leq \bar{U}^e(\bar{L}(\delta)) \text{ and } u_i^0 \geq \bar{v}^i(\bar{L}(\delta))\}. \tag{5}$$

Structure of the algorithm The algorithm to compute the sets of pure strategy PPE payoffs for all discount factors has the following structure.

First, we compute $\bar{U}^e(L)$ for all levels of L. To determine this upper envelope, it is typically not necessary to compute the function $U^e(L, a)$ for all pure action profiles $a \in A$. For example, if the joint equilibrium payoff $G(a)$ of an action profile a is lower than the joint payoff of a stage game Nash equilibrium, a is clearly not an optimal equilibrium state profile and we can dismiss it without any further calculation. In Appendix B, we discuss several heuristics that speed up the calculation of $\bar{U}^e(L)$.

The second step is to compute in a similar fashion the lower envelopes of punishment payoffs $\bar{v}^i(L)$. If the stage game is symmetric, $\bar{v}^i(L)$ will be the same for all players and it suffices to compute it for one player.

The third step is to compute $\bar{L}(\delta)$ for all discount factors. For this step, it is helpful to work with discount rates $r = \frac{1-\delta}{\delta}$ and to define the function

$$r^*(L) = \frac{\bar{U}^e(L) - \sum_{i=1}^{n} \bar{v}^i(L)}{L}, \tag{6}$$

which describes the discount rate for which a stationary equilibrium with action plan $(\bar{a}^k(L))_k$ generates liquidity L.[10] The numerator on the right hand side of (6) is a bounded, piece-wise linear function in L. One can obtain $\bar{L}(\delta)$ for all discount factors $\delta \in (0,1)$, by piece-wise inverting the function $r^*(L)$. We then know optimal action plans $(\overline{\alpha}^k(\bar{L}(\delta)))_k$, and the variables determining the payoff sets $\bar{U}^e(\bar{L}(\delta))$ and $\bar{v}^i(\bar{L}(\delta))$, i.e. we basically have solved the game for all discount factors. We illustrate this procedure in Sections 4 and 5.

Approximating the mixed strategy PPE payoff set

From any finite grid of possible mixed action profiles, we can derive an inner approximation for the set of mixed strategy PPE payoffs. For every mixed action α in the grid, we can exactly evaluate $U^e(L,\alpha)$ and $v_i(L,\alpha)$. These values are independent of the other mixed action profiles included in the grid, because it suffices to check deviations to a pure actions. The evaluated values from the grid generate a lower bound for the envelope $\bar{U}^e(L)$ over all mixed action profiles, and similarly an upper bound for $\bar{v}^i(L)$.[11] These bounds can be used to derive an inner approximation of the mixed strategy PPE payoff sets in a similar way as we derive the exact pure strategy PPE payoff set.

4 Perfect monitoring

With perfect monitoring, the played action profile is perfectly observable by all players. This means that we have a game with perfect monitoring if the signal

[10] If the stage game has a Nash equilibrium in pure strategy, we set $r^*(0) = \infty$, otherwise $r^*(L)$ is not defined for sufficiently low L.

[11] Parametric linear programming methods can be helpful for approximating $\bar{U}^e(L)$ and $\overline{v}_i(L,\alpha)$. For any fixed support of mixing probabilities, a continuous change in the mixing probabilities changes the coefficients of the corresponding linear programs continously. Yet, a change in the support also changes the set of action constraints.

space is equal to the pure action space, i.e. $Y = A$ and the signal distribution is

$$\phi(y|a) = \begin{cases} 1 & \text{if } y = a \\ 0 & \text{if } y \neq a \end{cases}.$$

This section provides a simple algorithm to compute the payoff sets of all pure strategy SPE for games of perfect monitoring. To implement an action profile a in the static problem, one can use a payment function $\hat{p}$ that requires each player i to pay $c_i(a) - g_i(a)$ following any signal (a_i', a_{-i}) with $a_i' \neq a_i$, and to pay nothing otherwise. The liquidity requirement is given by

$$\mathcal{L}(a) = \sum_{i=1}^{n} (c_i(a) - g_i(a)) . \tag{7}$$

That this liquidity suffices to implement a can be seen by considering the liquidity distribution $\lambda_i = \frac{c_i(a)-g_i(a)}{L(a)}$. That this liquidity is necessary follows from summing up the action and payment constraints over all players.

With the payment function $\hat{p}$, no money will be burned on the equilibrium path. Thus, for all $L \geq \mathcal{L}(a)$ we find that the maximal implementable joint payoffs are equal to the joint stage game payoffs:

$$U^e(L, a) = G(a).$$

To calculate the minimal punishment payoffs $v^i(L, a)$ for player i and an action profile a, consider a liquidity distribution λ that gives no liquidity to player i, i.e. $\lambda_i = 0$. It follows from Lemma 1 that a can then be implemented with the payment function $p + \lambda L(a) - (c(a) - g(a))$. We thus find that

$$v^i(L, a) = c_i(a),$$

i.e., player i's minimal punishment payoff is always equal to his stage game cheating payoff under his punishment profile a^i. These closed form solutions simplify the computation of optimal equilibria for all discount factors considerably.

Note that these results do not extend to the case of mixed strategies. In particular, money burning can occur under perfect monitoring once mixed strategies are allowed, as we show with an example in Appendix A. We also show there that money burning cannot be optimal as soon as there is one player who plays a pure strategy.

Illustrating the algorithm for perfect monitoring We now exemplify the perfect monitoring version of the algorithm for a simplified Cournot game taken from [1]. Two firms simultaneously choose either low (L), medium (M), or high (H) output and stage game payoffs are given by the following matrix:

	L	M	H
L	10,10	3,15	3,15
M	15,3	7,7	7,7
H	7,0	5,-4	5,-4

Step 1: The first step is to create a list of candidates for optimal equilibrium action profiles. We order all action profiles $a \in A$ decreasingly in their joint payoff $G(a)$ and break ties by putting action profiles with a lower liquidity requirement $\mathcal{L}(a)$ first. Then we remove all action profiles from the list that do not have a strictly lower liquidity requirement than all earlier action profiles in the list. In the example, we get the following list:[12]

No.	a^e	$G(a^e)$	$\mathcal{L}(a^e)$
1.	(L,L)	20	10
2.	(L,M)	18	4
3.	(M,M)	14	0

Note that if the stage game has at least one Nash equilibrium then the last profile of the list is always the Nash equilibrium with the highest joint payoffs. The table describes the highest implementable joint payoffs $\overline{U}(L)$, which is a step-wise function given by

$$\overline{U}(L) = \begin{cases} 20 & \text{if } 10 \leq L \\ 18 & \text{if } 4 \leq L < 10 \\ 14 & \text{if } L < 4 \end{cases} .$$

Step 2: In a similar way, we create for each punishment state $i = 1, ..., n$ a list of action profiles. We order action profiles increasingly in player i's cheating payoff $c_i(a)$. We break ties by putting those profiles with a lower liquidity requirement $\mathcal{L}(a)$ first. We remove action profiles that do not have a strictly lower liquidity requirement than all earlier action profiles. In the example, we get the following list for the punishment state of player 1:

No.	a^1	$c_1(a^1)$	$\mathcal{L}(a^1)$
1.	(M,H)	0	6
2.	(M,M)	7	0

[12]Instead of (L, M) as second element of the list, we could alternatively pick the profile (M, L).

which describes $\overline{v}_1(L)$ and $\overline{a}^1(L)$. As the stage game is symmetric, the lists of punishment profiles for the other players simply consists of the correspondingly permuted action profiles.

