

A Crisis-Bailout Game*

Bruno Salcedo Bruno Sultanum Ruilin Zhou[†]

The Pennsylvania State University

(Draft: November 8, 2014)

Abstract

This paper studies the optimal design of a liability-sharing arrangement as an infinitely repeated game. We construct a schematic, non-cooperative, 2-player model. The active agent can take a costly, unobservable action to try to avert a crisis. Whenever a crisis occurs, each agent decides unilaterally how much to contribute mitigating it. For the one-shot game, when the avoidance cost is too high relative to the expected loss of crisis for the active agent, the first-best is not achievable, i.e., the active agent cannot be induced to put in effort to minimize the incidence of crisis at any static Nash equilibrium. We show that with the same stage-game environment, the first-best cannot be implemented as a perfect public equilibrium (PPE) of the infinitely repeated game either. Instead, at any constrained efficient PPE, the active agent “shirks” infinitely often, and when crisis happens, the active agent is “bailed out” infinitely often. The frequencies of crisis and bailout are endogenously determined at equilibrium. This result of the welfare optimal equilibrium being characterized by recurrent crises and bailouts is consistent with historical episodes of financial crises with varying frequency and varied external responses for troubled institutions and countries in the real world. We explore some comparative statics of the PPEs of the repeated game numerically.

JEL classification: C73 · D82.

Keywords: Bailouts · Moral hazard · Repeated games · Imperfect monitoring · Second-best.

*We thank Ed Green and V. Bhaskar for helpful discussion and comments.

[†]Email addresses: bruno@psu.edu, sultanum@psu.edu, rzhou@psu.edu

1 Introduction

This paper studies the optimal design of an arrangement for sharing the costs of avoiding and mitigating bad outcomes in an infinitely repeated game. When a bad outcome causes losses to several players, it is a public bad, and cost sharing is necessary to induce the ex-post socially optimal level of mitigation. However, if it is ex-ante efficient for some player to make unobservable, costly investment to avert bad outcomes, then sharing the cost of mitigation might undermine that player's incentive to do so. This model is proposed as a schematic model of ongoing relationships among parties (for example, countries in the Euro zone today) that jointly suffer the adverse effects of financial crises that can result from imprudent behavior (that is, by shirking on investment in avoidance) by a member of the group.

In episodes of financial failure, a “*bailout*” occurs when one institution (or country) experiences some financial difficulties or turmoil that causes direct harm there, and when other institutions contribute to the cost of “cleaning up” the mess (for example, by recapitalizing a bank) because they suffer the externalities from the financial failure. Historically, in episodes of financial turmoil, some troubled institutions have been bailed out, and others have not. As a result, the fate of these troubled parties ranges from complete failure/bankruptcy to full recovery.¹ A good positive theory needs to account for this very wide spectrum of outcomes. In view of the diverse outcomes, and of the tension between ex-ante and ex-post efficiency, questions about economic efficiency are especially salient in this context.

The problem has to be modeled as a game between the “crisis-inflicting” party and the potential “help-to-clean-up” party. We construct a schematic, non-cooperative, 2-player model where an active agent takes costly unobservable action to try to avert a crisis, and whenever a crisis occurs, each agent decides unilaterally how much to contribute mitigating it. This is a moral hazard problem regarding the potential action taken by the active agent, and the implied negative externality on the passive agent. The one-shot game is able to generate any combination of avoidance/mitigation patterns as static Nash equilibrium. In particular, when the avoidance cost is too high relative to the expected loss of crisis for the active agent, the first-best cannot be supported as a static Nash equilibrium. That is, the active agent cannot be induced to put in effort

¹For example, while the majority of US companies sink or float on their own, the US government has consistently bailed out large corporations in auto-mobile industry such as GM and Chrysler. Among financial institutions, the Federal Reserve Bank gave a lot of assistance in recent economic crisis in the form of loans and guarantees to large banks like Citigroup and Bank of America, while letting other large financial institutions such as Lehman Brothers and Washington Mutual fail. Among sovereign countries, US government helped Mexico to survive the 1994 Tequila crisis, while many countries suffered huge losses during the 1997 Asian financial crisis with little help from the IMF. The current ongoing Euro area crisis is yet to be fully resolved, with multiple rounds, varying magnitudes of bailouts given to Greece, Italy, Spain, and all other potential problem countries.

to minimize the incidence of crisis.

We then consider the infinite repetition of the one-shot game. We study the perfect public equilibria (PPE) of the repeated game. We show that in the same environment where the active agent shirking is the only static Nash equilibrium of the stage game, the first-best cannot be implemented as a PPE of the infinitely repeated game either. This is because in order to induce the active agent to take the costly avoidance action, the expected mitigation cost for him in case of crisis must be even higher. With both high avoidance cost and high mitigation cost, the active agent is better off not doing anything. In order to compensate the active player for taking the avoidance action *sometimes*, he has to be allowed to shirk other times, and be bailed out (pay less than his share of the mitigation cost) when crises happen.

Based on this intuition, we show that at any constrained efficient PPE, the active agent shirks infinitely often, and when crises happen, the active agent is bailed out infinitely often. As a result, crisis occurs more often compared to the first-best outcome. We demonstrate this kind of equilibrium as well as study comparative statics with numerical examples. To do these, we approximate the constrained efficient allocation with equilibrium allocation of finite-stage automaton. We show that a particular PPE, where the active agent shirks sometimes and is bailed out other times, can yield a welfare level much higher than the repetition of static Nash equilibrium, and not much below that of the first-best. This result of the welfare optimal equilibrium being characterized by recurrent crises and bailouts is consistent with historical episodes of financial crises with varying frequency and varied external responses for troubled institutions and countries in the real world. As in [Green and Porter \(1984\)](#), such a phenomenon reflects an equilibrium that passes recurrently through several distinct states, rather than independent randomization by individual agents. It is a deliberate arrangement of using both occasional shirking and bailout as mechanisms to incentivize good behavior of the active agent as often as possible. We explore some comparative statics of the PPEs of the repeated game numerically.

2 The one-shot game

There are two agents, the active agent 1, and the passive agent 2. The terminology “active” versus “passive” reflects that player 2 will have no role in the first period of the two-period stage game. There are two periods. In period 1, the active agent chooses whether to take a costly action to avert a crisis, $a \in A = \{0, 1\}$. The cost of avoidance is $d > 0$, and not taking the action costs nothing. Agent 1’s first period action a affects the distribution of the state in period 2, $\xi \in X = \{0, 1\}$. If $\xi = 1$, a crisis occurs, and if $\xi = 0$, there is no crisis. The probability of a crisis in period 2 is $\pi^1 \in (0, 1)$ if agent 1

takes the the avoidance action ($a = 1$), and $\pi^0 \in (0, 1)$ otherwise ($a = 0$). That is,

$$\Pr(\xi = 1 | a) = \pi^a, \quad \text{and} \quad \Pr(\xi = 0 | a) = 1 - \pi^a. \quad (1)$$

We assume that $\pi^1 < \pi^0$ so that taking the avoidance action reduces the probability of having a crisis.

In period 2, the state ξ is realized. If there is a crisis ($\xi = 1$), the two agents simultaneously decide how much to contribute to mitigate it. Let m_i denote agent i 's contribution, $m_i \in M = \mathbb{R}_+$. The crisis is mitigated if the total contribution of the two agents $m_1 + m_2$ is no less than 1. Otherwise, each agent i suffers a loss c_i due to the crisis in addition to the cost of his contribution m_i . If there is no crisis, the two agents incur no loss due to crisis, but we still allow the agents to choose positive contributions, which would be completely wasted since there is no crisis to mitigate.² Figure 1 summarizes the structure of the game.

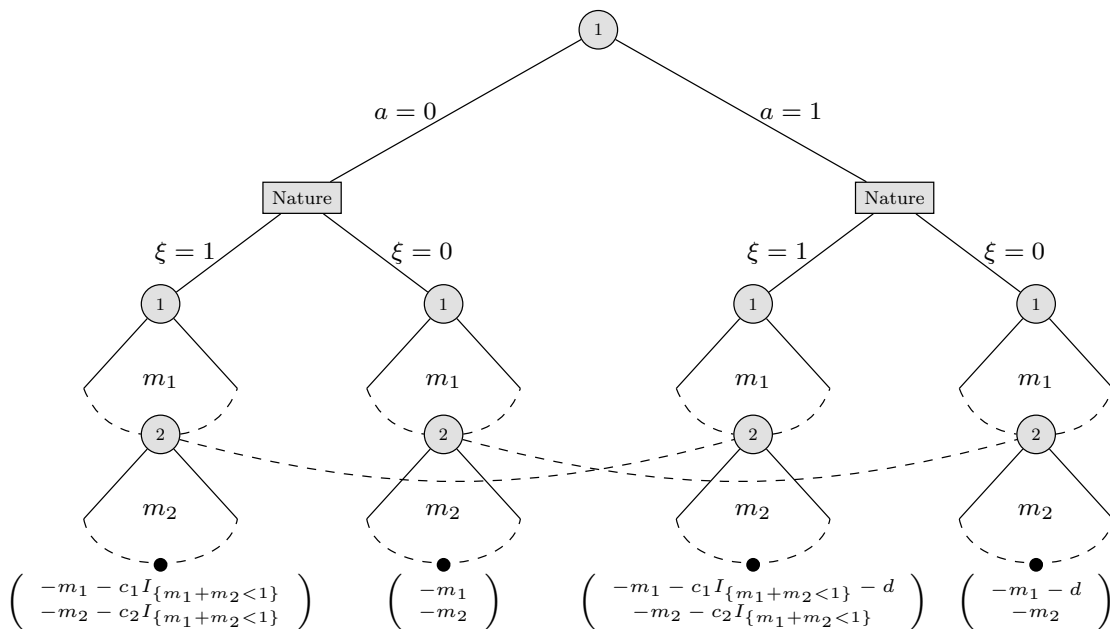


Figure 1: The one-shot game

We are interested in studying the case where mitigation after a crisis is ex-post efficient. So we impose the following assumption.

Assumption 1. For $i = 1, 2$, $c_i \in (0, 1)$, and $c_1 + c_2 > 1$.

²Note that we assume non-transferable utility explicitly here. That is, the two agents' contribution can be used only to mitigate the crisis, and can not go to each other's consumption. This rules out direct subsidy from agent 2 to agent 1 to pay his avoidance cost. In section 7, we show that relaxing this assumption would greatly reduce the difficulty of achieving better allocation in equilibrium.

The assumption implies that neither agent alone will be willing to mitigate the crisis, but together they should, since the cost of mitigation, 1, is less than the total loss if the crisis is not mitigated, $c_1 + c_2$.

2.1 Equilibrium

The structure of the game allows us to restrict attention to pure and public strategies without loss of generality. Players have no use for randomization of the mitigation action because they are risk neutral and M is convex. The active player has no reason to randomize the avoidance action because it is the only action chosen at the first stage. The only private information is the avoidance action of the active player, and he finds no benefit from conditioning his choice on this information. We thus solve for (pure-strategy, public-perfect) Nash equilibria.

Consider period 2 first. Let $u_i(\xi, m_1, m_2)$ denote agent i 's period-2 utility when the state is ξ , and the two agents' contributions are m_1 and m_2 . When there is a crisis ($\xi = 1$), the mitigation game is played. Agent i 's payoff is

$$u_i(1, m_1, m_2) = -m_i - c_i I_{\{m_1+m_2 < 1\}}. \quad (2)$$

Let $\rho_i(m_{-i})$ denote agent i 's response in a crisis state when the other agent's contribution is m_{-i} , $\rho_i: M \rightarrow M$. Then agent i 's optimal response correspondence is

$$\rho_i^*(m_{-i}) = \begin{cases} 0 & \text{if } c_i < 1 - m_{-i} \\ [0, 1 - m_{-i}] & \text{if } c_i = 1 - m_{-i} \\ 1 - m_{-i} & \text{otherwise} \end{cases}. \quad (3)$$

The fixed points of the product correspondence of the two agents' optimal response correspondences determines the Nash equilibria of the period-2 mitigation game, $(m_1(1), m_2(1))$. By assumption 1, there are two types of period-2 equilibria when there is a crisis ($\xi = 1$):

- *No-mitigation*: neither agent contributes anything, $m_1^N(1) = m_2^N(1) = 0$. The period-2 utility of agent i is then $u_i(1, m_1^N(1), m_2^N(1)) = -c_i$.
- *Mitigation*: the two agents jointly contribute 1 unit to mitigate the crisis, $m_1^M(1) \in [1 - c_2, c_1]$, $m_2^M(1) = 1 - m_1^M(1)$. The period-2 utility of agent i in this equilibrium is then $u_i(1, m_1^M(1), m_2^M(1)) = -m_i^M(1)$.

When there is no crisis ($\xi = 0$), there is no loss due to crisis, so period-2 utility of agent i is

$$u_i(0, m_1, m_2) = -m_i. \quad (4)$$

The optimal action is not to contribute anything. So at equilibrium, $m_1(0) = m_2(0) = 0$, and period-2 utility of agent i is $u_i(0, m_1(0), m_2(0)) = 0$.

Now consider agent 1's period-1 problem of whether to take the avoidance action, $a \in A$, assuming that the two agents are going to play their equilibrium strategies in period 2. Let $v_i(a, m_1, m_2)$ denote the expected value of agent i at the beginning of period 1 if agent 1 takes action a in period 1, and the strategy profile in period 2 is $(m_1(\xi), m_2(\xi))_{\xi \in X}$.

$$v_1(a, m_1, m_2) = \sum_{\xi \in X} \Pr(\xi|a)u_1(\xi, m_1(\xi), m_2(\xi)) - ad. \quad (5)$$

$$v_2(a, m_1, m_2) = \sum_{\xi \in X} \Pr(\xi|a)u_2(\xi, m_1(\xi), m_2(\xi)). \quad (6)$$

Agent 1's optimal period-1 action a depends on which of the period-2 equilibrium is to be played in case of crisis. If no-mitigation equilibrium $(m_1^N(1), m_2^N(1))$ is anticipated, then:

$$v_1(a, m_1^N, m_2^N) = \begin{cases} -\pi^0 c_1 & \text{if } a = 0 \\ -d - \pi^1 c_1 & \text{if } a = 1 \end{cases}.$$

and the optimal action a^N is:

$$a^N = \begin{cases} 1 & \text{if } c_1 \geq \hat{d} \\ 0 & \text{otherwise} \end{cases}, \quad (7)$$

where \hat{d} is a rescale of the avoidance cost d adjusting by its impact on the probability of a crisis,

$$\hat{d} \equiv \frac{d}{\pi^0 - \pi^1}. \quad (8)$$

By a similar argument, if a mitigation equilibrium $(m_1^M(1), m_2^M(1))$ is anticipated, then the optimal action a^M is:

$$a^M = \begin{cases} 1 & \text{if } m_1^M(1) \geq \hat{d} \\ 0 & \text{otherwise} \end{cases}. \quad (9)$$

Putting the two-period analysis together, Table 1 summarizes all the equilibria of the one-shot game for different parameter regions, which depend on the magnitude of the adjusted cost of avoidance \hat{d} relative to c_1 and $1 - c_2$. Note that, by Assumption 1, none of these regions is empty and they cover all possible cases. The table only specifies the

range of equilibrium values for agent 1's contribution $m_1(1)$. Agent 2's contribution is given by $m_2(1) = 1 - m_1(1)$ for equilibria with mitigation, and $m_2(1) = 0$, for equilibria without mitigation. The table indicates two important results. First, regardless of the parameter region, both mitigation and no-mitigation equilibrium always coexist.³ Second, every possible combination of avoidance and mitigation occurs in at least one region.

Table 1: Equilibria of the one-shot game

region	parameter values	a	$m_1(1)$	mitigation/not	ex-ante cost
I	$\hat{d} \leq 1 - c_2 < c_1$	1	$[1 - c_2, c_1]$	yes	$d + \pi^1$
		1	0	no	$d + \pi^1(c_1 + c_2)$
II	$1 - c_2 < \hat{d} \leq c_1$	1	$[\hat{d}, c_1]$	yes	$d + \pi^1$
		0	$[1 - c_2, \hat{d}]$	yes	π^0
III	$1 - c_2 < c_1 < \hat{d}$	1	0	no	$d + \pi^1(c_1 + c_2)$
		0	$[1 - c_2, c_1]$	yes	π^0
		0	0	no	$\pi^0(c_1 + c_2)$

2.2 Welfare comparison among coexisting equilibria

The last column of Table 1 lists the ex-ante expected total cost C at each equilibrium, which includes the cost of the avoidance action, and the expected cost of mitigation and/or the loss due to crisis if there is insufficient mitigation,

$$C = ad + \left((1 - a)\pi^0 + a\pi^1 \right) \left[m_1 + m_2 + (c_1 + c_2)I_{\{m_1 + m_2 < 1\}} \right]. \quad (10)$$

We use this cost to discuss the welfare ranking of different coexisting equilibria. An action profile is said to *ex-ante dominate* another one, if it has a lower expected total cost.

We have assumed that when there is a crisis, it is less costly to mitigate than to suffer the consequence ($c_1 + c_2 > 1$), i.e., ex-post mitigation dominates no-mitigation in all parameter regions. In parameter region I, the cost of avoidance is low ($\hat{d} \leq 1 - c_2$), agent 1 takes the avoidance action in any equilibrium regardless of subsequent mitigation decisions, and avoidance/mitigation dominates avoidance/no-mitigation ex-ante. In parameter region II, the cost of avoidance is in the intermediate range ($1 - c_2 < \hat{d} \leq c_1$). Among the three types of equilibria, avoidance/mitigation ex-ante dominates both the

³In each case, there is always a continuum of mitigation equilibrium where the two agents' split of 1 unit of contribution varies over an interval. This is due to the assumption that the mitigation choices are simultaneous. Other game form assumptions would yield a smaller set of equilibria.

no-avoidance/mitigation and the avoidance/no-mitigation equilibria. So in both regions, the ex-ante efficient outcome can be achieved as an equilibrium outcome.

Parameter region III is more problematic. At both types of equilibrium, agent 1 does not take the avoidance action. When $\hat{d} \geq 1$, the no-avoidance/mitigation equilibrium has the lower expected total cost than avoidance/mitigation, so it is ex-ante efficient. When $\hat{d} < 1$, the first-best outcome requires that agent 1 takes the avoidance action, which implies that the first-best outcome cannot be supported as an equilibrium. The rest of the paper is concerned on whether and how repetition of the one-shot game can help to mitigate inefficiency in this case.

3 The repeated game

Time is discrete, dates are indexed by $t \in \{1, 2, \dots\}$. At each date t , the two agents, who have a common discount factor $\delta \in (0, 1)$, play the two-period stage game described above. At the beginning of date t a payoff-irrelevant public signal $\theta_t \in \Theta = [0, 1]$ is observed. These signals are i.i.d. $U[0, 1]^4$. Agent 1 then decides whether to take the avoidance action, $a_t \in A$. The state $\xi_t \in X$ is then drawn independently according to the distribution $\Pr(\xi_t | a_t)$, as given in (1). Both agents observe the realization of state variable ξ_t before simultaneously deciding their respective contributions, $(m_{t1}, m_{t2}) \in M \times M$.

Denote the date- t public information by $h_t = (\theta_t, \xi_t, m_{t1}, m_{t2}) \in H \equiv \Theta \times X \times M \times M$, and let $h^t = (h_1, \dots, h_t) \in H^t$, where $H^0 = \{\emptyset\}$ and $H^t = H^{t-1} \times H$. The public history at stage 1 of date t is $(h^{t-1}, \theta_t) \in H^{t-1} \times \Theta$. The public history at stage 2 of date t is $(h^{t-1}, \theta_t, \xi_t) \in H^{t-1} \times \Theta \times X$.

We focus on perfect Bayesian equilibria where both agents play pure and public strategies. The structure of the game allows us to do this without loss of generality (see Proposition 6 in appendix C). Formally, a public strategy is a sequence of measurable functions $\sigma_1 = \{\alpha_t, \mu_{t1}\}_{t=1}^\infty$ for agent 1, and $\sigma_2 = \{\mu_{t2}\}_{t=1}^\infty$ for agent 2, where $\alpha_t: H^{t-1} \times \Theta \rightarrow A$, and $\mu_{ti}: H^{t-1} \times \Theta \times X \rightarrow M$, $i = 1, 2$. That is, at date t , given a public history h^{t-1} and public signal θ_t , agent 1 takes the avoidance action $\alpha_t(h^{t-1}, \theta_t)$. Then, given the realization of $\xi_t = 1$, the two agents contribute $(\mu_{t1}(h^{t-1}, \theta_t, 1), \mu_{t2}(h^{t-1}, \theta_t, 1))$ to mitigate crisis. As in the one-shot game, we assume that no contribution is made when there is no crisis, i.e., $\mu_{t1}(h^{t-1}, \theta_t, 0) = \mu_{t2}(h^{t-1}, \theta_t, 0) = 0$.

Associated with strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a stochastic stream of payoff vectors, which can be defined as the expected present discounted value of the continuation game after any history. Formally, at any date $t \geq 1$, for any history $h^{t-1} \in H^{t-1}$, let $\sigma|_{h^{t-1}}$ denote the partial function obtained by restricting the domain of σ to histories that extend h^{t-1} ,

⁴The signal θ_t is introduced to convexify the payoff set without explicitly considering randomized strategy for agent 1's action a_t .

which represents the strategy profile induced by σ after the history h^{t-1} . For any public history $h^{t-1} \in H^{t-1}$, any $\theta_t \in \Theta$, denote the expected discounted value for the two agents associated to strategy profile $\sigma|_{h^{t-1}}$ by $V(\sigma|_{h^{t-1}}; \theta_t) = (V_1(\sigma|_{h^{t-1}}; \theta_t), V_2(\sigma|_{h^{t-1}}; \theta_t))$,

$$V_i(\sigma|_{h^{t-1}}; \theta_t) = (1 - \delta) \mathbf{E} \left[\sum_{\tau=t}^{\infty} \delta^{\tau-t} v_i(\alpha_\tau, \mu_{1\tau}, \mu_{2\tau}) \mid h^{t-1}, \theta_t \right] \quad (11)$$

where v_i is the single period payoff function of agent i given by (5)–(6), and the expectation is taken with respect to ξ_t and all future realization of $\{\theta_\tau, \xi_\tau\}_{\tau=t+1}^{\infty}$.

Definition 1. A public strategy profile $\sigma^* = (\alpha^*, \mu_1^*, \mu_2^*)$ is a *perfect public equilibrium* (PPE) if and only if for any $t \geq 1$:

(i) For any $h^{t-1} \in H^{t-1}$, any $\theta_t \in \Theta$, and for any public strategy σ'_1 :

$$V_1(\sigma_1^*|_{h^{t-1}}, \sigma_2^*|_{h^{t-1}}; \theta_t) \geq V_1(\sigma'_1, \sigma_2^*|_{h^{t-1}}; \theta_t). \quad (12)$$

(ii) For any $h = (h^{t-1}, \theta_t, \xi_t) \in H^{t-1} \times \Theta \times X$ and for any pair of public strategies σ'_1 and σ'_2 :

$$\begin{aligned} (1 - \delta)u_1(\xi_t, \mu_{t1}^*(h), \mu_{t2}^*(h)) + \delta \mathbf{E}V_1(\sigma_1^*|_{h^*}, \sigma_2^*|_{h^*}; \theta_{t+1}) \\ \geq (1 - \delta)u_1(\xi_t, \mu'_{t1}(h), \mu_{t2}^*(h)) + \delta \mathbf{E}V_1(\sigma'_1|_{h'_1}, \sigma_2^*|_{h'_1}; \theta_{t+1}), \end{aligned} \quad (13)$$

$$\begin{aligned} (1 - \delta)u_2(\xi_t, \mu_{t1}^*(h), \mu_{t2}^*(h)) + \delta \mathbf{E}V_2(\sigma_1^*|_{h^*}, \sigma_2^*|_{h^*}; \theta_{t+1}) \\ \geq (1 - \delta)u_2(\xi_t, \mu_{t1}^*(h), \mu'_{t2}(h)) + \delta \mathbf{E}V_2(\sigma_1^*|_{h'_2}, \sigma'_2|_{h'_2}; \theta_{t+1}). \end{aligned} \quad (14)$$

where $h^* = (h, \mu_{t1}^*(h), \mu_{t2}^*(h))$, $h'_1 = (h, \mu'_{t1}(h), \mu_{t2}^*(h))$, and $h'_2 = (h, \mu_{t1}^*(h), \mu'_{t2}(h))$.