Step 3: The first action profiles in each list form our initial action plan. In the example, we have $(a^e = (L,L), a^1 = (M,H), a^2 = (H,M))$. By construction this is an optimal action plan whenever it can be implemented in a stationary equilibrium. Adapting equation (6), we find that the maximal discount rate for which a stationary equilibrium with this action plan exists, is given by

$$r^*(\max_k \mathcal{L}(a^k)) = \frac{G(a^e) - \sum_{i=1}^n c_i(a^i)}{\max_k \mathcal{L}(a^k)}. \tag{8}$$

In the current example, we find

$$r^* = \frac{20}{\max\{10, 6\}} = 200\%.$$

This corresponds to a critical discount factor of $\delta^* = \frac{1}{1+r^*} = \frac{1}{3}$.[13] Thus, by varying the up-front payments, we can implement for every discount factor $\delta \in (\frac{1}{3}; 1)$ every (weakly) individually rational distribution of the maximum joint stage game payoff of 20 as subgame perfect equilibrium payoff of the repeated game.

It is straightforward that for any finite stage game, the minimal discount factor δ^* for which every individually rational distribution of the maximum joint stage game payoff can be implemented is always strictly below 1. This result is a folk theorem for games with side payments. For games without side payments, it generally only holds true that every feasible and *strictly* individually rational payoff can be implemented for sufficiently large discount factors. Moreover, in games with more than 2 players, the folk theorem without side payments only holds under certain regularity conditions, e.g. if the NEU condition [2] holds. In our setting, we have a folk theorem even for stage games that violate the NEU condition; the reason being that once transfers are allowed the condition is always satisfied.

Step 4: In the next step, we replace the action profile a^k that has the highest liquidity requirement $\mathcal{L}(a^k)$ by the next action profile in the list for state k. If

[13] If no transfers are allowed then clearly (L, L) can also not be sustained on the equilibrium path if $\delta < \frac{1}{3}$. Using grim-trigger strategies, one can sustain it whenever $\delta \geq \frac{5}{8}$ and [1] shows that (L,L) can be sustained for all $\delta \geq \frac{4}{7}$. A stick-and-carrot strategy profile with the action plan above but no transfers constitutes a subgame perfect equilibrium even for all $\delta \geq \frac{5}{14}$. This means the critical discount factor to sustain (L, L) in every period without transfers lies in the interval $[\frac{1}{3}, \frac{5}{14}]$.

several action profiles of the action plan have the highest liquidity requirement, we replace all those action profiles. In our example, we replace the equilibrium action profile a^e, so that the new action plan becomes $a^e = (L, M), a^1 = (M, H), a^2 = (H, M)$. Using again formula (8), we find that this action plan can be implemented whenever

$$r \leq \frac{18}{\max\{4, 6\}} = 300\%.$$

Correspondingly, for every discount rate $\delta \in (\frac{1}{4}, \frac{1}{3})$ the actual action plan is optimal and the set of subgame perfect equilibrium payoffs is given by all (u_1^0, u_2^0) with $u_1^0 + u_2^0 \leq 18$ and $u_1^0, u_2^0 \geq 0$.

We repeat step 4 until we reach the end of the list of action profiles in every state k. The final action plan only consists of Nash equilibria of the stage game. In the example, we find the following critical discount factors, payoffs and action plans:

Step	δ^*	U^e	v_1	v_2	a^e, a^1, a^2
1	1/3	20	0	0	(L,L),(M,H),(H,M)
2	1/4	18	0	0	(L,M),(M,H),(H,M)
*3	1/2	18	7	7	(L,M),(M,M),(M,M)
4	0	14	7	7	(M,M),(M,M),(M,M)

Note that the critical discount factor δ^* does not necessarily decrease in every step. If δ^* it is not lower than in all previous steps, we simply ignore the corresponding action plan. This is the case in step 3 of our example.

The algorithm always delivers a list of all critical discount factors, corresponding payoff sets and optimal action plans. When using a heap sort algorithm to create the $n + 1$ ordered lists, which each have a maximal length of $|A|$ action profiles, the computational complexity of our algorithm in terms of elementary calculations and comparisons is of just log-linear order $\mathcal{O}(n|A| \log |A|)$. Even large stage games with more than a 100000 action profiles can be solved in less than a second.

Kranz (2010) explains how to use the software implementation of our algorithm and gives several examples. It is also illustrated how methods of adaptive grid refinement and random sampling of action profiles allow to effectively compute inner approximations to the sets of SPE payoffs for continuous stage games with high dimensional action spaces (like oligopolies with 10 or more firms).[14]

[14]Consider stage games with compact action spaces and continous payoff functions. If one can provide closed-form solutions of the cheating payoffs of the continuous stage game, one can

In comparison, we can note that allowing for monetary transfers allows much faster computation of the set of pure strategy equilibrium payoffs than in the framework studied by Judd, Yeltekin and Conklin (2003) with public randomization. That is because without monetary transfers no general closed-form solutions for the static problems could be obtained, and in each iteration the algorithm of Judd et. al. has to solve several linear programs.[15]

5 A Noisy Prisoners' Dilemma game

In this example we derive closed form solutions for the set of pure strategy PPE payoffs in a repeated noisy prisoners' dilemma game with imperfect public monitoring. There are two players. In the stage game, a player can either cooperate (C) or defect (D). Expected payoffs $g(a)$ are given by the following normalized payoff matrix:

	C	D
C	$1, 1$	$-s$,$1+d$
D	$1+d$,$-s$	0,0

with $d, s > 0$ and $d - s < 1$. Players do not publicly observe the played action profile, but only a realized signal y that can take four different values: y_C, y_D, y_1 and y_2. The signal distribution is as follows:

calculate the liquidity requirement $\mathcal{L}(a)$ for any action profile $a \in A$.

To compute inner approximations of the sets of SPE payoffs, we can draw a finite random sample of action profiles in order to calculate lower bounds of the function $\bar{U}^e(L)$ and upper bounds on $\bar{v}^i(L)$ in a similar way we calculated the step functions above. As the sample size grows large, these lower bounds converge in probability to the true functions.

The practical issue is to sample action profiles in a way that achieves relatively quick convergence for most stage games. Different methods are implemented in the software package and work well in examples.

[15]JYC report a computation time of almost 45 minutes (on a Pentium 500Mhz, PC) for the finest considered approximation for the payoff set of a discretized repeated Cournot duopoly with 15 x 15 action profiles and a given discount factor of $\delta = 0.8$

$\phi(\mathbf{y}\vert\mathbf{a})$	**CC**	**CD**	**DC**	**DD**
$\mathbf{y}_C$	$1-\alpha_A-2\alpha_P$	$1-\alpha_A-\beta_A-2\alpha_P-\beta_P$	$1-\alpha_A-\beta_A-2\alpha_P-\beta_A$	0
$\mathbf{y}_D$	α_A	$\alpha_A+\beta_A$	$\alpha_A+\beta_A$	1
$\mathbf{y}_1$	α_P	α_P	$\alpha_P+\beta_P$	0
$\mathbf{y}_2$	α_P	$\alpha_P+\beta_P$	α_P	0

with $0<\alpha_A\le\alpha_A+\beta_A$ and $0<\alpha_P\le\alpha_P+\beta_P$ and $1-\alpha_A-\beta_A-2\alpha_P-\beta_P\ge 0$. To interpret the signal structure, assume that mutual cooperation CC shall be implemented.[16] The signal y_D is an anonymous indicator for defection: y_D becomes more likely if some player unilaterally defects but its probability distribution does not depend on the identity of the deviator. The parameter α_A can be interpreted as the probability of a type-one error, i.e. the probability that y_D is observed although no player defected. The parameter β_A measures by how much the likelihood of y_D increases if some player unilaterally deviates.

The signal y_i is an indicator for unilateral defection by player i. Like α_A, the parameter α_P can be interpreted as the probability of a type-one error, i.e. the probability to wrongly get a signal for unilateral defection of player i. Similar to β_A, the parameter β_P measures by how much the likelihood of y_i increases if player i unilaterally deviates from mutual cooperation.