Condition (12) states that, at the beginning of period t , agent 1's continuation strategy $(\alpha^*|_{h^{t-1}}, \mu_1^*|_{h^{t-1}})$ is optimal from period t onward after any history h^{t-1} . Conditions (13) and (14) state that at period t , after the state ξ_t is realized, both agents' contribution strategies at period t , $(\mu_{t1}^*(h), \mu_{t2}^*(h))$, and continuation strategies $(\alpha^*|_{h^*}, \mu_1^*|_{h^*}, \mu_2^*|_{h^*})$ from period $t + 1$ onward, are optimal, for any history $h = (h^t, \theta_t, \xi_t)$.

Let \mathcal{V} denote the set of PPE payoff vectors.

$$\mathcal{V} = \{V(\sigma) \mid \sigma \text{ is a PPE}\}.$$

This set is non-empty since unconditional repetition of a static Nash equilibrium of the stage game is a PPE, and we have shown that the stage game always has equilibria. We show in appendix C that the payoff set \mathcal{V} is convex and compact and can be characterized

recursively using a modified version of the standard APS decomposition. We use this recursive formulation in the proof of our main results.

As discussed in section 2.2, the first-best outcome can be achieved as an equilibrium of the one-shot game whenever the adjusted cost of avoidance \hat{d} is either sufficiently low or sufficiently high. When \hat{d} is low than c_1 , the private cost of a crisis together with appropriate assigned mitigation cost can induce the active agent to take the avoidance action. When \hat{d} is greater than 1, the social benefit of avoiding a crisis is smaller than its cost. For the rest of the paper we focus on the case in which \hat{d} takes an intermediate value so that the first-best is not achievable in the static one-shot game.

Assumption 2. $c_1 < \hat{d} < 1$.

4 Optimal level of crises and bailouts

Under assumptions 1 and 2, the first-best requires the active agent to always take the avoidance action, and that crises are always mitigated. However, this is not achievable at all the static Nash equilibria at which the active agent do not take the avoidance action. In this section, we study how and to which extent repetition of the one-shot game can help to improve welfare by inducing some good behavior by the active agent. Our first result is that the first-best can *never* be attained: it is not possible to generate incentives for the active agent to behave *every* period.

Given that the first-best is not possible, we then turn our attention to constrained efficient allocation. We investigate the properties of those PPE that are constrained efficient within the set of PPE. We find that, in any constrained efficient PPE, crises are always mitigated, and hence all the welfare loss arises from the impossibility that the avoidance action is taken every period. Also, for sufficiently high discount factors, it is possible to reduce the inefficiency of the static equilibrium by having the active agent take the avoidance action *sometimes*. In order to do so, the passive agent has to bail him out, in a precise sense described further ahead. Both avoidance and bailouts occur infinitely often, and their optimal frequency is determined endogenously.

4.1 The impossibility of implementing the first-best

Under assumption 2, avoidance/mitigation is socially optimal, but agent 1 does not take the avoidance action in any equilibrium of the static game since $c_1 < \hat{d}$. The question is whether infinite repetition of the stage game can give rise to sufficient incentives for agent 1 to always take the avoidance action, despite the fact that he can cheat undetected. In other words, can avoidance/mitigation be sustained as a PPE of the repeated game? Unfortunately, the answer is no.

Proposition 1. *There is no PPE in which the avoidance action is taken at every period.*

Proof. Let $\sigma^* = (\alpha^*, \mu_1^*, \mu_2^*)$ be a PPE and fix a history (h^t, θ_t) with $\alpha_{1t}^*(h^t, \theta_t) = 1$. Define $m_i(\xi) = \mu_{it}^*(h^t, \theta_t, \xi)$, $u^\xi = u_1(\xi, m_1(\xi), m_2(\xi))$ as in equation (2), $V = V_1(\sigma^*|_{h^t, \theta_t, \cdot})$ as 1's expected discounted payoff, and $V^\xi = V_1(\sigma^*|_{h^t, \theta_t, \xi, m(\xi)})$ as the continuation values. So that:

$$V = -(1 - \delta)d + \pi^1 \left(\delta V^1 + (1 - \delta)u^1 \right) + (1 - \pi^1) \left(\delta V^0 + (1 - \delta)u^0 \right). \quad (15)$$

For $\alpha^*(h^t, \theta_t)$ to be optimal, choosing $a = 0$ must be no better than $a = 1$, which implies:

$$\left(\delta V^1 + (1 - \delta)u^1 \right) \leq \left(\delta V^0 + (1 - \delta)u^0 \right) - (1 - \delta)\hat{d}. \quad (16)$$

After $\xi = 1$ is realized, player 1 could guarantee himself a cost of not more than c_1 for this period, and his minimax payoff $-\pi^0 c_1$ from $t + 1$ onwards. Hence, in order for $\mu_{1t}^*(h^t, \theta_t, \xi)$ to be optimal, it must be the case that:

$$\left(\delta V^1 + (1 - \delta)u^1 \right) \geq -(1 - \delta + \pi^0 \delta)c_1. \quad (17)$$

Combining both incentive constraints we have that:

$$\delta \pi^0 \hat{d} \geq - \left(\delta V^0 + (1 - \delta)u^0 \right) - (1 - \delta + \pi^0 \delta) \left(c_1 - \hat{d} \right). \quad (18)$$

Substituting (15) in (16) and doing some algebra:

$$\delta \pi^0 \hat{d} \leq \left(\frac{\delta^2}{1 - \delta} \right) V^0 - \left(\frac{\delta}{1 - \delta} \right) V + \delta u^0. \quad (19)$$

Combining (18) and (19) and doing some more algebra we obtain:

$$V^0 \geq V + \gamma, \quad \gamma \equiv (1 - \delta + \pi^0 \delta) \left(\frac{1 - \delta}{\delta} \right) \left(\hat{d} - c_1 \right). \quad (20)$$

Hence, we have established that whenever the active agent chooses $a = 1$ and there is no crisis, his expected value must increase at least by a fixed factor γ . The assumption $\hat{d} > c_1$ guarantees that $\gamma > 0$. This implies that, if there is no crisis for n subsequent periods and the active agent keeps choosing $a = 1$ with probability 1, then we must have

$$V^n > V + n\gamma, \quad (21)$$

where V^n is agent 1's expected discounted value at period $t + n$. Since the set of feasible payoffs of agent 1 is bounded above by 0, it must be the case that after a long enough

history of no crisis agent 1 takes the non-avoidance action. Otherwise, V^n would be greater than 0. On the other hand, since $\pi^1, \pi^0 \in (0, 1)$, any finite sequence of no crisis occurs with positive probability. Therefore, taking the avoidance action at every period with probability 1 cannot be part of a PPE. \square

The reason that full avoidance cannot be supported as a PPE when the cost of avoidance is too high is informative. The expected discounted value for the active agent V is a mixture of his expected value if there is crisis V^1 and his expected value if there is no crisis V^0 . The only way to generate incentives for avoidance is to guarantee that V^0 is sufficiently greater than V^1 . On the other hand, the mitigation incentive constraint implies that V^1 cannot be too low. As a result, whenever agent 1 chooses $a = 1$ and there is no crisis, his continuation value must increase at least by a fixed amount. Therefore, if there is no crisis for a sufficiently long period, the implied continuation value stops being feasible.⁵

4.2 Efficient mitigation

It is never possible to have avoidance played on every period. However, except for very low discount factors, there exist PPE in which the active agent sometimes takes the avoidance action (Lemma 3). Since there is no static Nash equilibrium with avoidance, incentives must be provided in order for this to be an equilibrium, for instance, by punishing the active agent after a crisis, or by rewarding him if there is no crisis. Two possible ways to punish the active agent after a crisis are to let him suffer the cost of the crisis (no mitigation), or to ask him to contribute more than necessary to mitigate the crisis (money-burning). Our second result is that neither of these forms of punishment schemes are optimal. In every constrained efficient PPE, whenever a crisis happens, both agents contribute exactly as much as is needed to mitigate it.⁶

Proposition 2. *In any constrained efficient PPE, crises are efficiently mitigated, that is, $\mu_{1t}(h^t, \theta_t, 1) + \mu_{2t}(h^t, \theta_t, 1) = 1$ almost surely along the equilibrium path.*

The result is very natural since both of these forms of punishment are ex-post inefficient. However, the proof is far from trivial because, given that there is imperfect monitoring, some degree of inefficiency ex-post could be necessary to generate incentives ex-ante.⁷ In order to establish the result, we need to show that there are always better ways to punish the active agent. That is, given a strategy profile that uses one of these

⁵It is crucial for this proposition that the avoidance action is not observable, see section 6.1. Hence, this is a result of moral hazard and not of the structure of the payoffs.

⁶ We conjecture that a similar result holds when crises don't happen. That is, $\mu_{it}(h^t, \theta_t, 0) = 0$ almost surely along the equilibrium path. However, we don't have a formal proof.

⁷This is a common feature of models with imperfect monitoring that can be traced back to [Green and Porter \(1984\)](#).

forms of punishment, we need to find continuation strategies which improve the payoff to player 2 while generating exactly the same expected continuation payoff for player 1.

This is easy to do when either the total contributions are insufficient to mitigate the crisis, or when both players make positive contributions. In such cases, it is enough to rearrange the contributions in the current period. The only difficult case is when player 2 does not contribute anything and player 1 contributes more than 1. In that case, it is also necessary to change the continuation strategies from the next-period onward. The details of the proof are in appendix A.

4.3 Bailouts as an incentive mechanism for avoidance

The difficulty in inducing player 1 to take the avoidance action is that the avoidance cost is too high for him to pay it on his own. The solution seems to be that agent 2 should help to pay it. If the two agents had transferable utility, we could consider schemes in which agent 2 sometimes subsidizes agent 1.⁸ However, we have made the assumption that each agent's contribution can be used only to mitigate crises, the only way for agent 2 to compensate agent 1 is by sometimes paying more in mitigation cost after a crisis has occurred. When this happens, we call it a 'bailout'.

Definition 2. A *bailout* is the situation where a crisis occurs and the active party — agent 1 — pays less than his crisis loss c_1 in mitigation. In other words, we say that there is a bailout if $\mu_{1t}(h^t, \theta_t, 1) < c_1$ and $\mu_{1t}(h^t, \theta_t, 1) + \mu_{2t}(h^t, \theta_t, 1) \geq 1$.

If agent 1 expects to be bailed out in the future, he may be willing to take the avoidance action, at least in some instances. On the other hand, bailouts are the only form of compensation available. And if agent 1 is not compensated then he has no reason to choose avoidance. It follows that bailouts are necessary to induce avoidance.

Proposition 3. *In any PPE where avoidance action is taken with positive probability, bailout occurs with positive probability.*

Proof. We will show that in a PPE in which there are no bailouts, the active agent always takes the no-avoidance action almost surely. Consider a strategy profile with no bailouts, i.e., such that for almost every history (h^t, θ_t) , either $\mu_{1t}(h^t, \theta_t, 1) \geq c_1$ or $\mu_{1t}(h^t, \theta_t, 1) + \mu_{2t}(h^t, \theta_t, 1) < 1$. This implies that $v_1(\sigma_t(h^t, \theta_t)) \leq -d - \pi^1 c_1$ if $\alpha_t(h^t, \theta_t) = 1$, and $v_1(\sigma_t(h^t, \theta_t)) \leq -\pi^0 c_1$ if $\alpha_t(h^t, \theta_t) = 0$. Assumption 2 implies that $-d - \pi^1 c_1 < -\pi^0 c_1$, and thus $V_1(\sigma) \leq -\pi^0 c_1$, with strict inequality whenever $\alpha_t(h^t, \theta_t) = 1$ for some history set of histories $\{h^t\}$ which is reached with positive probability along the equilibrium path. Individual rationality requires $V_1(\sigma) \geq -\pi^0 c_1$,

⁸Section 6.2 shows that, if subsidies are allowed (agent 2 directly pay agent 1 in consumption goods), and if both agents are sufficiently patient, then the first-best of avoidance/mitigation can be achieved.

and thus it can be satisfied if and only if $\alpha_t(h^t, \theta_t) = 0$ almost surely along the equilibrium path. \square

Proposition 3 shows that bailouts are necessary in order to support avoidance actions, but it says nothing about efficiency. Assumption 2 implies only that more avoidance reduces total expected cost, but this benefit could be outweighed by the cost of generating incentives for it. This issue is addressed by the following proposition.

Proposition 4. *There exists some $\delta^* \in (0, 1)$ such that:*

- (i) *If $\delta < \delta^*$, then every PPE (and therefore every constrained efficient PPE) has avoidance played with probability zero at all periods.*
- (ii) *If $\delta > \delta^*$, then every constrained efficient PPE has avoidance played infinitely often, and bailouts happen infinitely often.*

These results so far provide a relatively complete picture of efficient PPE. For low discount factors, there is nothing to be done, and every PPE essentially reduces to repetition of static Nash equilibria of the stage game. For higher discount factors, in every efficient PPE: crises are always mitigated, there exist both periods of avoidance and periods of non avoidance infinitely often, and the active agent is bailed out infinitely often. Hence the level of avoidance and bailouts are endogenously determined to generate the right incentives that maximize efficiency. In the next section we investigate concrete equilibria to further illustrate how bailouts help to generate incentives for avoidance.

The proof of Proposition 4 makes use of the three lemmas below. The proofs of the lemmas are in appendix B. Lemma 1 simply indicates that whenever it is possible to have avoidance, it is possible to do it in a way that dominates the best static Nash equilibria of the stage game in terms of total cost.

Lemma 1. *If it is possible for agent 1 to choose the avoidance action with positive probability at least once in at least one PPE, then there exists a PPE with total cost less than π^0 , i.e. there exists $(V_1, V_2) \in \mathcal{V}$ such that $V_1 + V_2 > -\pi^0$.*

Lemma 2, is the crucial step of the proof. It states that, if it is possible to have avoidance then, given any non-avoidance PPE, it is possible to improve the payoff of player 2 without affecting the payoff of player 1. Since continuation values of efficient PPE must always lay on the upper boundary of \mathcal{V} , this implies that avoidance must happen infinitely often. By Proposition 3, this implies that bailouts must happen infinitely often as well.

Lemma 2. *If there exists a PPE with total cost less than π^0 , then every PPE without avoidance is Pareto dominated by a different PPE which increases the payoff of agent 2 while keeping the payoff of agent 1 unchanged, i.e. for every PPE σ such that $\alpha_t(h^t) = 0$ almost surely, there exist $(V_1, V_2) \in \mathcal{V}$ such that $V_1 = V_1(\sigma)$ and $V_2 > V_2(\sigma)$.*

Finally, lemma 3 shows that it is not possible to have avoidance when the discount factor is very low, and it is possible to do so when it is very high. This, together with the monotonicity of the set of PPE payoffs with respect to the discount factor, implies the existence of the threshold δ^* –strictly between 0 and 1– separating a region where no avoidance is possible, from a region where avoidance and bailouts happen infinitely often in all efficient equilibria.

Lemma 3. *There exists numbers $0 < \underline{\delta} < \bar{\delta} < 1$ such that, if $\delta \leq \underline{\delta}$ then there is no avoidance in any PPE, and if $\delta > \bar{\delta}$ then there exist PPE in which the avoidance action is played with positive probability at least once.*

Now we are in a position to prove Proposition 4.

Proof of Proposition 4. Let $D \subseteq (0, 1)$ be the set of discount factors for which it is possible to play avoidance. Lemma 3 implies that $D \neq \emptyset$. Lemma 1 implies that for any $\delta \in D$, there exists a PPE with total cost below π^0 . Since the set of PPE payoffs is monotone with respect to the discount factor, the same is true for every $\delta' \geq \delta$. Since obtaining a cost below π^0 requires avoidance, this implies that D is an interval.

Let $\delta^* = \inf D$. We have established that $(\delta_0, 1) \subseteq D$. Lemma 3 implies that $0 < \underline{\delta} \leq \delta^* \leq \bar{\delta} < 1$. For the remainder of the proof fix any $\delta > \delta^*$ and let σ be a constrained efficient PPE given δ .

Suppose towards a contradiction that there exists a set of histories which is reached with positive probability after which there is no avoidance. For every history in the set let σ' be the continuation strategies. By Lemma 1 we know that the consequence of Lemma 2 apply, and hence there exists some PPE σ'' such that $V_1(\sigma'') = V_1(\sigma')$ and $V_2(\sigma'') > V_2(\sigma')$. If we replace σ' with σ'' in every such history, all the incentive constraints of previous periods are (weakly) relaxed. Hence we obtain a PPE which Pareto dominates σ . Therefore we can conclude that if σ is constrained Efficient, there is avoidance infinitely often almost surely. By Proposition 3, this implies that there are also bailouts infinitely often almost surely. \square

5 Numerical examples: equilibrium with endogenously determined crises and bailouts

In this section we approximate the second best by considering PPEs where equilibrium behavior can be described by finite-state automata. An automaton consists of four components: states, an initial distribution over states, a transition rule, and a mapping from states to actions. Let Ω be a finite set of states. States are mapped into actions by $\alpha_1 : \Omega \rightarrow A_i$ and $\mu_i : \Omega \times X \rightarrow M$, $i = 1, 2$. Given the current state ω_t , the active agent takes the action $\alpha_1(\omega_t)$, and, after a crisis state ξ_t is realized, the agents' contributions

are given by $\mu_i(\omega_t, \xi_t)$. After actions are realized, the state ω_{t+1} for the next period is randomly drawn according to the transition rule $\eta : \Omega \times X \times M \times M \rightarrow \Delta(\Omega)$. The initial state for period 1 is drawn according to the initial distribution $\eta_0 \in \Delta(\Omega)$.

An automaton, describes a profile of public strategies for the repeated game. In fact, if we didn't restrict attention to *finite* automata, every profile of public strategies could be described by an automaton (Mailath and Samuelson, 2006, pp. 230). However, for computational reasons, in all of our numerical exercises we restrict attention to finite automata with a fixed upper bound on the number of states in Ω . In what follows, we provide an example in the form of an automaton, where avoidance takes place at some and not at other times.

5.1 An illustrative example of equilibrium mechanism

The goal of this subsection is to illustrate how bailouts can be used to induce avoidance and, thus, improve efficiency. Consider the set of parameters

$$\delta = 0.95, \quad \pi^1 = 0.2, \quad \pi^0 = 0.9, \quad d = 0.5, \quad c_1 = 0.6, \quad c_2 = 0.5.$$

Note that with this set of parameters, the adjusted avoidance cost is higher than the cost for agent 1 and avoidance is socially efficient,

$$c_1 = 0.6 < \hat{d} \approx 0.7143 < 1.$$

The first best has an expected total cost of 0.7, but by Proposition 1 this is not obtainable in any PPE. In contrast, the first-best could be sustained as a PPE if the avoidance action was observable. The lower total cost that can be obtained with unobservable action in a static Nash equilibria is 0.9, which implies a welfare loss (relative to the first-best) of 28.57%.

Figure 2 describes a PPE which minimizes the total expected long-run cost among all PPEs which can be described by a 4-state automata, with one state being the minmax equilibrium. The equilibrium is the following. Agents start in state ω_1 , where the strategy profile is $(a, m_1, m_2) = (1, 0.61, 0.39)$. In this state agent 1 is supposed to take the avoidance action. However, he does not have private incentives to do so since $c_1 = 0.6 < \hat{d} \approx 0.7143$. In order to generate incentives, in case there is no crisis agents go to the state ω_2 , which is a better state for agent 1 than ω_1 because his contribution in mitigating crisis is 0.00 instead of 0.61. In ω_2 the strategy profile is $(1, 0.00, 1.00)$. Again, agent 1 is supposed to avoid the crisis but he does not have private incentives to do it. This time, in order to generate incentives, in case there is no crisis agents go to the state ω_3 , which is a better state for agent 1 than ω_1 because his contribution in mitigating crisis is the same but he does not need to avoid crisis. There is a fifth

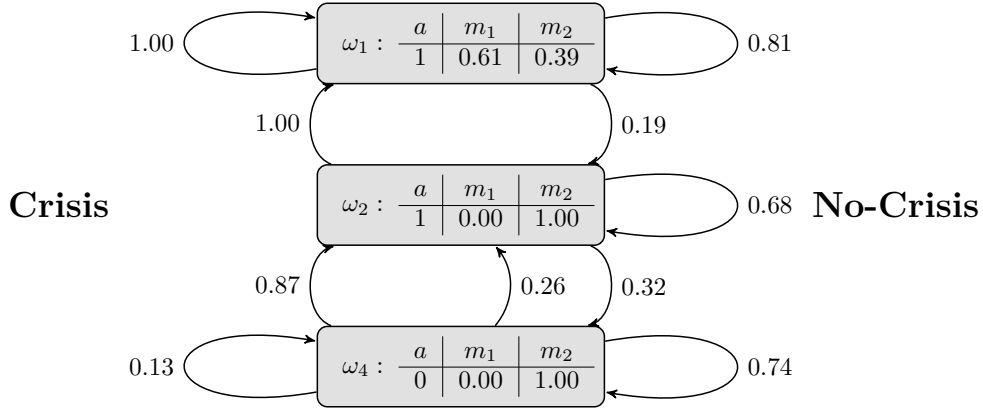


Figure 2: Automaton

state ω_4 , which is not in the figure, with strategy profile of non-avoidance/no-mitigation (the minmax equilibrium). This state is out of the equilibrium path and works as a punishment state in case of a detectable deviation.

Table 2: Summary statistics of the sample PPE

State ω	Invariant distribution	$u_1(\omega)$	$u_2(\omega)$	$V_1(\omega)$	$V_2(\omega)$	Welfare	Welfare loss
ω_1	0.50	-0.622	-0.078	-0.539	-0.177	-0.748	2.23%
ω_2	0.38	-0.500	-0.200	-0.478	-0.250	-0.727	3.87%
ω_3	0.12	-0.000	-0.900	-0.420	-0.328	-0.716	6.86%
LRA	—	-0.501	-0.222	-0.501	-0.222	-0.724	3.41%

Notes: LRA refers to the long run averages, which correspond to the expected values evaluated using the invariant distribution. $u_i(\omega)$ denotes agent i 's expected payoff for the current period when the state is ω . $V_i(\omega)$ denotes agent i 's total discounted expected payoff when the current state is ω .

Although this PPE may not be in the Pareto frontier, it provides a lower bound on what can be achieved by a constrained efficient allocation. In state ω_1 , the expected discounted total cost is 0.748, which is only 2.23% greater than the minimum feasible one 0.7. Furthermore, in states ω_2 and ω_3 , agent 1 contributes $m_0 = 0$ to mitigate the crisis, rather than 0.6 — his loss if there is no mitigation. By our earlier definition, this is a bailout. On average, crisis occurs 28.4% of the time, compared to 20% at the first-best. When crisis does happen, bailout occurs 70.3% of the time.

Player 1 takes the avoidance action in all states except for ω_3 , and, if a crisis occurs, he is bailed out in all states except for ω_1 . In state 1, player 1's discounted expected payoff is close to his minimax payoff, which is $-\pi^0 c_1 = -0.54$. Notice that all transitions happen from one state to the immediate follower. Also, the probability of moving to a state preferred by player 1 is always higher when there is no crisis. Hence, the automata

is reminiscent of the revision strategies used in Rubinstein and Yaari (1983) and Radner (1985). If the likelihood of recent events is higher under the avoidance action, player 1 is rewarded by moving into states in which he does not have to take the avoidance action and is bailed out. Otherwise, he is punished by moving into states in which he has to take the avoidance action and pay a high share of the mitigation cost. Similar patterns are observed with different specifications of payoffs, and different bounds on the number of states.

5.2 Comparative statics

The equilibrium displayed in figure 2 illustrates how bailouts can be used in order to generate avoidance in equilibrium. In this subsection we study how properties of this equilibrium change with key parameters of the model: the avoidance cost, d , the private costs, (c_1, c_2) , and the probabilities of crises, (π^0, π^1) . For each set of parameters we found the PPE which minimizes the total discounted long-run cost among 6 states automata. Then we compare the implied long run probabilities of avoidance, crisis and bailouts, as well as the long run average cost of avoidance and the agents' mitigation payments.

Table 3 displays features of the equilibrium outcome for different values of the avoidance cost d . As one could expect, when the avoidance cost increases, avoidance is played less frequently and, therefore, crisis happens more often. The average avoidance cost also goes up, even though the avoidance action is taken less often. What is more surprising is that the incidence of bailouts increases and agent 1 mitigation cost decreases. This changes reflect the structure of the equilibrium. Bailouts are the mechanism by which agent 2 compensates agent 1 for bearing the avoidance cost alone. As d increases the compensation needs to be higher in order to generate enough incentives for avoidance to be played.