To calculate the required liquidity to implement mutual cooperation in the static problem, consider an equal liquidity distribution $\lambda_1=\lambda_2=\frac{1}{2}$. Clearly, incentives to deviate for each player i are minimized if he is required to make the maximum payments $\frac{1}{2}L$ after signals y_D and y_i. Since the problem is symmetric, it is disadvantageous to impose on some player a payment after signal y_C. Whether player i has to make a payment or receives a payment after signal y_{-i} has no effect on his incentives to deviate in the static problem. Mutual cooperation can thus be implemented with total liquidity L if and only if

$$0\ge d-(\beta_A+\beta_P)\frac{1}{2}L,$$

which yields a liquidity requirement of

$$\mathcal{L}(CC)=\frac{2d}{\beta_A+\beta_P}.$$

[16]For notational convenience, we abbreviate action profiles (a_1,a_2) by a_1a_2.

This formula is quite intuitive. If actions could be perfectly monitored, the liquidity requirement would be $2d$. This value is divided by the increase in the likelihood to get a signal y_i or y_D if player i defects.

To minimize the amount of money burning, it is optimal that after signal y_1 player 1 transfers all of his liquidity to player 2, and vice versa. Money burning can only be optimal after signal y_D. We find that for $L \geq \frac{2d}{\beta_P}$, mutual cooperation can be implemented without any money burning and that for $L \in [\mathcal{L}(CC), \frac{2d}{\beta_P})$, a total amount of $\frac{2d-\beta_P L}{\beta_A}$ must be burned after signal y_D. The maximum implementable joint payoffs are thus given by

$$U^e(L, CC) = \begin{cases} 2 & \text{if } L \geq \frac{2d}{\beta_P} \\ 2(1 - \frac{\alpha_A}{\beta_A} d) + \frac{\alpha_A}{\beta_A} \beta_P L & \text{if } \frac{2d}{\beta_A + \beta_P} \leq L \leq \frac{2d}{\beta_P} \end{cases} . \tag{9}$$

Let us now consider the asymmetric action profile CD. Its liquidity requirement can be most easily calculated by assuming that the whole liquidity is allocated to player 1. The minimum required payment $p_1(y_D)$ after signal y_D that removes player 1's incentives to defect satisfies

$$s + (\alpha_A + \beta_A) p_1(y_D) = p_1(y_D).$$

If after signal y_D player 1 makes this payment $p_1(y_D)$ to player 2 and no other payments are made, then no player has an incentive to deviate and no money is burned. We thus find

$$\mathcal{L}(CD) = \frac{s}{1 - \alpha_A - \beta_A}$$

and

$$U^e(L, CD) = G(CD) = 1 + d - s.$$

For the action profile DC the same results hold and for the stage game Nash equilibrium it holds that $\mathcal{L}(DD) = 0$ and $U^e(L, DD) = 0$.

For every level of total liquidity L, the profile DD is an optimal punishment profile for both players, since the Nash equilibrium payoffs are min-max payoffs for both players. Hence, we find $\overline{v}_i(L) = 0$ for all $L \geq 0$.

Recall that in games with perfect monitoring, $\overline{U}^e(L) - \overline{V}(L)$ is always a step function. The algorithm for perfect monitoring calculates the critical discount rate $r^*(L)$ at every jump point. With imperfect monitoring, $\overline{U}^e(L) - \overline{V}(L)$ is in general an increasing piece-wise linear function with jumps. Figure 1 illustrates the function $\overline{U}^e(L) - \overline{V}(L)$ for the noisy prisoners' dilemma game for a parameter constellation that satisfies $\beta_P > 0$ and $0 < G(CD) < U^e(L(CC), CC)$.

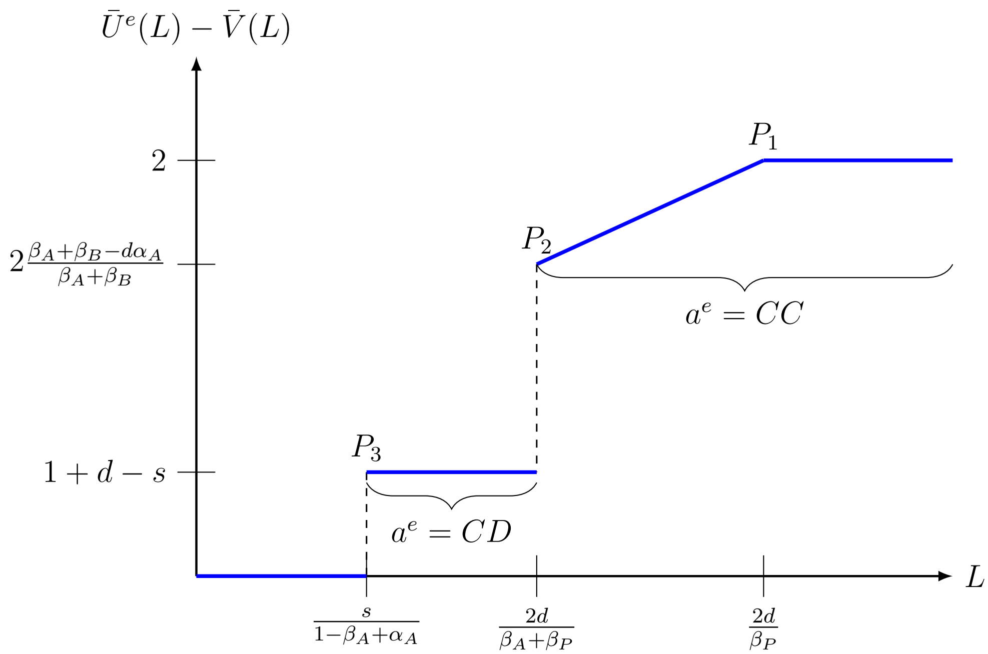

Figure 1: Optimal action profiles and payoffs of the noisy prisoners' dilemma game

The graph has a kink P_1 and two jump points P_2 and P_3. We can calculate the critical discount rate at every jump point, kink and increasing linear segment of $\overline{U}^e(L) - \overline{V}(L)$ by using the formula

$$r^*(L) = \frac{\overline{U}^e(L) - \overline{V}^e(L)}{L}. \tag{10}$$

For the points P_1 and P_2, we find

$$r^*\left(\frac{2d}{\beta_P}\right) = \frac{\beta_P}{d}$$

and

$$r^*\left(\frac{2d}{\beta_A+\beta_P}\right) = \frac{\beta_P+\beta_A-d\alpha_A}{d}.$$

On the line segment between the two points, i.e. for $L \in (\frac{2d}{\beta_A+\beta_P}; \frac{2d}{\beta_P}]$, the maximum discount rate is given by

$$r^*(L) = \frac{2(1-\frac{\alpha_A}{\beta_A}d)}{L} + \frac{\alpha_A}{\beta_A}\beta_P.$$

Money burning facilitates the implementation of CC if the maximum discount rate increases when moving from P_1 to P_2. This is the case if and only if $d \leq \frac{\beta_A}{\alpha_A}$.

Given a plot of $\overline{U}^e(L) - \overline{V}(L)$, as in Figure 1, there is a simple graphical rule to find out whether the maximal discount rate increases or decreases along a line segment. Consider the intercept at $L = 0$ of the line going through P_1 and P_2. The critical discount rate increases from P_1 to P_2 if and only if this intercept is positive. With a sharp glance, one can establish that this is indeed the case in Figure 1.

Similarly, one can check graphically whether the maximal discount rate is higher in point P_3 than in point P_2. In Figure 1, the intercept of the line through P_2 and P_3 is negative. This means that in the depicted case there is no discount rate for which CD or DC are optimal equilibrium state profiles: playing CC with appropriate amounts of money burning yields higher payoffs and can be implemented for a larger range of discount factors.