Table 3: $\delta = 0.9$, $\pi^1 = 0.2$, $\pi^0 = 0.9$, $c_1 = 0.6$, $c_2 = 0.5$

d	$\mathbb{P}(a = 1)$	$\mathbb{P}(\xi = 1)$	$\mathbb{P}(m_1 < c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost ²
0.45	0.9569	0.2302	0.4647	0.4306	0.0784	0.1518	101.66%
0.50	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08%
0.55	0.7598	0.3682	0.8942	0.4179	0.0182	0.3500	104.81%
0.60	0.7115	0.4019	0.9301	0.4269	0.0143	0.3877	103.61%
0.60	0.6851	0.4204	0.8727	0.4453	0.0227	0.3977	101.85%

Note: The probabilities and expectations are evaluated using the implied invariant distribution.

² Expressed as percentage of the first-best expected total cost.

Table 4 displays features of the equilibrium outcome for different values of (c_1, c_2) . Agent's 2 private cost of crisis, c_2 , has a very small effect on the equilibrium outcome. On

the other hand, the impact of c_1 on the equilibrium outcome is substantial. Increasing c_1 from 0.55 to 0.65 reduces long run average cost from about 106.4% to 101.8% of the first-best; the incidence of crisis is reduced in approximately one third (from 0.36 to 0.24); and bailouts happens only 60% of the time instead of 89%. When c_1 and c_2 increases, both agents' min max payoff decreases. But decreasing the min max payoff of agent 2 has little effect on the equilibrium outcome since he is the passive agent — the difficult is to generates incentives for agent 1.

Table 4: $\delta = 0.9$, $\pi^1 = 0.2$, $\pi^0 = 0.9$, $d = 0.5$

c_1	c_2	$\mathbb{P}(a = 1)$	$\mathbb{P}(\xi = 1)$	$\mathbb{P}(m_1 < c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost ²
0.55	0.5	0.7763	0.3566	0.8870	0.3882	0.0177	0.3388	106.39%
	0.7	0.7729	0.3590	0.8875	0.3864	0.0176	0.3414	106.49%
0.60	0.5	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08%
	0.7	0.8529	0.3029	0.7119	0.4265	0.0391	0.2638	104.20%
0.65	0.5	0.9363	0.2446	0.6136	0.4681	0.0655	0.1791	101.82%
	0.7	0.9349	0.2456	0.5420	0.4675	0.0722	0.1733	101.86%

Note: The probabilities and expectations are evaluated using the implied invariant distribution.

² Expressed as percentage of the first-best expected total cost.

Table 5 displays the long run expected total cost above the first-best for different combinations of (π^0, π^1) . The other parameter are set to $\delta = 0.95$, $d = 0.5$, $c_1 = 0.6$ and $c_2 = 0.5$. The symbol – represents the combinations of (π^0, π^1) that do not satisfy assumption 2. From section 2, we know that the first-best is achieved in the static game when $c_1 \geq \hat{d} = d/(\pi^0 - \pi^1)$. When $\hat{d} \geq 1$, non avoidance is optimal and the first-best is also achieved in the static game by non-avoidance/mitigation. These possibilities are represented in the left lower corner and right higher corner of table 5. The effect of (π^0, π^1) on the welfare cost is not uniform. Combinations of (π^0, π^1) , with \hat{d} either closer to c_1 or 1, lead to lower cost. This means that sometimes decreasing π^0 reduces the welfare cost, sometimes increasing π^0 reduces the welfare cost. The same is true for π^1 . On the other hand, for combinations of (π^0, π^1) with the same \hat{d} (*i.e.* $\pi^0 - \pi^1$ constant), higher π^0 and π^1 always lead to a lower welfare cost. The interpretation of this results is not simple. One could think that higher π^1 means that avoidance is less effective in preventing crisis, which could imply a higher cost, but this is not true. The correct measure is $\pi^0 - \pi^1$, how much the probability of crisis decreases by the avoidance action. With $\pi^0 - \pi^1$ hold constant, the only impact is of increasing π^0 . Higher π^0 implies that the min max utility of agent 1 is lower, hence, it is easier to generate incentives for avoidance.

Table 6 displays features of the equilibrium outcome for different values of the discount rate δ . As one could expect, lower δ is associated with lower welfare. Increasing δ from 0.6 to 0.9 decreases the total cost in about 4% of the first-best. This pattern

Table 5: Total cost above the first-best (%)¹

$\pi^1 \backslash \pi^0$	0.65	0.70	0.75	0.80	0.85	0.90	0.95
0.10	3.25	5.71	8.49	9.05	8.00	2.94	–
0.15	–	2.70	4.65	6.46	6.59	4.41	1.28
0.20	–	–	2.20	3.68	4.67	4.54	2.74
0.25	–	–	–	1.75	2.89	3.30	2.89
0.30	–	–	–	–	1.38	2.19	2.26

¹ Long run average as percentage of the first-best.

reflects that when δ is high agents are more willing to cooperate since they care more about punishments in the future.

Table 6: $\pi^1 = 0.2$, $\pi^0 = 0.9$, $c_1 = 0.6$, $c_2 = 0.5$, $d = 0.5$

δ	$\mathbb{P}(a = 1)$	$\mathbb{P}(\xi = 1)$	$\mathbb{P}(m_1 < c_1)$	$\mathbb{E}(d)$	$\mathbb{E}(m_1)$	$\mathbb{E}(m_2)$	Expected Total Cost ²
0.6	0.7111	0.4022	0.5978	0.3555	0.1063	0.2960	108.25%
0.7	0.7630	0.3659	0.8345	0.3815	0.0636	0.3023	106.77%
0.8	0.7979	0.3415	0.7491	0.3989	0.0444	0.2970	105.78%
0.9	0.8571	0.3000	0.7937	0.4285	0.0340	0.2660	104.08%

Note: The probabilities and expectations are evaluated using the implied invariant distribution.

² Expressed as percentage of the first-best expected total cost.

6 Alternative mechanisms

We have shown that the first-best cannot be achieved as a PPE of the repeated game, and that every constrained efficient PPE has bailouts infinitely often. In this section we show that these results depend crucially on the assumptions that the avoidance action is not perfectly observed, and on the assumption that the only way to compensate the active agent is through mitigation. In particular, we consider two alternative models, one in which the avoidance action is perfectly observed, and one in which the passive agent can directly subsidize the active agent. In both cases, when agents are sufficiently patient, there is a PPE which achieves the first-best.

6.1 The avoidance action is observable

Suppose that a is perfectly observable, so that mitigation contributions and continuation strategies can depend not just on the realization of the crisis state, but also on the action taken by the active player. Consider the grim trigger strategy profile described as follows. Along the equilibrium path, active agent always chooses $a = 1$, and the agents contribute m_2 and $m_1 = 1 - m_2$ if there is a crisis and 0 otherwise. If there is

any deviation, then the active agent chooses $a = 0$ forever and both agents never again make positive mitigation contributions.

We will show that for sufficiently high discount factor, there exists one such grim trigger strategy which is a PPE. Since there is always avoidance and mitigation along the equilibrium path, this strategy profile implements the first-best.

The single-deviation principle still applies, and hence it suffices to search for deviations at a single period. The strategies after a deviation are repetition of a static Nash equilibrium and thus constitute a PPE. Along the equilibrium path, it is optimal for the active agent to choose the avoidance action if and only if:

$$-(d + \pi^1 m_1) \geq -\pi^0 c_1. \quad (22)$$

The mitigation incentive constraint for player 1 is:

$$-(1 - \delta)m_1 - \delta(d + \pi^1 m_1) \geq -(1 - \delta)c_1 - \delta\pi^0 c_1, \quad (23)$$

and for player 2 it is:

$$-(1 - \delta)m_2 - \delta\pi^1 m_2 \geq -(1 - \delta)c_2 - \delta\pi^0 c_2. \quad (24)$$

The mitigation constraint for player 3 is satisfied if and only if:

$$m_2 \leq \left(\frac{1 - \delta + \delta\pi^0}{1 - \delta + \delta\pi^1} \right) c_2 \quad (25)$$

Hence, we can choose m_2 as to satisfy this condition with equality, and incentives for player 2 are automatically satisfied. Then, using the fact that $m_1 + m_2 = 1$, we have that:

$$m_1 = 1 - \left(\frac{1 - \delta + \delta\pi^0}{1 - \delta + \delta\pi^1} \right) c_2 < 1 - c_2 < c_1 \quad (26)$$

This implies that $c_1 - m_1 > 0$, and thus, after some simple algebra, the mitigation constraint for 1 is satisfied whenever the avoidance constraint is satisfied. This implies that the proposed strategy profile is a PPE if and only if:

$$d + \pi^1 - \pi^1 \left(\frac{1 - \delta + \delta\pi^0}{1 - \delta + \delta\pi^1} \right) c_2 \leq \pi^0 c_1 \quad (27)$$

Which, after some algebra, is satisfied if and only if:

$$\delta \geq \frac{d + \pi^1(1 - c_2) - \pi^0 c_1}{d + \pi^1(1 - c_2) - \pi^0 c_1 + \pi^1(d + \pi^0(c_1 + c_2) - \pi^1)} \quad (28)$$

6.2 Monetary transfers

Up to now, we have assumed that the two agents have non-transferable utility. Specifically, by (2), (4), (5) and (6), each period if agent 1 takes the avoidance action, he has to pay the cost d by himself. If there is a crisis, the two agents' contributions m_1 and m_2 are only for mitigation, and if there is no crisis, any positive contribution by the two agents is wasted. Suppose we relax this assumption allowing agents to make transfer to each other when there is no crisis. That is, at any date t , after any history h^t , when there is no crisis, if agent i makes strictly positive contribution $\mu_{ti}(h^t, 0)$, it goes to agent $-i$ for consumption.

Now consider the following strategy profile. For any $t \geq 1$, any $h^t \in H^{t-1}$ in the equilibrium path,

$$\begin{aligned}\alpha(h^t) &= 1, \\ \mu_{t1}(h^t, 1) &= m_1, \quad \mu_{t2}(h^t, 1) = 1 - m_1 \\ \mu_{t1}(h^t, 0) &= 0, \quad \mu_{t2}(h^t, 0) = b\end{aligned}\tag{29}$$

where $m_1 \geq 0$, and $b \geq 0$ will be specified later. That is, along the equilibrium path, agent 1 always takes the avoidance action, and contributes m_1 when there is a crisis, and agent 2 compensates agent 1 with b units of consumption when there is no crisis. In other words, b can be viewed as the subsidy to agent 1 from agent 2 in no-crisis time. In case of a detectable deviation, the agents switch to play the one-shot Nash equilibrium with no-avoidance and no mitigation forever. This strategy profile implies the following discounted expected lifetime utility for the two agents,

$$\begin{aligned}V_1 &= -(d + \pi^1 m_1) + (1 - \pi^1)b, \\ V_2 &= -\pi^1(1 - m_1) - (1 - \pi^1)b\end{aligned}$$

For the strategy profile (α, μ_1, μ_2) described to be an equilibrium strategy, agent 1's payoff V_1 has to satisfy the following incentive constraints,

$$-d - \pi^1 m_1 + (1 - \pi^1)b \geq -\pi^0 m_1 + (1 - \pi^0)b\tag{30}$$

$$-(1 - \delta)m_1 + \delta V_1 \geq -(1 - \delta)c_1 - \delta\pi^0 c_1\tag{31}$$

and agent 2's payoff V_2 has to satisfy the following incentive constraints,

$$-(1 - \delta)b + \delta V_2 \geq -\delta\pi^0 c_2\tag{32}$$

$$-(1 - \delta)(1 - m_1) + \delta V_2 \geq -(1 - \delta)c_2 - \delta\pi^0 c_2\tag{33}$$

It is easy to show that (m_1, b) can be chosen such that all incentive constraints

(30)–(33) are satisfied if δ is sufficient high. In particular, let

$$m_1 = \pi^0 c_1, \quad b = \hat{d} - \pi^0 c_1 \quad (34)$$

That is, in crisis time agent 1 contributes $\pi^0 c_1$, and in no-crisis time agent 2 compensates agent 1 exactly the amount that makes him indifferent between avoidance and non-avoidance action. At this (m_1, b) , all incentive constraints (30)–(33) are satisfied for high δ , hence the corresponding strategy implements avoidance/mitigation as an equilibrium of the repeated game.