By solving equation (10) for L and plugging into the formula for $\overline{U}^e(L)$, one can find the maximal joint equilibrium payoff $U^e(r)$ as a function of the discount rate r. For the case depicted in Figure 1, we find:

$$U^e(r) = \begin{cases} 2 & \text{if } r \leq \frac{d}{\beta_P} \\ 2(1 - \frac{\alpha_A}{\beta_A} d) \left[1 + \frac{\alpha_A}{r\beta_P\beta_A - \alpha_A}\right] & \text{if } \frac{d}{\beta_P} \leq r \leq \frac{\beta_P + \beta_A - d\alpha_A}{d} \\ 0 & \text{otherwise} \end{cases} . \tag{11}$$

Together with the fact that one can always implement punishment payoffs of zero, condition (11) characterizes the set of pure strategy sequential equilibrium payoffs for the considered case. Alternative cases, e.g. parameter constellations in which CD is an optimal equilibrium state profile for some discount rates, can be characterized in a similar fashion.

6 Repeated games without money burning

In this section we explore what can be achieved in a repeated game with transferable utility if money burning is not allowed. In particular, we investigate to which extent money burning can be replaced by the use of a public correlation device. We consider a variant of the previous set-up in which payments are required to add up to zero, and in which players observe the outcome of a public correlation device at the beginning of each period.

To characterize the set of PPE payoffs in this class of games, we extend action and payment plans by a collective punishment state, indexed with $k = b$. The public correlation device allows strategies to specify transition probabilities to this collective punishment state. The proof of Theorem 2 below shows that all PPE payoffs can be implemented by a class of stationary equilibria that move with positive probability to the collective punishment state instead of the equilibrium state if all payments are conducted.

We develop a more convenient characterization of equilibrium payoffs by considering stationary equilibria that have an endogenous restriction on the amount of money burning. Consider a stationary strategy profile σ of the game with money burning and add a collective punishment state $k = b$ with action profile α^b and payment function p^b. We define an additional constraint on the amount of money burning

$$\sum_{i=0}^{n} p_i^k(y) \leq \frac{\delta}{(1-\delta)} \left(U(\sigma) - U^b\right) \quad \text{for all } y \tag{MBC-k}$$

with $U^b = \sum_{i=1}^{n} u_i^b$ and

$$u_i^b = (1-\delta)(g_i(\alpha^b) - E[p_i^b|\alpha^b]) + \delta u_i(\sigma).$$

Then we consider the following maximization problem over action plans and payment plans that are extended in this way

$$\max_{(\alpha^k, p^k)_{k=e,b,1,...,n}} \quad U - V - U^b \tag{LP-MB}$$

s.t. (PC-k),(AC-k),(BC-k) and (MBC-k) for all $k = e, b, 1, ..., n$.

Theorem 2 *If the linear program LP-MB is solved by a stationary equilibrium σ of the game with money burning and collective punishment state action profile α^b and payment function p^b, it holds that the set of PPE payoffs in the game without money burning is given by*

$$\left\{ u^0 \in \mathbb{R}^n | U^b \leq \sum_{i=1}^{n} u_i^0 \leq U(\sigma), u_i^0 \geq v_i(\sigma) \right\}. \tag{12}$$

For an algorithm that allows a faster computation of the payoff set, we can derive similar links to static problems as in games with unlimited money burning. Consider the static problem of Section 3.2, with the additional restriction that there is an upper bound $B \geq 0$ on the amount of money that is allowed to be burned after any signal y. We denote by $L(\alpha, B)$ the liquidity requirement of an action profile with that upper bound on money burning:

$$L(\alpha, B) = \min_{p(.)} L \text{ s.t. } (PC), (AC), (BC) \text{ and} \tag{LP-BL}$$

$$\sum_{i=0}^{n} p_i(y) \leq B \text{ for all } y \in Y. \tag{MBC}$$

Similarly, we define for all $L \geq L(\alpha, B)$ and $0 \leq B \leq L$ the highest joint equilibrium payoff in the static problem by

$$U^e(L, B, \alpha) = \max_{p(.)} \left(G(\alpha) - \sum_{i=1}^{n} E[p_i|\alpha] \right) \tag{LP-Be}$$
$$\text{s.t.(PC),(BC), (AC), (MBC),}$$

the lowest collective punishment payoff by

$$U^b(L, B, \alpha) = \min_{p(.)} \left(G(\alpha) - \sum_{i=1}^{n} E[p_i|\alpha] \right) \tag{LP-Bb}$$
$$\text{s.t.(PC),(BC), (AC), (MBC),}$$

and player i's punishment payoff by

$$v^i(L, B, \alpha) = \min_{p(.)} \left(g_i(\alpha) + \lambda_i L - \sum_{i=1}^{n} E[p_i|\alpha] \right) \tag{LP-Bi}$$
$$\text{s.t.(PC),(BC), (AC) and (MBC).}$$

The corresponding upper and lower envelopes over all action profiles are denoted by

$$\begin{aligned}
\bar{U}^e(L, B) &= \max_{\alpha \in \mathcal{A}} U^e(L, B, \alpha), \\
\bar{U}^b(L, B) &= \min_{\alpha \in \mathcal{A}} U^b(L, B, \alpha), \\
\overline{v}_i(L, B) &= \min_{\alpha \in \mathcal{A}} v^i(L, B, \alpha).
\end{aligned}$$

The profiles at which these values are attained are denoted by $\overline{\alpha}^k(L, B)$. We say a pair (L, B) of liquidity and bound on money burning can be generated by stationary equilibria for a discount factor δ if

$$\begin{aligned}
L &\leq \frac{\delta}{1-\delta} \left(\bar{U}^e(L, B) - \bar{V}(L, B) \right), \\
B &\leq \frac{\delta}{1-\delta} \left(\bar{U}^e(L, B) - \bar{U}^b(L, B) \right).
\end{aligned}$$

Let $(\bar{L}(\delta), \bar{B}(\delta))$ denote the pair of largest liquidity and bound on money burning that can be generated for a discount factor δ. Whenever some pair (L, B) can be generated, the pair $(\bar{L}(\delta), \bar{B}(\delta))$ also can be generated, since larger levels of B allow to generate larger levels of L and vice versa.

Proposition 3 *In the case without money burning, the set of PPE payoffs under discount factor δ is given by*

$$\{u^0 \in \mathbb{R}^n \mid \bar{U}^b(\bar{L}(\delta), \bar{B}(\delta)) \leq \sum_{i=1}^{n} u_i^0 \leq \bar{U}^e(\bar{L}(\delta), \bar{B}(\delta)) \tag{13}$$
$$\text{and } u_i^0 \geq v^i(\bar{L}(\delta), \bar{B}(\delta)) \text{ for all } i\}.$$

To compute the functions $U^e(L, B, \alpha)$, $U^b(L, B, \alpha)$ and $v^i(L, B, \alpha)$ for all $L \geq L(\alpha, B)$ and $B \leq L$ one can exploit the fact that their surface is described by a finite number of planar segments, which can be characterized by methods of parametric linear programming and sensitivity analysis (see, e.g., Gal and Nedoma [9]). The computations can take considerably longer than computing the one-dimensional functions for the case of unlimited money burning. Still, one may be able to obtain closed-form solutions for simple signal structures.[17] Once $\bar{U}^e(L, B)$ and $\bar{v}^i(L, B)$ are fully characterized, optimal action structures for all discount factors can be very quickly obtained.

A sufficient condition for the equilibrium payoff set not to be affected by the possibility to burn money is that a single stage game Nash equilibrium α^b is an optimal punishment profile for all players. Both the collective punishment payoff U^b and the sum of individual punishment payoffs V are then equal to $G(\alpha^b)$ and the payment constraints imply the money burning constraints. Hence, our characterization of the payoff sets in the noisy prisoners' dilemma game remains valid even if no money burning is allowed. In addition, we have already found that the restriction not to burn money does not shift the Pareto frontier of the set of equilibrium payoffs in games with perfect monitoring.

7 Conclusion

In this paper, we presented a characterization of PPE payoff sets for infinitely repeated games with public monitoring and monetary transfers. Monetary transfers

[17] For example, in the noisy prisoner's dilemma game and the action profile $a = CC$, we find $L(a, B) = \frac{1}{\beta_P}(2d - B\beta_A)$ and $U^e(L, B, a)$ is given as in equation (9).

are a realistic assumption and at the same time greatly simplify the analysis. Our results can be used to numerically compute the pure strategy equilibrium payoff sets for any finite stage game and they also facilitate development of closed-form analytical solutions.

One interesting direction for future work is to study to which extent monetary transfers, in conjunction with communication, allow a tractable characterization of payoff sets for games with private monitoring or for the set of mixed strategy equilibrium payoffs in games with public monitoring. The problem becomes considerably more complicated, since it is not necessarily optimal to use a payment plan that induces full information revelation in every period (see, e.g. Fuchs [7] for an analysis in a principal agent framework).

Another direction for future research is to study optimal renegotiation-proof equilibria in a framework with monetary transfers and imperfect public monitoring. If we considered only stationary equilibria, a natural, minimal renegotiation-proofness requirement is that after no history there shall be money burning. An interesting question is whether there is a concept of renegotiation-proofness for which every renegotiation-proof payoff can be implemented with a stationary equilibrium without money burning.

Appendix A: Perfect monitoring, Mixed Strategies and Money burning

This appendix exemplifies that under perfect monitoring, optimal stationary equilibria with mixed strategies may require inefficient continuation play before payment stages, which can be achieved with money burning. Consider the following stage game

	C	D	E	F
C	$1,1$	$-d,1+d$	$-x,x$	$-x,-x$
D	$1+d,-d$	$0,0$	$-x,-x$	$-x,x$
E	$x,-x$	$-x,-x$	$0,0$	$0,0$
F	$-x,-x$	$x,-x$	$0,0$	$0,0$

Positive joint payoffs can only be achieved if at least some player chooses C with positive probability and the other player chooses C or D with positive probability. We denote with $\alpha_i(a_i)$ the probability with which player i plays the pure action

a_i. A necessary condition that a player $i \neq j$ who is asked to play C or D in the static problem will not deviate to E is

$$1 + d \geq x(\alpha_j(C) - \alpha_j(D)) - L$$

which is equivalent to

$$L \geq x(\alpha_j(C) - \alpha_j(D)) - 1 - d \tag{14}$$

Similarly, a necessary condition that player i will not deviate to F is given by

$$L \geq x(\alpha_j(D) - \alpha_j(C)) - 1 - d \tag{15}$$

For ease of exposition, we will consider in the following analysis the limit case $x \to \infty$. The main result that money burning on the equilibrium may be necessary in optimal stationary equilibria will hold via continuity for any sufficiently large finite x.

In that limit case, conditions (14) and (15) can be jointly satisfied only if player j mixes with equal and positive probability between C and D. This means any equilibrium in which C or D is played with positive probability by one player requires

$$\alpha_1(C) = \alpha_1(D) \text{ and } \alpha_2(C) = \alpha_2(D)$$

Critical discount factor without money burning We now derive a lower bound on the minimal discount factor for every stationary equilibrium with positive expected payoffs in which no money is burned. Assume liquidity is equally distributed across players and every player is required to transfer all his liquidity to the other player in case he plays D and the other player plays C. Player i is then indifferent between playing C and D if and only if

$$\alpha_j(C)+\alpha_j(D)(\frac{L}{2}-d)-x(\alpha_j(E)+\alpha_j(F)) = \alpha_j(C)(1+d-\frac{L}{2})-x(\alpha_j(E)+\alpha_j(F))$$

Solving this condition for L with the constraint $\alpha_j(C) = \alpha_j(D)$, yields a minimal liquidity requirement of

$$\mathcal{L} = 2d$$

The endogenously generated liquidity (using optimal penal codes) is given by

$$L = \frac{\delta}{1-\delta}G(\alpha)$$

The joint payoff $G(\alpha)$ is maximized if no player chooses E or F with positive probability, i.e. every player chooses C and D with probability $\frac{1}{2}$. We then have $G(\alpha) = 1$. Equalizing $\mathcal{L}$ and L, yields a minimal discount factor of

$$\bar{\delta}_0 = \frac{2d}{1 + 2d}$$

A discount factor of at least $\bar{\delta}_0$ is necessary for there to be a stationary equilibrium without money burning that has a positive joint payoff.

Allowing for money burning We now show that by using money burning on the equilibrium path, positive equilibrium payoffs can also be implemented for discount factors below $\bar{\delta}_0$.

We consider equilibria with the structure that each player plays C and D with probability $\frac{1}{2}$, transfers $\frac{L}{2}$ units of money to the other player if he chooses D and the other player C, and burns $b \leq \frac{L}{2}$ units of money under the outcome DD. Player i is then indifferent between playing C or D if and only if

$$\frac{1}{2}(1 + \frac{L}{2} - d) = \frac{1}{2}(1 + d - b)$$

which corresponds to a liquidity requirement of

$$\mathcal{L}(b) = 2d - 2b$$

The endogenously generated liquidity is given by

$$L(b) = \frac{\delta}{1 - \delta}\left[1 - \frac{b}{2}\right]$$

Money burning under DD reduces the liquidity requirement, since players have smaller incentives to deviate from C to D. On the other hand, money burning happens with positive probability on the equilibrium path and therefore also reduces the endogenously generated liquidity. Setting $\mathcal{L}(b)$ equal to $L(b)$ yields a minimal discount factor of

$$\bar{\delta}(b) = \frac{4d - 4b}{2 + 4d - 5b}$$

For $b = 0$, the discount factor $\bar{\delta}(b)$ coincides with the minimal discount factor $\bar{\delta}_0$ required to have positive payoffs without money burning. Its derivative with respect to b is given by

$$\bar{\delta}'(b) = -\frac{4(2 - d)}{(4d - 5b + 2)^2}$$

Thus, as long as $d < 2$, small amounts of money burning under outcome DD allow to implement positive expected payoffs for some discount factors below the minimal discount factor for stationary equilibria without money burning.

A sufficient condition for not requiring money burning In the presented example, both players mix in the optimal stationary equilibrium. This turns out to be a necessary condition for money burning to be part of an optimal stationary equilibrium. It can be shown that money burning is never necessary as soon as one player plays a pure strategy, even if the perfect monitoring assumption only holds for that player.[18] The argument goes a follows: Let us assume that the actions of player 1 are perfectly observable to the other players, such that the signal space has the form $Y = A_1 \times Y_{-1}$ with

$$\phi((a_1, y_{-1})|a) = \phi_{-1}(y_{-1}|a_{-1})$$

and

$$\phi((\hat{a}_1, y_{-1})|a) = 0 \text{ if } \hat{a}_1 \neq a_1.$$

Now consider an action profile α in which $\alpha_1 = a_1$ is a pure action. We claim that for any liquidity with which α can be implemented, it can also be implemented without money burning. In order to see this, take a payment plan p that implements α for liquidity L, and then define the payments $\tilde{p}$ as

$$\tilde{p}_1(a_1, y_{-1}) = p_1(a_1, y_{-1}) - B(a_1, y_{-1}),$$

where $B(a_1, y_{-1})$ denotes the total amount of money burned after signal (a_1, y_{-1}). That is, we simply give player 1 all the money that was supposed to be burned. Furthermore, define

$$\tilde{p}_2(\hat{a}_1, y_{-1}) = p_2(\hat{a}_1, y_{-1}) - B(\hat{a}_1, y_{-1}) \text{ if } \hat{a}_1 \neq a_1.$$

All other payments stay the same. Then the payment constraints still hold, because payments have become weakly lower, and the budget constraint is satisfied with equality. If an action constraint for player 1 changes at all, then it is relaxed, and the other players' action constraints are not affected. Hence, it is always weakly better to use payments that are budget-balanced after every signal.