This subsidy in no-crisis-time scheme is simple theoretically, but may not be easy to implement in reality. For example, it might be difficult to justify paying Greece’s government every period—subsidy in normal time and mitigation in crisis time—to the public!

References

- Abreu, D., Pearce, D., and Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063.
- Cronshaw, M. B. and Luenberger, D. G. (1994). Strongly symmetric subgame perfect equilibria in infinitely repeated games with perfect monitoring and discounting. *Games and Economic Behavior*, 6(2):220–237.
- Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100.
- Mailath, G. J. and Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. Oxford University Press.
- Radner, R. (1985). Repeated principal-agent games with discounting. *Econometrica*, 53(5):1173–1198.
- Rubinstein, A. and Yaari, M. E. (1983). Repeated insurance contracts and moral hazard. *Journal of Economic Theory*, 30(1):74–97.

A Proof of proposition 2

Crises are mitigated First we show that in all constrained efficient PPE crises are mitigated, i.e., $\mu_{1t}(h^t, 1) + \mu_{2t}(h^t, 1) \geq 1$, almost surely along the equilibrium path. To see this, let σ^* be a PPE such that $\mu_{1t}^*(h^t, 1) + \mu_{2t}^*(h^t, 1) < 1$ for every history h^t in some set of histories $\{h^t\}$ reached with positive probability. Consider the alternative strategy profile σ' with $\mu'_{1t}(h^t, 1) = \mu_{1t}^*(h^t, 1) + c_1$ and $\mu'_{2t}(h^t, 1) = \max\{0, 1 - \mu'_{1t}(h^t, 1)\}$ for all such histories, and which coincides with σ^* everywhere else. We claim that σ' is also a PPE and Pareto dominates σ^* .

Notice that the cost for player 1 according to σ' and σ^* coincide in every period after any history. The cost for 2 is also the same for every period and history other than h^t . If there is a crisis in period t and given h^t , the cost for 2 corresponding to σ^* is $-\mu_{2t}^*(h^t, 1) - c_2$, and according to σ' it is:

$$-\mu'_{2t}(h^t, 1) = -\max\{0, 1 - \mu'_{1t}(h^t, 1)\} \leq -\max\{0, 1 - c_1\} < -c_2 \leq -c_2 - \mu_{2t}^*(h^t, 1).$$

This implies that σ' Pareto dominates σ^* , and that, in order to verify that σ' is an equilibrium, we only need to check that there are no profitable deviations after $(h^t, 1)$.

For 2 there are no profitable deviations because $\mu'_{2t}(h^t, 1)$ is a static best response to $\mu'_{1t}(h^t, 1)$. For 1 there are two possibilities. If $\mu'_{2t}(h^t, 1) > 0$ then $\mu'_{1t}(h^t, 1)$ is already a static best response to $\mu'_{2t}(h^t, 1)$. If $\mu'_{2t}(h^t, 1) = 0$, then the best possible deviations where also available given σ^* . Since σ^* is an equilibrium, then so is σ' .

No excessive mitigation Now we will show that there is no money burning after a crises occurred in any constrained efficient PPE, i.e., $\mu_{1t}(h^t, 1) + \mu_{2t}(h^t, 1) \leq 1$ almost surely along the equilibrium path. Let σ^* be a PPE such that $\mu_{1t}(h^t, 1) + \mu_{2t}(h^t, 1) > 1$ after a set of positive histories $\{h^t\}$ that is reached with positive probability. Fix any such history h^t and let $\sigma^{**} = \sigma^*|h^t$. Further ahead, we will show that there exists some $V' \in \mathcal{V}$ such that $V'_1 = V_1(\sigma^{**})$ and $V'_2 > V_2(\sigma^{**})$. Now consider the alternative strategy profile σ' that plays according to σ^* before any history in $\{h^t\}$ is reached, but after one such history h^t is reached, it switches to the strategy profile which generates V' . Since the payoffs for player 1 remain unchanged, and the payoffs for player 2 are increased, σ' is also a PPE and it Pareto dominates σ^* .

It remains to show that, given σ^{**} , there exists some $V' \in \mathcal{V}$ such that $V'_1 = V_1(\sigma^{**})$ and $V'_2 > V_2(\sigma^{**})$. For that purpose we use the APS decomposition. From proposition 5, it follows that $V \in \mathcal{V}$ if and only if there exists an action profile $(a, (m_i^\xi))$, and a profile of continuation values $W^\xi = (W_i^\xi)_{i \in I} \in \mathcal{V}$ satisfying mitigation and avoidance incentive constraints, and such that $V = (1 - \delta)u(a, m_1, m_2) + \delta w$. Notice that we can write:

$$V_i = -(1 - \delta)ad + \pi^a [\delta W_i^1 - (1 - \delta)m_i^1] + (1 - \pi^a) [\delta W_i^0 - (1 - \delta)m_i^0].$$

Since mitigation deviations are observable they can be punished radically and hence the mitigation constraints can be written as:

$$-(1 - \delta)m_i^1 + \delta W_i^1 \geq -(1 - \delta + \delta\pi^0)c_i, \quad \text{and} \quad -(1 - \delta)m_i^0 + \delta W_i^0 \geq -\delta\pi^0 c_i.$$

If $a = 1$, then the avoidance constraint can be written as:

$$[\delta W_1^0 - (1 - \delta)m_1^0] - [\delta W_1^1 - (1 - \delta)m_1^1] \geq (1 - \delta)\hat{d},$$

and if $a = 0$ the only difference is that the sign of the inequality must be reversed. The thing to notice here is that both the total expected discounted value and the incentive constraints, only depend on $\delta W_i^\xi - (1 - \delta)m_i^\xi$.

Now let $a = \alpha_1^{**}(\emptyset)$ and $m_i^\xi = \mu_{i1}^{**}(\emptyset, \xi)$ denote the profile of first period strategies. And let $w_i^\xi = V_i(\sigma^{**}|\xi, \sigma_1^*(\xi))$ denote the equilibrium continuation values.

First suppose that $m_2^1 > 0$ or $m_2^0 > 0$, then it suffices to replace with $\tilde{m}_2^1 = \max\{0, 1 - m_1^1\}$. Hence for the rest of the proof we assume that $m_1^1 > 1$ and player 2 doesn't make any contributions. Since the crisis is mitigated this implies that $u_2 = 0 > V_2$, and, consequently:

$$V_2 > W_2 \geq \min\{W_2^1, W_2^0\}. \quad (35)$$

We will use this fact to construct new continuation values that support the desired outcomes. The construction depends on three different cases.

The first case is that both $m_1^1 > 1$ and $m_1^0 > 0$. This gives us freedom to adjust the contribution of player 1 independently of ξ . Let ξ be a crisis state such that $V_2 > W_2^\xi$, and consider the alternative strategy profile obtained by setting:

$$\tilde{W}^\xi = (1 - \varepsilon)W^\xi + \varepsilon V \in \mathcal{V}, \quad \text{and} \quad \tilde{m}_1^\xi = m_1^\xi + \frac{\delta\varepsilon}{1 - \delta}(V_1 - W_1^\xi). \quad (36)$$

Simple algebra shows that:

$$(1 - \delta)m_1^\xi + \delta W_1^\xi = (1 - \delta)\tilde{m}_1^\xi + \delta\tilde{W}_1^\xi. \quad (37)$$

Which implies that neither the expected payoff nor the incentives change for player 1. However, the continuation value for player 2 has improved.

The second case is when $m_1^0 = 0$, and $a = 1$. In this case, the fact that $m_1^1 > 1 > \hat{d} = \hat{d} + m_1^0$ implies that incentives won't change if we set $\tilde{W}^\xi = W$. After this transformation, we know that $u_2 > V_2 > W = W_2^\xi$ for both ξ , and the previous transformation with $\xi = 1$ does the trick.

Finally, suppose that $m_1^0 = 0$, and $a = 0$. If $V_2 > W_2^1$ then the previous transformation with $\xi = 1$ works. Since we need $\tilde{m}_1^0 \geq 0 = m_1^0$, applying the previous transformation for $\xi = 0$ is only feasible if and only if $V_1 \geq W_1^0$. Hence we only need to worry about the case when $W_1^0 > V_1$ and $W_2^1 \geq V_2$.

What we can do here is try to change W^0 and m_1^1 in such a way that V_1 is preserved and incentives are preserved. Moving W_1^0 towards V_1 makes player 1 worse off, hence, to preserve 1's payoffs we would need to make him better off by reducing m_1^1 . Reducing m_1^1 does not affect the mitigation constraint, and makes avoiding the crisis less important for player 1, hence incentives are preserved. Hence in this last sub-case it suffices to set

\tilde{m}_1^1 so that:

$$\tilde{m}_1^1 = m_1^1 - \frac{\delta(1 - \pi^0)\varepsilon}{(1 - \delta)\pi^0}(W_1^0 - V_1). \quad (38)$$

Which implies that $(1 - \delta)\Delta u = -\delta\Delta W_1$ and thus $\Delta V_1 = 0$ \square

B Missing steps in proof of proposition 4

We start with some notation and preliminaries. Throughout the proof we only use strategy profiles with efficient mitigation, i.e. strategy profiles such that $\mu_{1t}(h^t, 0) = \mu_{2t}(h^t, 0) = 0$ and $\mu_{1t}(h^t, 1) + \mu_{2t}(h^t, 1) = 1$.⁹ A strategy profile for the pure game without money burning can be described by a tuple (a, m_1, m_2) where m_i represents i 's contribution in the event of a crisis. We use the notation $\langle a, m_1, m_2 \rangle$ to denote the strategy of the repeated game which consists of repeating (a, m_1, m_2) after any arbitrary history. Notice that $\langle 0, m_1, 1 - m_1 \rangle$ is a PPE as long as $m_1 \in [1 - c_2, c_1]$.

Given a strategy profile $\sigma = (\alpha, \mu_1, \mu_2)$, in order to analyze the incentive constraints at period t given h^t , we use the notation $m_i = \mu_{it}(h^t, 1)$ to denote the mitigation contributions if there is a crisis at period t , and $V_i^\xi = V_i(\sigma|_{h^t, \xi})$ for the continuation values. Also, we only consider strategies which go to the minimax equilibrium $\langle 0, 0, 0 \rangle$ after any observable deviation. With this in mind, and with mitigation and no money burning, the incentive constraints for mitigation if $m_i > 0$ can be written as:

$$-(1 - \delta)m_i + \delta V_i^1 \geq -(1 - \delta + \delta\pi^0)c_i \quad (39)$$

If $a = 1$, then the incentive constraint for avoidance can be written as:

$$\delta(V_1^0 - V_1^1) \geq (1 - \delta)(\hat{d} - m_1) \quad (40)$$

Proof of lemma 1. Suppose that σ is a PPE and after some history we have $\alpha_t(h^t) = 1$. There are two possibilities, if the continuation values are such that either $V_1^0 + V_2^0 > -\pi^0$ or $V_1^1 + V_2^1 > -\pi^0$ then the continuation strategies satisfy the conclusion of the lemma and we are done. Otherwise, we must have $V_1^0 + V_2^0 \leq -\pi^0$ and $V_1^1 + V_2^1 \leq -\pi^0$. This implies that V^1 and V^0 can be achieved by some strategy profiles $\langle 0, m_1^0, m_2^0 \rangle$ and $\langle 0, m_1^1, m_2^1 \rangle$ respectively, i.e. we must have $V_i^\xi = -\pi^0 m_i^\xi$ for $i = 1, 2$ and $\xi = 0, 1$. Also, since we are trying to minimize total cost and increasing the payoff of agent 2 only relaxes the incentive constraints, we can assume without loss of generality that $m_2 = 1 - m_1$, and $m_2^\xi = 1 - m_1^\xi$. Substituting in (39) and (40), the incentive constraints

⁹We don't change the game as to preclude the possibility of burning money or no-mitigation. We simply don't use such strategies in our constructions.

for agent 1 at period t given h^t yield:

$$\left. \begin{aligned} (1 - \delta)m_1 + \delta\pi^0 m_1^1 &\geq (1 - \delta)\hat{d} + \delta\pi^0 m_1^0 \\ (1 - \delta)m_1 + \delta\pi^0 m_1^1 &\leq (1 - \delta + \delta\pi^0)c_1 \end{aligned} \right\} \Rightarrow (1 - \delta)\hat{d} + \delta\pi^0 m_1^0 \leq (1 - \delta + \delta\pi^0)c_1$$

Individual rationality requires that $V_2^0 = -\pi^0(1 - m_1^0) \geq -\pi^0 c_2$, which in turn implies that $m_1^0 \geq 1 - c_2$ and thus:

$$(1 - \delta)\hat{d} + \delta\pi^0(1 - c_2) \leq (1 - \delta + \delta\pi^0)c_1$$

After some simple algebra, this implies that $\delta \geq \underline{\delta}$ and therefore, by Lemma 3, σ^0 is a PPE and $(V_1(\sigma^0), V_2(\sigma^0)) \in \mathcal{V}$. Finally, notice that $V_1(\sigma^0) + V_2(\sigma^0) = -(1 - \delta)(d + \pi^1) - \delta\pi^0 > -\pi^0$, where the inequality follows because assumption 2 implies that $d + \pi^1 < \pi^0$. \square

Proof of lemma 2. We will show that there exist two PPE σ^1 and σ^2 such that $V_1(\sigma^1) = -\pi^0 c_1$ and $V_2(\sigma^1) > -\pi^0(1 - c_1)$, and $V_1(\sigma^2) > -\pi^0(1 - c_2)$ and $V_2(\sigma^2) = -\pi^0 c_2$. Now consider any $m_1 \in [1 - c_2, c_1]$ and the strategy profile $\sigma = \langle 0, m_1, 1 - m_1 \rangle$. The corresponding payoff vector is $V_i(\sigma) = (-\pi^0 m_1, -\pi^0(1 - m_1))$. Since $1 - c_2 \leq m_1 \leq c_1$, it follows that $V_1(\sigma^1) = -\pi^0 c_1 \leq -\pi^0 m_1 \leq -\pi^0(1 - c_2) < V_1(\sigma^2)$, and hence we can write $V_1(\sigma) = \lambda V_1(\sigma^1) + (1 - \lambda)V_1(\sigma^2)$ with $\lambda \in (0, 1]$. Now consider the strategy profile $\tilde{\sigma}$ that plays σ^1 with probability λ and σ^2 with probability $(1 - \lambda)$. By construction we have that $V_1(\tilde{\sigma}) = V_1(\sigma)$. Also we know that $V_1(\tilde{\sigma}) + V_2(\tilde{\sigma}) > -\pi^0 = V_1(\sigma) + V_2(\sigma)$, and therefore $V_2(\tilde{\sigma}) > V_2(\sigma)$. Hence it only remains to show the existence of σ^1 and σ^2 .

Let σ^* be a PPE with total expected discounted cost less than π^0 , i.e. such that $V_1(\sigma^*) + V_2(\sigma^*) > -\pi^0$. Individual rationality implies that $V_i(\sigma^0) \geq -\pi^0 c_i$ for $i = 1, 2$, and we assume that both the inequalities are strict (the proof can be easily adapted to accommodate the general case). Hence we can write $V_1(\sigma^*) = -\pi^0(m_1 - \Delta_1)$ and $V_2(\sigma^*) = -\pi^0(1 - m_1 - \Delta_2)$, for some $m_1 \in [1 - c_2, c_1]$ and $\Delta_1, \Delta_2 > 0$. Also, notice that by mixing σ^* with $\langle 0, m_1, 1 - m_1 \rangle$ we obtain a new PPE with payoffs between those of σ^* and $\langle 0, m_1, 1 - m_1 \rangle$. Hence we can assume without loss of generality that Δ_1, Δ_2 are arbitrarily small.

For $\varepsilon \in (0, 1 - c_1)$, consider the strategy profile σ^ε defined as follows:

- $\sigma_1^\varepsilon = (0, c_1 + \varepsilon, 1 - c_1 - \varepsilon)$
- It $\xi_1 = 1$ then $\sigma_t^\varepsilon = \sigma_{t-1}^1$ for all $t > 1$
- It $\xi_1 = 0$ then $\sigma_t^\varepsilon = (0, 1 - c_2, c_2)$ for all $t > 1$

The continuation strategies after period 1 are PPE. In period 1, agent 2 plays a static best response, and his contributions do not affect the continuation value. Hence we

only need to check the incentives for agent 1 in period 1. The corresponding mitigation constraint (39) can be written as:

$$\begin{aligned} & -(1 - \delta)(c_1 + \varepsilon) - \delta\pi^0(m_1 - \Delta_1) \geq -(1 - \delta)c_1 - \delta\pi^0c_1 \\ \Leftrightarrow \quad & \varepsilon \leq \varepsilon^* \equiv \left(\frac{\delta}{1 - \delta}\right) \pi^0(\Delta_1 + c_1 - m_1) > 0 \end{aligned}$$

The corresponding avoidance constraint for $\alpha_1 = 0$ can be written as the opposite of (40):

$$\delta(V_1^0 - V_1^1) = -\delta\pi^0(m_1 - \Delta_1) + \delta\pi^0(1 - c_2) \leq (1 - \delta)(\hat{d} - m_1) \quad (41)$$

Using the fact that $m_1 \in [1 - c_2, c_1]$, it is straightforward to verify that:

$$\delta(V_1^0 - V_1^1) \leq \delta\pi^0\Delta_1 \quad \wedge \quad (1 - \delta)(\hat{d} - m_1) \geq (1 - \delta)(\hat{d} - c_1)$$

Hence, a sufficient condition in order to satisfy the avoidance constraint (41) is that:

$$\delta\pi^0\Delta_1 \leq (1 - \delta)(\hat{d} - c_1) \quad \Leftrightarrow \quad D_1 \leq \frac{(1 - \delta)(\hat{d} - c_1)}{\delta\pi^0} > 0$$

Which can be satisfied because we have already established that we can assume without loss of generality that Δ_1 is small enough. Hence $\sigma^1 = \sigma^{\varepsilon^*}$ is a PPE. The corresponding payoff for player 1 is:

$$V_2(\sigma^1) = \pi^0 \left[-(1 - \delta)(c_1 + \varepsilon^*) - \delta\pi^0(m_1 - \Delta_1) \right] + (1 - \pi^0) \left[-\delta\pi^0c_1 \right] = -\pi^0c_1$$

Finally, since the total cost of σ^* is less than π^0 , it follows that so is the total cost of σ^1 and therefore:

$$V_2(\sigma^1) > -\pi^0 - V_1(\sigma^*) = -\pi^0 + \pi^0c_1 = -\pi^0(1 - c_1)$$

Hence σ^1 is a PPE satisfying the desired properties.

The proof for σ^2 is almost analogous. The only difference is that the relevant avoidance constraint in the construction can also be satisfied when δ is low enough. This is not a problem because it can be satisfied whenever $\delta \leq \bar{\delta}$ (and in particular when $\delta = \bar{\delta}$ the conditions of the Lemma are also satisfied (by lemmas 3 and 1)). Since the set of equilibrium payoffs is monotone with respect to δ , this implies that whenever $\delta > \bar{\delta}$ there will be some PPE which achieves the same payoff vector. \square

proof of Lemma 3. Let $\bar{\delta}, \underline{\delta} \in (0, 1)$ be the following bounds for the discount factor:

$$\bar{\delta} \equiv \frac{\hat{d} - c_1}{\hat{d} - c_1 + \pi^0(c_1 + c_2 - 1)}, \quad \underline{\delta} \equiv \frac{\hat{d} - c_1}{\hat{d} - c_1 + \pi^0 c_1}. \quad (42)$$

The assumptions that $\hat{d} > c_1$, $c_1 + c_2 > 1$, and $c_2 < 1$ imply that $0 < \underline{\delta} < \bar{\delta} < 1$.

Let σ^0 be the public strategy profile described as follows:

- $\sigma_1^0 = (1, c_1, 1 - c_1)$
- It there is a crisis in period one, then $\sigma_t^0 = (0, c_1, 1 - c_1)$ for $t > 1$.
- It there is no crisis in period one, then $\sigma_t^0 = (0, 1 - c_2, c_2)$ for $t > 1$.

For $t > 1$, σ^0 is just repetition of static Nash equilibria and hence the continuation strategies are PPE. If there is a crisis in period 1, the agent's contributions are static best responses and do not affect the continuation value, hence the mitigation incentive constraints are satisfied. We only need to check the avoidance incentives for agent 1 in period 1. We have $V_1^0 = -\pi^0(1 - c_2)$ and $V_1^1 = -\pi^0 c_1$. And hence $V_1^0 - V_1^1 = \pi^0(c_1 + c_2 - 1)$. Therefore the avoidance constraint (40) can be written as:

$$\begin{aligned} \delta(V_1^0 - V_1^1) \geq (1 - \delta)(\hat{d} - c_1) &\Leftrightarrow \delta\pi^0(c_1 + c_2 - 1) \geq (\hat{d} - c_1) - \delta(\hat{d} - c_1) \\ &\Leftrightarrow \delta[\pi^0(c_1 + c_2 - 1) + (\hat{d} - c_1)] \geq (\hat{d} - c_1) \Leftrightarrow \delta \geq \bar{\delta} \end{aligned}$$

On the other hand, suppose that σ is a PPE and after some history the active agent chooses $\alpha_t(h^t) = 1$. If $\mu_{1t}(h^t) > 0$ then the avoidance and mitigation constraints (39)–(40) require:

$$\left. \begin{aligned} (1 - \delta)m_1 - \delta V_1^1 &\leq (1 - \delta + \delta\pi^0)c_1 \\ (1 - \delta)m_1 - \delta V_1^1 &\geq (1 - \delta)\hat{d} - \delta V_1^0 \end{aligned} \right\} \Rightarrow (1 - \delta)\hat{d} - \delta V_1^0 \leq (1 - \delta + \delta\pi^0)c_1$$

If $\mu_{1t}(h^t) = 0$ then there is no mitigation constraint, but individual rationality requires $-\delta V_1^1 \leq \delta\pi^0 c_1 \leq (1 - \delta + \delta\pi^0)c_1$ and we obtain the same inequality. Furthermore, feasibility implies that $V_1^0 \leq 0$. After some algebra this implies that $\delta \geq \underline{\delta} \in (0, 1)$. Hence if δ is low enough there is no avoidance. \square

C APS decomposition

With agents playing public strategies only, this repeated game has a recursive structure. After an arbitrary history, the continuation strategy profile of a PPE is an equilibrium profile of the original game. The standard way to characterize the set of PPE values is

to use the self-generation procedure introduced in [Abreu et al. \(1990\)](#). This appendix establishes an analogous procedure, and shows that our restriction to pure and public strategies is without loss of generality.

The model here differs from APS in three aspects. The stage game is a multi-stage game. We allow for public randomization but exclude individual mixed strategies. Our game is not purely a game with perfect monitoring (actions are perfectly observable), nor one purely with imperfect public monitoring (agents observe only noisy information of past actions). It is a mixture of the two: at any date t , agent 1's avoidance action a_t is his private information, agent 2 sees only the a_t -induced realization of state ξ_t , yet both agents observe each other's contribution m_{1t} and m_{2t} . Nonetheless, the self-generation procedure can be used to generate the set of PPE values \mathcal{V} .

The expected payoff for player $i = 1, 2$ of action profile $s = (a, m_1, m_2)$ when the continuation values are given by $w: X \times M^2 \rightarrow \mathbb{R}^2$ is¹⁰

$$g_i(s, w) \equiv (1 - \delta)v_i(s) + \delta \sum_{\xi \in X} \Pr(\xi|a) w_i(\xi, m_1(\xi), m_2(\xi)) \quad (43)$$

Definition 3. An action profile $s^* = (a^*, m_1^*, m_2^*)$ together with a continuation payoff function $w: X \times M^2 \rightarrow \mathbb{R}^2$ is *admissible* with respect to a set $W \subset \mathbb{R}^2$ if and only if:

- (i) for all $\xi \in X$, $w(\xi, m_1^*(\xi), m_2^*(\xi)) \in W$.
- (ii) $a^* \in \arg \max_{a \in A} g(a, m_1^*, m_2^*, w)$
- (iii) For any $\xi \in X$:

$$m_1^*(\xi) \in \arg \max_{m_1 \in M} (1 - \delta)u_1(\xi, a^*, m_1, m_2^*(\xi)) + \delta w_1(\xi, m_1', m_2^*(\xi))$$

$$m_2^*(\xi) \in \arg \max_{m_2 \in M} (1 - \delta)u_2(\xi, a^*, m_1^*(\xi), m_2) + \delta w_2(\xi, m_1^*(\xi), m_2')$$

Condition (i) says that the continuation payoff takes value in the W set. Conditions (ii) and (iii) require that the action profile (a, m_1, m_2) is a Nash equilibrium if the two agents' payoff function is given by g_i .