Appendix B: Computing $U^e(L, a)$ and $\overline{U}^e(L)$

This appendix illustrates for the case of pure strategy equilibria, how $U^e(L, a)$ and $\overline{U}^e(L)$ can be exactly computed and describes heuristics to reduce computation time. Similar methods can be applied to the computation of $v_i(L, a)$ and $\overline{v}_i(L)$.

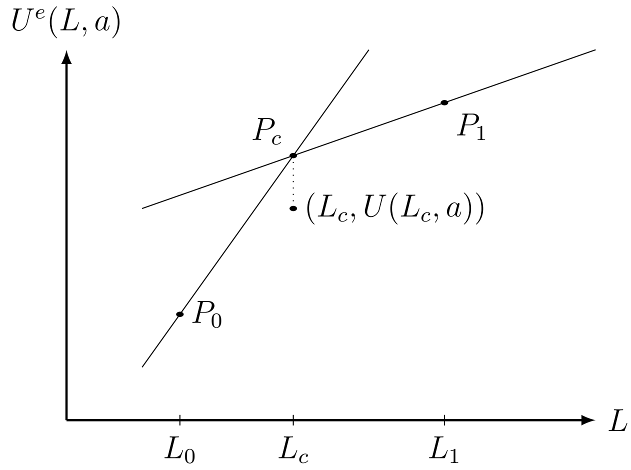

Figure 2: Constructing $U^e(L,a)$

Calculating $\mathbf{U}^e(\mathbf{L},\mathbf{a})$ Assume that we have calculated $U^e(L,a)$ at two different levels $L_0 < L_1$ illustrated by the points P_0 and P_1 in Figure 2. We describe a procedure that fully computes $U^e(L,a)$ on the interval $[L_0, L_1]$. From the dual values of the solution of the problem (LP-e) we can get the slope of $U^e(L,a)$ at L_0 and L_1.[19] Figure 2 illustrates the corresponding tangents. The two tangents either coincide or have a cut point $P_c = (L_c, U_c)$ with $L_0 < L_c < L_1$ and $U_0 < U_c < U_2$. In the first case, $U^e(L,a)$ is given on the interval $[L_0, L_1]$ by the line $\overline{P_0P_1}$. In the second case the line $\overline{P_0P_cP_1}$ constitutes an upper bound on $U^e(L,a)$. We calculate $U^e(L_c,a)$. If $U^e(L_c,a) = U_c$ then $U^e(L,a)$ coincides with this upper bound $\overline{P_0P_cP_1}$. Otherwise, we proceed recursively by calculating $U^e(L,a)$ on the two intervals $[L_0, L_c]$ and $[L_c, L_1]$. If there are $n_k \geq 2$ kinks between L_0 and L_1, this procedure fully characterizes the function $U^e(L,a)$ on the interval by solving at most $2\,(n_k - 1) + 1$ times the linear program (LB-e). To quickly solve (LP-e) at different levels of L, one can use standard re-optimization techniques, e.g. based on the dual simplex algorithm.[20]

[18] We thank Jon Levin for this suggestion.

[19] If $U^e(L|a)$ has a kink at L, it depends on the way the linear program is set up, whether the dual values delivers the right hand or left hand slope. It is no problem to calculate, the correct slope, however.

[20] Moreover, using a simplex algorithm, the case $U^e(L_c,a) = U_c$ can sometimes be verified without the need of solving the linear program (LP-e) at L_c. A sufficient condition for

The lowest possible level of L is given by the liquidity requirement $L(a)$. The right hand starting point of our procedure is given by the minimal liquidity $\overline{L}^e(a)$ above which $U^e(L,a)$ does not anymore increase in L. We can calculate $\overline{L}^e(a)$ by adding a restriction on the maximal allowed expected amount of money burning in the problem (LP-L).[21]

Calculating the upper envelope $\overline{U}^e(L)$ For the calculation of the upper envelope $\overline{U}^e(L)$, let us define by

$$U^e(L,\widetilde{A}) = \max_{a\in\widetilde{A}} U^e(L,a)$$

the upper envelope with respect to a subset of action profiles $\widetilde{A} \subseteq A$. Hence, we have

$$U^e(L,\widetilde{A}\cup\{a\}) = \max\{U^e(L,\widetilde{A}), U^e(L,a)\}.$$

We can calculate $\overline{U}^e(L)$ by subsequently adding all action profiles to the set $\widetilde{A}$. To calculate the new envelope $U^e(L,\widetilde{A}\cup\{a\})$, it is often not necessary to compute the whole function $U^e(L,a)$. Recall, that the method to calculate $U^e(L,a)$ delivers in each step an upper bound on $U^e(L,a)$. It suffices to proceed the calculation of $U^e(L,a)$ only for those values of L for which the upper bound exceeds $U^e(L,\widetilde{A})$. If an upper bound of $U^e(L,a)$ lies everywhere below $U^e(L,\widetilde{A})$, we can immediately dismiss the action profile a. Since $U^e(L,a)$ is bounded by $G(a)$, a sufficient condition to dismiss a is that $G(a) \leq U^e(L(a),\widetilde{A})$. A weaker sufficient condition is $G(a) \leq U^e(\widetilde{L}(a),\widetilde{A})$, where $\widetilde{L}(a) \equiv \sum_{i=1}^{n}(c_i(a) - g_i(a))$ is the liquidity requirement under perfect monitoring, which always satisfies $\widetilde{L}(a) \leq L(a)$. The last condition can be checked very quickly since no linear program has to be solved for a.

The order in which action profiles are added to $\widetilde{A}$ can influence the total computation time, because action profiles can be more quickly dismissed if $U^e(L,\widetilde{A})$ is already large. One should first add all Nash equilibria of the stage game, which

---

$U^e(L_c,a) = U_c$ is that the the optimal (dual) basis of the solved problem at L_0 (or L_1) remains an optimal basis at L_c. This condition can be checked with standard formulas used to calculate sensitivity bounds. However, it can happen that the optimal basis changes between L_0 and L_c even though the function $U^e(L|a)$ has no kink between L_0 and L_c.

[21] If the full-dimensionality condition of the folk theorem by Fudenberg, Levine and Maskin [8] holds we must impose zero money burning to calculate L. Otherwise, we first have to solve the problem (LB-e) with unlimited liquidity to calculate the minimally required amount of money burning.

satisfy $U^e(L,a) = G(a)$ for all $L \geq 0$. An educated guess about which action profiles are likely to be optimal, e.g. symmetric ones, can be furthermore helpful.
Punishment states Similar methods can be used to calculate $v_i(L,a)$ and $\overline{v}_i(L)$. For the computation of $\overline{v}_i(L)$, it is helpful to first add to $\widetilde{A}$ all those action profiles a in which a_i is a best-reply to a_{-i}, since these action profiles satisfy $v_i(L,a) = g_i(a)$ for all $L \geq L(a)$.

Appendix C: Proofs

Proposition 1 and Lemma 1 are straightforward and proofs are omitted.
Proof of Theorem 1: We rely on the recursive structure of public perfect equilibria and compactness of the equilibrium value set (see e.g. the result in APS, which straightforwardly extend to our setting, and the corresponding results for PPE with mixed strategies in the book by Mailath and Samuelson [19]). Let $\bar{U}$ denote the highest joint payoff that can be implemented with some PPE and $\bar{v}_i$ the lowest payoff for player i that can be implemented with some PPE. There must exist a PPE σ^e without payments in the first payment stage whose joint payoffs are given by

$$\sum_{i=1}^{n} u_i^0(\sigma^e) = \bar{U}.$$

Furthermore, for every player $i = 1, ..., n$, there exists a PPE σ^i without payments in the first payment stage that gives player i a payoff of

$$u_i^0(\sigma^i) = \bar{v}_i.$$

For all $k \in K$ let α^k be the first action profile played on the equilibrium path of σ^k. Let $w^k(y)$ be the vector of continuation payoffs of σ^k in the first period after signal y has been realized (but before the second payment stage), i.e. we have

$$u_i^0(\sigma^k) = (1-\delta)g_i(\alpha^k) + E[w_i^k|\alpha^k].$$

We define

$$p_i^k(y) = \frac{\delta u_i^0(\sigma^e) - w_i^k(y)}{1-\delta},$$

and will show that the stationary strategy profile σ defined by action plan $(\alpha^k)_k$ and payment plan $(p^k)_k$ is a stationary equilibrium. The budget constraints of σ are equivalent to

$$\delta\bar{U} \geq \sum_{i=1}^{n} w_i^k(y),$$

which holds due to the definition of $\bar{U}$ as the highest possible sum of payoffs and the fact that the sum of payments cannot be negative. Second, for the action constraints, we have to show that

$$g_i(a_i, \alpha_{-i}^k) - E[p_i^k|a_i, \alpha_{-i}^k] \geq g_i(\widehat{a}_i, \alpha_{-i}^k) - E[p_i^k|\widehat{a}_i, \alpha_{-i}^k],$$

holds for all $i \in \{1, ..., n\}$, $a_i \in supp(\alpha_i)$ and $\widehat{a}_i \in A_i$. This condition is equivalent to

$$(1-\delta)g_i(a_i, \alpha_{-i}^k) + E[w_i^k|a_i, \alpha_{-i}^k] \geq (1-\delta)g_i(\widehat{a}_i, \alpha_{-i}^k) + E[w_i^k|\widehat{a}_i, \alpha_{-i}^k],$$

which describes the incentive constraints for playing α^k in the first period of σ^k. Third, for the payment constraints we have to show that

$$p_i^k(y) \leq \delta(g_i(\alpha^e) - E[p_i^e|\alpha^e] - g_i(\alpha^i) + E[p_i^i|\alpha^i]).$$

With our definition of payments $p_i^k(y)$ this reads

$$\delta u_i^0(\sigma^e) - w_i^k(y) \leq \delta((1-\delta)g_i(\alpha^e) + E[w_i^e|\alpha^e] - (1-\delta)g_i(\alpha^i) - E[w_i^i|\alpha^i]),$$

which is equivalent to

$$w_i^k(y) \geq \delta \bar{v}_i.$$

Because $\bar{v}_i$ is the lowest player i payoff in the action stage, this condition obviously holds if player i receives a net payment after signal y in the corresponding continuation equilibrium of σ^i. It also holds for signals which require player i to make a net transfer, because otherwise player i would have an incentive not to make the payment and σ^i would not be a PPE. Player i's expected payoff in the stationary equilibrium σ is

$$g_i(\alpha^e) - \frac{1}{1-\delta}E[\delta\bar{u}_i^e - w_i^k(y)|\alpha^e] = u_i(\sigma^e),$$

and his punishment payoff is

$$(1-\delta)g_i(\alpha^i) - E[\delta\bar{u}_i^e - w_i^k(y)|\alpha^i] + \delta\bar{u}_i^e = \bar{v}_i.$$

It then follows from Proposition 1 that we can define incentive compatible up-front payments for σ to implement any PPE equilibrium payoff.■

Proof of Lemma 2: Since $p_i(y) \leq \lambda_i L$, the action constraints (AC) for player i imply

$$g_i(a_i, \alpha_{-i}) + \lambda_i L - E[p_i|a_i, \alpha_{-i}] \geq g_i(\widehat{a}_i, \alpha_{-i}) \; \forall a_i \in supp(\alpha_i), \widehat{a}_i \in A_i,$$

which implies $v^i(L,\alpha) \geq c_i(\alpha)$. In the case that $g_i(\alpha) = c_i(\alpha)$, it follows that $g_i(a_i, \alpha_{-i}) = c_i(\alpha_i)$ for all $a_i \in supp(\alpha_i)$. One can then take $\lambda_i = 0$ and $p_i(y) = 0$ for all y to implement α.■

Proof of Proposition 2: Let λ be an arbitrary liquidity distribution and for every state $k = e, 1, ..., n$, let $\tilde{p}^k$ be a payment function that solves the static problems (LP-k) given liquidity allocation $\lambda \bar{L}(\delta)$ and action profile $\overline{\alpha}^k(\bar{L}(\delta))$. We now construct a liquidity distribution λ^* and a stationary equilibrium σ with action plan $(\overline{\alpha}^k(\bar{L}(\delta)))_k$ such that the payment constraints of the stationary contract will coincide with the payment constraints of the static problems (LP-k) given a liquidity allocation $\lambda^* \bar{L}(\delta)$.

We define λ^* by

$$\lambda_i^* = \delta \left(\lambda_i + \frac{g_i(\overline{\alpha}^e) - E[\tilde{p}_i^e|\overline{\alpha}^e] - \bar{v}^i(\bar{L}(\delta))}{\bar{L}(\delta)} \right).$$

λ^* is a liquidity distribution, since i) by the definition of $\tilde{p}_i^e$ and $\bar{v}^i$, we have $g_i(\overline{\alpha}^e) - E[\tilde{p}_i^e|\overline{\alpha}^e] \geq \bar{v}^i(\bar{L}(\delta))$, i.e. $\lambda_i^* \geq 0$, and ii)

$$\sum_{i=1}^{n} \lambda_i^* = \delta \left(1 + \frac{\bar{U}^e(\bar{L}(\delta)) - \sum_{i=1}^{n} \bar{v}^i(\bar{L}(\delta))}{\bar{L}(\delta)} \right) = 1.$$

It follows from Lemma 1 that the payment function

$$p^k = \tilde{p}^k + (\lambda^* - \lambda)\,\bar{L}(\delta)$$

solves the static problem (LP-k) with liquidity distribution λ^*.

For a stationary strategy profile σ with action plan $(\overline{\alpha}^k(\bar{L}(\delta)))_k$ and payment plan $(p^k)_k$, we have $U(\sigma) = \bar{U}^e(\bar{L}(\delta))$, $v_i(\sigma) = \bar{v}^i(\bar{L}(\delta))$ and

$$\begin{aligned} \lambda_i^* \bar{L}(\delta) &= \delta \left(\lambda_i + \frac{g_i(\overline{\alpha}^e) - E[p^e|\overline{\alpha}^e] + (\lambda_i^* - \lambda_i)\bar{L}(\delta) - \bar{v}^i(\bar{L}(\delta))}{\bar{L}(\delta)} \right) \\ &= \delta \lambda_i^* \bar{L}(\delta) + \delta \frac{u_i(\sigma) - v_i(\sigma)}{\bar{L}(\delta)}. \end{aligned}$$

This is equivalent to

$$\lambda_i^* \bar{L}(\delta) = \frac{\delta}{1-\delta} \left(\frac{u_i(\sigma) - v_i(\sigma)}{\bar{L}(\delta)} \right).$$

All payments constraints of σ coincide with the payment constraints in the static problems given liquidity allocation $\lambda_i^* \bar{L}(\delta)$. So by construction the payment, budget and action constraints of σ are satisfied.■

Proof of Theorem 2: First we show that the set described in (12) is a subset of the set of PPE payoffs without money burning.