In our setting with public randomization, the adequate self-generating operator is the one in [Cronshaw and Luenberger \(1994\)](#). For any set $W \subseteq \mathbb{R}^2$ and every action

¹⁰ The sunspots do not appear explicitly the sequential formulation. The action profile s specifies the pure actions chosen after θ_t is realized, and the continuation value w is taken to be the average continuation values integrating over θ_{t+1} . Public randomization enters implicitly in the the convex hull operation in (45).

profile $s = (a, m_1, m_2)$ define:

$$B_s(W) = \left\{ v \in W \mid \begin{array}{l} \exists w: X \times M^2 \rightarrow W \text{ such that} \\ (s, w) \text{ is admissible w.r.t. } W \text{ and } v = g(s, w) \end{array} \right\} \quad (44)$$

and:

$$B(W) = \text{co} \left(\bigcup_{s \in S} B_s(W) \right) \quad (45)$$

A set $W \subseteq \mathbb{R}^2$ is said to be self-generating if and only if $W \subseteq B(W)$. The following proposition states that using this notion of self-generation, the characterization of APS applies to our setting.

Proposition 5. *\mathcal{V} is compact and convex and it is the largest-self generating set. Furthermore, if W is bounded and $\mathcal{V} \subseteq W$ then $B^n(W) \rightarrow \mathcal{V}$.*

Lemma 4 (Single deviation principle). *An individually rational strategy profile is a PPE if and only if it admits no profitable single-deviations.*

Proof. The result follows can be proven in a similar way as Proposition 2.2.1 in (Mailath and Samuelson, 2006, pp 25), The general lines of the argument are as follows. Fix an individually rational strategy profile and suppose that there is a profitable deviation and let V be the difference in values. Since the set of individually rational payoffs is bounded, we know that there is some T such that the payoffs after T contribute to less than $\|V\|/2$ and thus there is also a profitable deviation of length at most T . If the deviation in the last period is profitable, then the proof is complete. If not, then there is a profitable deviation of length at most $T - 1$ periods. By induction, this implies that there is a profitable deviation of length 1. \square

Lemma 5 (Self-generation). *If W is bounded and self-generating then $W \subseteq \mathcal{V}$*

Proof. For each $v \in W$, by Caratheodory's theorem, there exist some $b_v^1, b_v^2, b_v^3 \in \cup_s B_s(W)$ and $(\lambda_v^1, \lambda_v^2, \lambda_v^3) \in \Delta^3$ such that $v = \sum_{n=1}^3 \lambda_v^n b_v^n$. For each b_v^n , we can define s_v^n and w_v^n so that $b_v^n = g(s_v^n, w_v^n)$ and (s_v^n, w_v^n) is admissible w.r.t. W . Now fix some $v^* \in W$. We will construct a PPE σ^* such that $v^* = V(\sigma^*)$. For that purpose we define a sequence of (public) history-dependant continuation values $\{w_t^*\}$ with $w_t^*: H^t \rightarrow W$. σ^* is defined as a function of w^* :

$$\sigma_t^*(h^t, \theta_t) = \begin{cases} s_{w_t^*(h^t)}^1 & \text{if } \theta_t \leq \lambda_{w_t^*(h^t)}^1 \\ s_{w_t^*(h^t)}^2 & \text{if } \lambda_{w_t^*(h^t)}^1 < \theta_t \leq \lambda_{w_t^*(h^t)}^2 \\ s_{w_t^*(h^t)}^3 & \text{if } \theta_t > \lambda_{w_t^*(h^t)}^2 \end{cases}$$

w^* is defined recursively with $w_1^* = v^*$ and:

$$w_{t+1}^*(h^{t+1}) = \begin{cases} w_{w_t^*(h^t)}^1(m_t, \xi_t) & \text{if } \theta_1 \leq \lambda_{w_t^*(h^t)}^1 \\ w_{w_t^*(h^t)}^2(m_t, \xi_t) & \text{if } \lambda_{w_t^*(h^t)}^1 < \theta_1 \leq \lambda_{w_t^*(h^t)}^2 \\ w_{w_t^*(h^t)}^3(m_t, \xi_t) & \text{if } \theta_1 > \lambda_{w_t^*(h^t)}^2 \end{cases}$$

It is straightforward to see that w_T^* and thus σ_t^* are measurable. Now notice that for every h^t :

$$\begin{aligned} w_t^*(h^t) &= \sum_{n=1}^3 \lambda_{w_t^*(h^t)}^n g(s_{w_t^*(h^t)}^n, w_{w_t^*(h^t)}^n) = \mathbf{E}_t \left[g(\sigma_t^*(h^t, \theta_t), w_{t+1}^*(h^{t+1})) \right] \\ &= \mathbf{E}_t \left[(1 - \delta)v(\sigma_t^*(h^t, \theta_t)) + \delta w_{t+1}^*(h^{t+1}) \right] \\ &= (1 - \delta)\mathbf{E}_t \left[v(\sigma_t^*) + \delta v(\sigma_{t+1}^*) + \frac{\delta}{1 - \delta} w_{t+2}^*(h^{t+2}) \right] = \dots = V(\sigma^*|_{t^t}) \end{aligned}$$

And hence we have that $v^* = w_1^* = V(\sigma^*)$. Finally, since actions and continuation values are admissible at every period, we know that there are no profitable single-deviations. By Lemma 4 this implies that σ^* is a PPE. \square

Lemma 6 (Factorization). \mathcal{V} is self-generating

Proof. Fix an arbitrary point $v^* \in \mathcal{V}$ and let σ^* be the PPE that generates it. For each $\theta \in [0, 1]$, $m_1, m_2 \in M$ and $\xi \in X$ set $w_\theta(m_1, m_2, \xi) = V(\sigma^*|_{\theta, \xi, m_1(\xi), m_2(\xi)}) \in \mathcal{V}$. Since σ^* is measurable it follows that w_θ is measurable for $\theta \in [0, 1]$. Since σ^* is a PPE we know that for each realization of θ_1 we have that $(\sigma_1^*(\theta_1), w_{\theta_1})$ is admissible w.r.t. \mathcal{V} and thus $g(\sigma_1^*(\theta_1), w_{\theta_1}) \in B_{\sigma_1(\theta_1)}(\mathcal{V})$. This implies that:

$$v^* = \int_0^1 g(\sigma_1^*(\theta_1), w_{\theta_1}) d\theta_1 \in \text{co} \left(\bigcup_s B_s(\mathcal{V}) \right) = B(\mathcal{V})$$

\square

Lemma 7. If W is compact then $B(W)$ is compact

Proof. Fix some $a \in A$. We will start by showing that $B_a(W) \equiv \cup_m B_{a,m}(W)$ is compact. Consider any sequence $\{v^n\}$ in $B_a(W)$ converging to some $v^* \in \mathfrak{R}^2$. By construction there exist sequences $\{m^n\}$ and $\{w^n\}$ such that $v^n = g(a, m^n, w^n)$ and (a, m^n, w^n) is admissible w.r.t. W . Since it is contained in a compact space, the sequence $\{m^n, w^n(m^n)\}$ has a subsequence converging to some limit $m^*, w^*(m^*)$. Since W and M are closed, we know that $m^*: X \rightarrow M$ and $w^*(m^*) \in W$. Let $w_i^*(m, m_{-i}^*(\xi), \xi) = \min_{w \in W} w_i$ for $m \neq m_i^*(\xi)$ and define other values of w^* arbitrarily. Since g is continuous we know that $v^* = g(a, m^*, w^*)$. Since the incentive constraints are defined by continuous functions

we know that and (a, m^*, w^*) is admissible w.r.t. W . Hence $v^* \in B_a(W)$. Since this was for arbitrary convergent sequences, this means that $B_a(W)$ is closed.

Now, since the payoffs of the stage game are all non-positive and W is bounded, then $B_a(W)$ is bounded above. Since admissibility implies that the values have to be conditionally individually rational, it is also bounded below. Hence $B_a(W)$ is compact.

Since a finite union of compact sets is compact, we have that $\cup_s B_s(W) = \cup_a B_a(W)$ is compact. The result then follows from the fact that the convex hulls of compact sets are compact. \square

Proof of proposition 5. Since B is \subseteq -monotone by construction, Lemma 5 implies that \mathcal{V} contains the union of all self-generating sets. By Lemma 6 this implies that \mathcal{V} is the largest self-generating set. Since $B(W)$ is convex for any W by construction, Lemma 6 also implies that \mathcal{V} is convex. Now fix any bounded set W such that $\mathcal{V} \subseteq W$. Let \bar{W} be the closure of W and define the sequence $\{W^n\}_{n=1}^\infty$ by $W^0 = \bar{W}$ and $W^{n+1} = B(W^n)$ for $n = 1, 2, \dots$. By definition of B and Lemma 7 we know that W^n is a \subseteq -decreasing sequence of compact sets and therefore has a (Hausdorff) limit $W^* = \cap_n W^n$ and this limit is compact. Since B is \subseteq -monotone and \mathcal{V} is self-generating, we know that $\mathcal{V} = B^n(\mathcal{V}) \subseteq B^n(W^0) = W^n$ for all n , and thus $\mathcal{V} \subseteq W^*$.

It only remains to show that W^* is self-generating, for this purpose we combine the proofs from Lemmas 5 and 7. Consider any $v^* \in W^*$. By construction we know that $v^* \in B(W)$ for all n . Therefore there exists sequences $\left\{ (b^{nk}, \lambda^{nk}, s^{nk}, w^{nk})_{k=1}^3 \right\}_{n=1}^\infty$ such that $v^* = \sum_{k=1}^3 \lambda^{nk} b^{nk}$, $b^{nk} = g(s^{nk}, w^{nk})$ and (s^{nk}, w^{nk}) is admissible w.r.t. W^n for all n . Since it is contained in a compact space, the sequence $\{b^n, \lambda^n, s^n, w^n(m^n)\}$ has a subsequence converging to some limit $b^*, \lambda^*, m^*, w^*(m^*)$. Since all the sets are closed, we know that the limit is where we want it to be. Let $w^{k*}(m, m_{-i}^{k*}(\xi), \xi) = \min_{w \in W^*} w_i$ for $m \neq m_i^{k*}(\xi)$ and define other values of w^{k*} arbitrarily. Since g is continuous we know that $b^{k*} = g(s^{k*}, w^{k*})$. Since the incentive constraints are defined by continuous functions we know that (s^{k*}, w^{k*}) is admissible w.r.t. W^* . It is straightforward to see that $v^* = \sum_{k=1}^3 \lambda^{k*} b^{k*} \in B(W^*)$. Therefore W^* is self generating and, by Lemma 5, $W^* \subseteq \mathcal{V}$. \square

We are now in a position to prove our claim about the restriction to pure-public strategies being without loss of generality. One could extend the definition of equilibrium in the obvious way to allow for mixed strategies that depend on private information. The following proposition states that the set of equilibrium payoffs would not change. The reason for this is because the new set would be self-generating in the original sense, and thus it would be contained in \mathcal{V} .

Proposition 6. *The set of equilibrium payoffs would remain unchanged if we allowed player 1 to use private strategies and we allowed both players to use mixed strategies.*

Proof of proposition 6. We present a sketch of the proof, a formal proof can be provided upon request. The definitions of PPE and $V(\sigma)$ can be easily extended to allow for mixed strategies in the obvious way. Let $\tilde{\mathcal{V}}$ be the set of payoffs resulting from pure or mixed PPE. Fix some $v^* \in \tilde{\mathcal{V}}$ and let σ^* be the a mixed strategy PPE and that generates it. Now delegate all the randomization to t_1 , define continuation values in the obvious way and show that the resulting pairs $(\sigma|\theta_1, w|\theta_1)$ are admissible w.r.t. the set of mixed PPE payoffs. Intuitively, this occurs because M is convex and thus there is no need to randomize contributions, and m_1 and m_2 are chosen after observing a and thus there is no need to randomize the avoidance action. This implies that the set of mixed PPE payoffs is self-generating and is thus contained in \mathcal{V} . \square