Let $(\alpha^k, p^k)_{k=e,b,1,\ldots,n}$ be a solution of LP-MB. The profile $\sigma = (\alpha^k, p^k)_{k\in K}$ is a stationary equilibrium in the game with money burning with joint equilibrium payoff U and punishment payoffs v_i. It is augmented by a collective punishment state with joint payoff U^b. We now connect σ to the collective punishment state to get a PPE $\widetilde{\sigma}$ without money burning but with the same payoffs as σ. This is done by replacing the money burning by an appropriate choice of transition probabilities between the equilibrium state and the collective punishment state. That is, the structure of the strategy $\widetilde{\sigma}$ differs from the one of σ only in so far as that if in state $k = e, b, 1, \ldots n$ signal y has been realized and no player deviated from the required payments $p^k(y)$, the state changes with a probability $\beta_P^k(y)$ to the collective punishment state and with probability $1 - \beta_P^k(y)$ to the equilibrium state. We define this probability as

$$\beta_P^k(y) = \frac{1-\delta}{\delta} \frac{\sum_{i=1}^n p_i^k(y)}{U - U^b}. \tag{16}$$

Constraints (BC-k) and (MBC-k) tell us that $\beta_P^k(y)$ indeed is a probability. Note that on the equilibrium path of $\widetilde{\sigma}$ there can be repeated stochastic transitions between the equilibrium state and the collective punishment state.

We define the payment function of the strategy $\widetilde{\sigma}$ in state $k = e, b, 1, \ldots, n$ by

$$\widetilde{p}_i^k(y) = p_i^k(y) - \frac{\delta}{1-\delta} \beta_P^k(y)(u_i - u_i^b).$$

Up-front transfers are set to zero. The probabilities $\beta_P^k(y)$ have been chosen such that the payments $\widetilde{p}_i^k(y), i = 1, \ldots, n$ add up to zero. With this definition of payments we have that

$$u_i(\widetilde{\sigma}) = (1-\delta)(g_i(\alpha^e) - E[\widetilde{p}^e(y)|\alpha]) + \delta u_i(\widetilde{\sigma}) + \delta E[\beta_P^e|\alpha^e](u_i^b(\widetilde{\sigma}) - u_i(\widetilde{\sigma}))$$
$$u_i^b(\widetilde{\sigma}) = (1-\delta)(g_i(\alpha^b) - E[\widetilde{p}^b(y)|\alpha]) + \delta u_i(\widetilde{\sigma}) + \delta E[\beta_P^b|\alpha^b](u_i^b(\widetilde{\sigma}) - u_i(\widetilde{\sigma}))$$

reduces to

$$u_i(\widetilde{\sigma}) = u_i \text{ and } u_i^b(\widetilde{\sigma}) = u_i^b.$$

After signal y in state k, continuation payoffs in $\widetilde{\sigma}$ are equal to

$$-(1-\delta)p_i^k(y)+\delta u_i(\sigma).$$

Hence, actions in $\widetilde{\sigma}$ are incentive compatible and the individual punishment payoffs of $\widetilde{\sigma}$ are equal to $v_i(\sigma)$. It is also straightforward to show that payments are incentive compatible. By varying the up-front payments in $\widetilde{\sigma}$ all divisions of the surplus $U(\widetilde{\sigma})$ in which each player gets at least v_i can be achieved. Moreover, the correlation device can be used in the up-front payment state to achieve all joint payoffs between U and U^b.

Second, we show that the set of PPE payoffs without money burning is a subset of the set defined in (12).

Let $\bar{U}$ and $\bar{U}^b$ denote the highest and lowest joint payoff that can be implemented with some PPE in the repeated game without money burning. Similarly, let $\bar{v}_i$ denote the lowest payoff for player i that can be implemented with some PPE. Let σ^e be a PPE with $U(\sigma^e)=\bar{U}$, σ^b a PPE with $U(\sigma^b)=\bar{U}^b$ and for every player i, let σ^i denote a PPE with $u_i(\sigma^i)=\bar{v}_i$. For all $k=e,b,1,...,n$ let α^k be the first (mixed) action profile played on the equilibrium path of σ^k. Note that it always holds true that

$$G(\alpha^b)\leq\bar{U}^b \text{ and } \bar{U}\leq G(\alpha^e).$$

Let $w^k(y)$ denote the vector of continuation payoffs after signal y has been realized in the first period according to $\widetilde{\sigma}^k$ and define

$$p_i^k(y)=\frac{\delta u_i(\sigma^e)-w_i^k(y)}{1-\delta}.$$

That the action, payment and budget constraints are satisfied follows as in the proof of Theorem 1. To see that money burning constraints (MBC-k) are satisfied note that

$$\sum_{i=1}^n p_i^k(y)=\frac{\delta\bar{U}-\sum_{i=1}^n w_i^k(y)}{1-\delta}\leq\frac{\delta}{1-\delta}\left(U-U^b\right).$$

Hence, $(\alpha^k,p^k)_{k=e,b,1,\ldots,n}$ solves LP-MB with value $\bar{U}-\Sigma_{i=1}^n\bar{v}_i-\bar{U}^b$.■

<u>Proof of Proposition 3:</u> The proof proceeds similarly as the proof of Proposition 2 and is therefore omitted.

References

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56:383–396, 1988.

Dilip Abreu, Prajit K. Dutta, and Lones Smith. The folk theorem for repeated games: A neu condition. *Econometrica*, 62(4):939–948, 1994.

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–63, 1990.

Andrew Atkeson. International lending with moral hazard and risk of repudiation. *Econometrica*, 59(4):1069–89, 1991.

Mark B. Cronshaw and David G. Luenberger. Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring and discounting. *Games and Economic Behavior*, 6(2):220–237, 1994.

Katherine Doornik. Relational contracting in partnerships. *Journal of Economics & Management Strategy*, 15(2), 2006.

William Fuchs. Contracting with repeated moral hazard and private evaluations. *The American Economic Review*, 97(4):1432–1448, 2007.

Drew Fudenberg, David I Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, 1994.

Tomas Gal and Josef Nedoma. Multiparametric linear programming. *Management Science*, 18(7):406–422, 1972.

Joseph E. Harrington and Andrzej Skrzypacz. Collusion under monitoring of sales. *The RAND Journal of Economics*, 2007.

Kenneth L. Judd, Sevin Yeltekin, and James Conklin. Computing supergame equilibria. *Econometrica*, 71(4):1239–1254, 2003.

Kenneth M. Kletzer and Brian D. Wright. Sovereign debt as intertemporal barter. *The American Economic Review*, 90(3):621–639, 2000.

Mikhail Klimenko, Garey Ramey, and Joel Watson. Recurrent trade agreements and the value of external enforcement. *Journal of International Economics*, 74(2):475–499, 2008.

Sebastian Kranz. Interactively solving repeated games: A toolbox. 2010.

Sebastian Kranz and Susanne Ohlendorf. Renegotiation-proof relational contracts. *mimeo*, 2009.

Jonathan Levin. Multilateral contracting and the employment relationship. *Quarterly Journal of Economics*, 117:1075–1103, 2002.

Jonathan Levin. Relational incentive contracts. *American Economic Review*, 93:835–857, 2003.

George J. Mailath, Ichiro Obara, and Tadashi Sekiguchi. The maximum efficient equilibrium payoff in the repeated prisoners' dilemma. *Games and Economic Behavior*, 40(1):99–122, 2002.

George Joseph Mailath and Larry Samuelson. *Repeated games and reputations: long-run relationships*. Oxford University Press, 2006.

John Malcomson and W. Bentley MacLeod. Implicit contracts, incentive compatibility and involuntary unemployment. *Econometrica*, 56(2):447–480, 1989.

Luis Rayo. Relational incentives and moral hazard in teams. *Review of Economic Studies*, 74(3):937–963, 2007.